

Bayesian estimation of the low-energy constants up to fourth order in the nucleon-nucleon sector of chiral effective field theory

Isak Svensson,^{*} Andreas Ekström, and Christian Forssén 

Department of Physics, Chalmers University of Technology, SE-412 96 Göteborg, Sweden



(Received 23 June 2022; accepted 19 December 2022; published 20 January 2023)

We use Bayesian methods and Hamiltonian Monte Carlo (HMC) sampling to infer the posterior probability density function (PDF) for the low-energy constants (LECs) up to next-to-next-to-next-to-leading order (N3LO) in a chiral effective field theory (χ EFT) description of the nucleon-nucleon interaction. In a first step, we condition the inference on neutron-proton and proton-proton scattering data and account for uncorrelated χ EFT truncation errors. We demonstrate how to successfully sample the 31-dimensional space of LECs at N3LO using a revised HMC inference protocol. In a second step we extend the analysis by means of importance sampling and an empirical determination of the neutron-neutron scattering length to infer the posterior PDF for the leading charge-dependent contact LEC in the 1S_0 neutron-neutron interaction channel. While doing so we account for the χ EFT truncation error via a conjugate prior. We use the resulting posterior PDF to sample the posterior predictive distributions for the effective range parameters in the 1S_0 wave as well as the strengths of charge-symmetry breaking and charge-independence breaking. We conclude that empirical point-estimate results of isospin breaking in the 1S_0 channel are consistent with the PDFs obtained in our Bayesian analysis and that, when accounting for χ EFT truncation errors, one must go to next-to-next-to-leading order to confidently detect isospin breaking effects.

DOI: [10.1103/PhysRevC.107.014001](https://doi.org/10.1103/PhysRevC.107.014001)

I. INTRODUCTION

In chiral effective field theory (χ EFT) [1–4] the nuclear interaction is parametrized in terms of low-energy constants (LECs) that capture unresolved physics and must be determined from data. The number of LECs grows with the order of the chiral expansion, and the parameter estimation becomes a challenging inference problem that is directly connected with the precision of the theory. Moreover, in χ EFT the long-range pion-nucleon (πN) interaction appears as a subprocess of the nuclear interaction. We can therefore constrain the long-range part of the nuclear interaction rather well using measured πN scattering data [5], albeit less so in the delta-full sector of χ EFT [6]. In a Bayesian context, this knowledge can be straightforwardly accounted for as a prior when learning more about the nuclear interaction from new data, as demonstrated in, e.g., Refs. [7,8]. In fact, Bayesian inference methods allow us to account for any prior beliefs about χ EFT, most importantly its truncation error [9]. This highlights some of the central synergies of χ EFT and Bayesian methods. In this paper, we focus on two challenges tied to a Bayesian analysis

of the nucleon-nucleon (NN) interaction: (i) How to reliably sample the high-dimensional Bayesian posterior probability density functions (PDFs) of the LECs up to next-to-next-to-next-to-leading order (N3LO) in χ EFT, and (ii) how to extend these posterior PDFs, here inferred from neutron-proton (np) and proton-proton (pp) scattering data, to account for the uncertainty in the isospin breaking (IB) and leading neutron-neutron (nn) short-range LEC acting in the 1S_0 partial wave. Throughout the paper we will often use the short-hand labels *posterior* and *prior* to indicate posterior and prior PDFs, respectively.

Sampling a high-dimensional PDF poses a significant challenge in any Bayesian analysis. In a previous paper [8] we explicitly demonstrated the efficiency and manageable dimensional scaling of the Hamiltonian Monte Carlo (HMC) [10] algorithm applied to χ EFT up to next-to-next-to-leading order (NNLO). HMC exploits the geometry of the parameter space and Hamiltonian dynamics to draw virtually independent Markov chain Monte Carlo (MCMC) samples concentrated to the bulk of the probability mass. In this paper, we revise our HMC protocol to sample the LEC posterior at the next chiral order, i.e., N3LO, where we encounter a significantly more challenging inferential problem in 31 dimensions, one per LEC.

Furthermore, the fundamental effect of IB leads to slight differences in the strong interaction between neutrons, protons, and between protons and neutrons. It originates from differences in the masses and electromagnetic charges of the up- and down-quarks [11] and is expected to be weaker in the irreducible interaction between three nucleons [12,13].

^{*}isak.svensson@chalmers.se

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by [Bibsam](https://www.bibsam.se/).

Although weak, the IB of the strong nuclear interaction plays an important role in *ab initio*, and mean-field, analyses of infinite nuclear matter and finite nuclei, in particular towards the drip lines with pronounced proton-to-neutron ratios (see, e.g., Refs. [14–23]). It is therefore important to quantify the uncertainties pertaining to the LECs governing the strengths of the IB effects. Unfortunately, inferences conditioned solely on the world database of NN scattering data [24,25], which does not contain nn cross sections, leaves the posterior of LECs acting (only) in the nn isospin channel unconstrained. To handle this, a point estimate of the leading charge-dependent LEC in the nn channel is typically obtained using the empirical value for the corresponding scattering length, primarily in the 1S_0 partial wave. This latter quantity parametrizes the total nn cross section in the limit of zero scattering energy and its value is estimated from data on hadronic reactions that involve two neutrons in the initial and/or final state [26]. A future prospect is to employ results from realistic lattice quantum chromodynamics to directly infer the values of the LECs [27].

In χ EFT, isospin is an exact symmetry at leading order (LO) [28]. At next-to-leading order (NLO), one introduces realistic charged-to-neutral pion mass splittings in the one-pion exchange potential (OPEP). The $\approx 3\%$ mass splitting of the rather light pions induces IB in the S waves and beyond. We also have charge-dependent and nonderivative LECs in the 1S_0 partial wave. Higher-order IB effects can be accounted for systematically by introducing, e.g., charge-dependent πN LECs, (charged) pion-photon interactions, and charge-dependent LECs in partial waves with nonzero angular momentum. However, many of those IB effects are estimated to be negligible compared to the OPEP mass splitting and nonderivative S -wave LECs [2–4]. Indeed, the leading IB πN LECs were recently [29] inferred using NN scattering data and fifth-order χ EFT, i.e., N4LO, and found to exhibit no significant charge dependence. In fact, higher-order IB effects are often neglected in quantitative chiral interactions [30–34].

In this paper, we employ NN interactions from χ EFT up to N3LO in Weinberg power counting as defined in Ref. [3], and with the IB effects due to pion mass-splitting in the OPEP and charge-dependent leading S -wave contacts. We use non-local regulators in relative momenta p according to $f(p) = \exp(-p^{2n}/\Lambda^{2n})$ with a cutoff $\Lambda = 450$ MeV and $n = 3$. Two-pion exchanges are spectral-function regulated [35,36] with a cutoff of 700 MeV. At N3LO, to complete the subleading two-pion exchange, we include all two-loop diagrams with some of them evaluated using numerical integration. To remedy the on-shell redundancy [32,37] in the S -wave contact potential at N3LO we set the contact fourth-order contact LECs \widehat{D}_{1S_0} , \widehat{D}_{3S_1} , and $\widehat{D}_{3S_1-^3D_1}$ (in the notation of Ref. [3]) to zero. For describing pp and np low-energy scattering data we append the standard electromagnetic interactions up to second order in the fine-structure constant as outlined in, e.g., Ref. [38], at all chiral orders.

Within this χ EFT framework, we perform a Bayesian study of the NN interaction, up to N3LO, conditioned on NN scattering data and subsequently extend the LEC posterior using an empirical value for the 1S_0 nn scattering length a_{exp}^{nn} . This allows us to quantify the uncertainties of the IB effects

due to the short-range LECs in χ EFT. In the process of doing so, we also give an example of the flexibility of the Bayesian framework to straightforwardly expand existing results by introducing new parameters and conditioning on new data. We also test the robustness of a commonly employed model [9,37] for estimating truncation errors in χ EFT.

This paper is organized as follows. In Sec. II, we outline the statistical model upon which we base all inferences and demonstrate how to draw samples from the posterior PDFs and analyze the consistency of our inferences. In Sec. III, we draw samples from the posterior predictive distributions (PPDs) for scattering lengths and effective ranges in the 1S_0 partial wave. We summarize our findings in Sec. IV.

II. STATISTICAL METHOD

In this section we explain our method for inferring the joint posterior PDFs for *all* the LECs in the NN sector of χ EFT up to N3LO. After the specification of the prior and likelihood, our inference is performed in two stages. In a first step, we apply the HMC algorithm from Ref. [8] to quantify the posterior PDFs for the NN LECs in the np and pp sectors conditioned on np and pp scattering data. The posteriors presented in this paper account for an uncorrelated χ EFT error model conditioned on the order-by-order convergence pattern up to N3LO. In a second step, we marginalize-in the nn non-derivative LEC $\widehat{C}_{1S_0}^{nn}$ and condition the inference on a single datum: The nn scattering length in the 1S_0 partial wave, $a_{\text{exp}}^{nn} = -18.9 \pm 0.4$ fm [26,39]. We note that this is one of the currently accepted values for this scattering length and that there are conflicting experimental values for which the experimental uncertainties are not fully understood. However, we do not account for this additional level of uncertainty. See, e.g., Ref. [40] for a summary of the present status on this topic and the proposal of a novel method to measure the nn scattering length. In this paper we only operate with empirical scattering lengths and effective ranges for which electromagnetic (EM) effects have been removed by the originators of those values. Similarly, our theoretical predictions for effective range parameters do not include any EM effects either. We do however include EM effects when we compute scattering cross sections during the HMC sampling.

A. Finding an expression for the joint posterior

Using the notation $\text{pr}(A|B)$ for the conditional probability that proposition A is true given B , the posterior PDF of interest $\text{pr}(\vec{\alpha}|D, I)$ can, according to Bayes' theorem, be evaluated as

$$\text{pr}(\vec{\alpha}|D, I) \propto \text{pr}(D|\vec{\alpha}, I) \times \text{pr}(\vec{\alpha}|I), \quad (1)$$

i.e., as a product of the data likelihood $\text{pr}(D|\vec{\alpha}, I)$ times the prior PDF $\text{pr}(\vec{\alpha}|I)$. Here, $\vec{\alpha}$ denotes the vector of LECs to be inferred, D is employed data, and I encompasses all other assumptions and given information. The given information includes, for example, the chiral order, the regulator cutoff, masses, etc., and we make assumptions about the truncation error, data selection, and so on. Throughout this paper, we omit the overall normalization factor $\text{pr}(D|I)$ as it does not play a central role in parameter estimation.

In a first step we use the np and pp scattering data $D_{np,pp}$ contained in the Granada database [24,25]. In doing so we obtained posteriors for the πN and contact NN LECs, excluding the \tilde{C}_{150}^{nn} LEC. To simplify the notation we often denote this latter, explicitly charge-dependent, nn LEC as α_{nn}

while all other LECs are collectively denoted $\tilde{\alpha}_{np,pp}$. In this notation, the main objective in this paper is to extract the joint posterior $\text{pr}(\tilde{\alpha}|D, I)$ of $\tilde{\alpha} \equiv (\alpha_{nn}, \tilde{\alpha}_{np,pp})$ conditional on $D = (a_{\text{exp}}^{nn}, D_{np,pp})$.

Using the product rule of probabilities we can write

$$\begin{aligned} \text{pr}(\tilde{\alpha}|D, I) &\equiv \text{pr}(\alpha_{nn}, \tilde{\alpha}_{np,pp} | a_{\text{exp}}^{nn}, D_{np,pp}, I) \text{pr}(\alpha_{nn} | \tilde{\alpha}_{np,pp}, a_{\text{exp}}^{nn}, D_{np,pp}, I) \times \text{pr}(\tilde{\alpha}_{np,pp} | a_{\text{exp}}^{nn}, D_{np,pp}, I) \\ &= \text{pr}(\alpha_{nn} | \tilde{\alpha}_{np,pp}, a_{\text{exp}}^{nn}, I) \times \text{pr}(\tilde{\alpha}_{np,pp} | D_{np,pp}, I), \end{aligned} \quad (2)$$

where we have assumed conditional independence in the final equality. Thus far, no strong assumptions have been made regarding the relation between α_{nn} and $\tilde{\alpha}_{np,pp}$, and the analysis is quite general.

We use HMC to sample $\text{pr}(\tilde{\alpha}_{np,pp} | D_{np,pp}, I)$, and the strategy that we employ for this is explained in Ref. [8]. In brief, we employ order-by-order differences up to N3LO, omitting the LO results, to estimate the variance of the (uncorrelated) χ EFT truncation error for describing scattering data. We place a multivariate Gaussian prior on the subleading πN LECs at NNLO and N3LO using the results from a Roy-Steiner analysis of πN scattering amplitudes [6]. Furthermore, we place a rather weak prior on the NN contact LECs at all orders using an uncorrelated Gaussian PDF with zero mean and standard deviation of $5 \times 10^4 \text{ GeV}^{-(k+2)}$ for the LECs belonging to the $k = 0, 2, 4$ (LO, NLO, N3LO) contact Lagrangian as defined in Weinberg power counting. The full prior factorizes into independent NN and πN PDFs since we assume no correlation between these sectors. In Sec. II B 1 we present further details about the sampling and how we revised the sampling protocol to reach N3LO. In the remainder of this section we assume that $\text{pr}(\tilde{\alpha}_{np,pp} | D_{np,pp}, I)$ is known to us.

Next, we extend this posterior by incorporating α_{nn} . Using Bayes' theorem, we rewrite the first factor in the final row of (2) according to

$$\begin{aligned} \text{pr}(\alpha_{nn} | \tilde{\alpha}_{np,pp}, a_{\text{exp}}^{nn}, I) &\propto \text{pr}(a_{\text{exp}}^{nn} | \alpha_{nn}, \tilde{\alpha}_{np,pp}, I) \\ &\times \text{pr}(\alpha_{nn} | I), \end{aligned} \quad (3)$$

where we again used that the prior for α_{nn} is conditionally independent of $\tilde{\alpha}_{np,pp}$. To evaluate the likelihood $\text{pr}(a_{\text{exp}}^{nn} | \alpha_{nn}, \tilde{\alpha}_{np,pp}, I)$ we numerically compute the scattering length at the specified chiral order in χ EFT and given values for $(\alpha_{nn}, \tilde{\alpha}_{np,pp})$. Equation (2) for the sought posterior PDF of the LECs thus becomes

$$\begin{aligned} \text{pr}(\tilde{\alpha}|D, I) &\propto \text{pr}(a_{\text{exp}}^{nn} | \alpha_{nn}, \tilde{\alpha}_{np,pp}, I) \\ &\times \text{pr}(\alpha_{nn} | I) \times \text{pr}(\tilde{\alpha}_{np,pp} | D_{np,pp}, I). \end{aligned} \quad (4)$$

One could certainly argue, using previous knowledge of IB, that we are in the right to place a narrow prior for α_{nn} based on the marginal PDFs for \tilde{C}_{150}^{np} and \tilde{C}_{150}^{pp} . However, to avoid building in such expectations on IB we selected a rather weak, and normally distributed, prior according to

$$\text{pr}(\alpha_{nn} | I) = \mathcal{N}(0, \bar{\alpha}_{nn}^2) \quad (5)$$

of width $\bar{\alpha}_{nn} = 5 \times 10^4 \text{ GeV}^{-2}$.

We develop the likelihood in (4) by relating a theoretically computed value a_{th}^{nn} of the nn scattering length to the experimental result via a stochastic model:

$$a_{\text{exp}}^{nn} = a_{\text{th}}^{nn} + \delta a_{\text{exp}}^{nn} + \delta a_{\text{th}}^{nn}. \quad (6)$$

This relation introduces the experimental error $\delta a_{\text{exp}}^{nn}$ and the χ EFT error $\delta a_{\text{th}}^{nn}$, which we assume is dominated by the truncation of the χ EFT series. Equation (6) implies that other errors—such as numerical errors—are negligible. The assumption that $\delta a_{\text{exp}}^{nn}$ and $\delta a_{\text{th}}^{nn}$ are independent and normally distributed random variables leads to a Gaussian likelihood for the scattering length,

$$\text{pr}(a_{\text{exp}}^{nn} | \alpha_{nn}, \tilde{\alpha}_{np,pp}, I) \propto \exp\left(-\frac{(a_{\text{exp}}^{nn} - a_{\text{th}}^{nn})^2}{2(\sigma_{\text{exp}}^2 + \sigma_{\text{th}}^2)}\right), \quad (7)$$

where σ_{exp}^2 and σ_{th}^2 denote the variances of the experimental and theoretical errors, respectively. We use $\sigma_{\text{exp}} = 0.4 \text{ fm}$ [26,40,41] as the experimental error.

To model the truncation error, we use the procedure from Refs. [9,37,42]. We therefore assume that we can express the order-by-order predictions for a_{th}^{nn} as the sum

$$a_{\text{th}}^{nn} = a_{\text{ref}} \sum_{i=0}^k c_i Q^i, \quad (8)$$

where k is the chiral order, i.e., $k = 0$ is LO and $k = 2, 3, 4, \dots$ corresponds to NLO, NNLO, N3LO, a_{ref} is a dimensionful reference value for the scattering length, c_i are dimensionless χ EFT expansion coefficients, and the χ EFT expansion parameter is assumed to be given by $Q = m_\pi/\Lambda_b$, which is reasonable for a quantity defined in the zero-momentum limit and analyzed in a pionful theory. To simplify notation, we omit explicit reference to the k dependence of a_{th}^{nn} . We employ a breakdown scale $\Lambda_b = 600 \text{ MeV}$ in accordance with the first analysis we performed in [8]. Assuming that the expansion coefficients c_i (between and within chiral orders) are independent and normally distributed yields the following form for the truncation error:

$$\text{pr}(\delta a_{\text{th}}^{nn} | \bar{c}^2, Q, a_{\text{ref}}, I) = \mathcal{N}(0, \sigma_{\text{th}}^2), \quad (9)$$

where \bar{c}^2 is the variance that characterizes the magnitude of the χ EFT expansion coefficients for the scattering length and effective range. One can show [9,37] that

$$\sigma_{\text{th}}^2 = \bar{c}^2 a_{\text{ref}}^2 \frac{Q^{2(k+1)}}{1 - Q^2}. \quad (10)$$

For the analysis of the scattering length, we place a (conjugate) inverse-gamma (\mathcal{IG}) prior \bar{c}^2 , leading to a marginal posterior for \bar{c}^2 which also follows an \mathcal{IG} distribution but with updated values for the parameters [43]. Marginalization of the variance leads to a Student's t distribution [44] for the truncation error (see Appendix A). Thus, we have that the \mathcal{IG} prior for \bar{c}^2 , with initial parameters α_0 and β_0 ,

$$\text{pr}(\bar{c}^2|I) = \mathcal{IG}(\alpha_0, \beta_0) \propto \frac{1}{\bar{c}^{2(\alpha_0+1)}} \exp\left(-\frac{\beta_0}{\bar{c}^2}\right), \quad (11)$$

is updated by a set of observed expansion coefficients \bar{c} according to

$$\text{pr}(\bar{c}^2|\bar{c}, I) = \mathcal{IG}(\alpha, \beta) \propto \frac{1}{\bar{c}^{2(\alpha+1)}} \exp\left(-\frac{\beta}{\bar{c}^2}\right) \quad (12)$$

with

$$\begin{aligned} \alpha &= \alpha_0 + \frac{n_c}{2}, \\ \beta &= \beta_0 + \frac{\bar{c}^2}{2}, \end{aligned} \quad (13)$$

where n_c is the length of the vector \bar{c} . In practice, we are of course rather limited in the amount of data that we have available to infer α_{nn} and learn about the corresponding χ EFT error for a_{th}^{nn} . To that end we here exploit the order-by-order convergence pattern of a_{th}^{pp} and a_{th}^{np} to learn about the χ EFT truncation error in a_{th}^{nn} . As will be further discussed in Sec. II B, we consider the NNLO-N3LO shift of a_{th}^{pp} to be an outlier and therefore omit this from the order-by-order data used to learn about the magnitude of the EFT truncation error. In detail, we have $n_c = 5$ expansion coefficients for informing the marginal posterior for \bar{c}^2 ; these are calculated (see, e.g., Ref. [8]) from the LO-NLO, NLO-NNLO shifts of a_{th}^{np} and a_{th}^{pp} , and the NNLO-N3LO shift of a_{th}^{pp} , all evaluated at the maximum *a posteriori* (MAP) locations for the PDF $\text{pr}(\bar{\alpha}_{np,pp}|D_{np,pp}, I)$; see Sec. II B for details. However, with our priors, we find that the inference does not change dramatically if we include the NNLO-N3LO shift of a_{th}^{pp} . This is described further in Sec. III and Appendix D.

We choose $(\alpha_0, \beta_0) = (1.0, 12.0)$, which yields the prior and resulting marginal posterior shown in Fig. 1. Our prior is predicated on our previous experience of these expansion coefficients, albeit for NN scattering data: We see it as unlikely that $\bar{c}^2 < 4$, but otherwise acknowledge our ignorance of the size of the truncation error. The mode of the prior, given by $\beta_0/(\alpha_0 + 1)$, is located at $\bar{c}^2 = 6$. Exposure to data \bar{c} shifts the bulk of the PDF towards larger truncation errors, with the mode of the marginal posterior falling at $\bar{c}^2 = 11.5$.

To account for the marginal PDF on the χ EFT truncation error, the posterior in (4) is expanded, viz.,

$$\text{pr}(\bar{\alpha}|\bar{c}, D, I) = \int \text{pr}(\bar{\alpha}, \bar{c}^2|\bar{c}, D, I) d\bar{c}^2, \quad (14)$$

where we now explicitly indicate that the posterior is conditional on \bar{c} . In slightly more detail, the joint posterior of $\bar{\alpha}$




FIG. 1. The prior $\text{pr}(\bar{c}^2|I)$ (dashed line) for \bar{c}^2 and the resulting marginal posterior $\text{pr}(\bar{c}^2|\bar{c}, I)$ (solid line). The latter is obtained from allowing the order-by-order predictions of a_{th}^{np} and a_{th}^{pp} to flow through the prior via the observed expansion coefficients \bar{c} . The individual (squared) elements of \bar{c} are shown as vertical lines, with black (orange) representing np (pp) elements. The center dot in the y-axis label $\text{pr}(\bar{c}^2|\cdot)$ acts as a placeholder for “ I ” and “ \bar{c}, I ”.

and \bar{c}^2 ,

$$\begin{aligned} \text{pr}(\bar{\alpha}, \bar{c}^2|\bar{c}, D, I) &\propto \text{pr}(a_{\text{exp}}^{nn}|\alpha_{nn}, \bar{\alpha}_{np,pp}, \bar{c}^2, I) \times \text{pr}(\alpha_{nn}|I) \\ &\times \text{pr}(\bar{\alpha}_{np,pp}|D_{np,pp}, I) \times \text{pr}(\bar{c}^2|\bar{c}, I), \end{aligned} \quad (15)$$

conditioned on scattering data, scattering lengths, and order-by-order information, is the object of interest that we end up evaluating numerically. Concluding this section we list possible extensions to our analysis that we leave for future work: (1) Incorporating a finite correlation length between the χ EFT expansion coefficients as function of energy, (2) allowing for LEC variability, $c_i = c_i(\alpha_{nn})$, (3) modifying the χ EFT expansion parameter Q to a slightly greater value $m_\pi/500$ MeV [7,8] or, better yet, (4) account for the uncertainty in Q by using an accompanying prior that ensures smooth matching to external (soft) momenta $> m_\pi$ in χ EFT.

B. Evaluating posteriors

In this section we expound on our sampling of the posterior for $\bar{\alpha}_{np,pp}$ using HMC and how we combine this posterior with α_{nn} to produce a joint posterior (15) for all LECs at a given order. For clarity, we now make explicit that the inference of $\bar{\alpha}_{np,pp}$ is conditional on a fixed variance $\bar{c}_{np,pp}^2$ of the χ EFT expansion coefficients. A detailed account of how to efficiently sample $\text{pr}(\bar{\alpha}_{np,pp}|D_{np,pp}, \bar{c}_{np,pp}^2, I)$ using HMC was given in Ref. [8]. Here, we will mainly remark on new developments and results. Our procedure for extracting the joint posterior can be considered a two-step process:

- (1) Sample the posterior $\text{pr}(\bar{\alpha}_{np,pp}|D_{np,pp}, \bar{c}_{np,pp}^2, I)$ using HMC.
- (2) Numerically evaluate the joint posterior $\text{pr}(\bar{\alpha}, \bar{c}^2|D, \bar{c}, I)$ in (15) using importance sampling.

1. Step 1: Sampling $\text{pr}(\tilde{\alpha}_{np,pp}|D_{np,pp}, \tilde{c}_{np,pp}^2, I)$ using HMC

We follow the steps laid out in Ref. [8], with two modifications. First, we consistently use all available chiral orders, i.e., up to N3LO, to learn about the variance of the truncation error. Second, we employ a new method for rapidly tuning the parameters of the HMC algorithm such that we achieve sufficiently high sampling efficiency to perform reliable sampling at N3LO. This yields an N3LO posterior that passes all imposed MCMC convergence tests.

Let us first focus on how we learn about the truncation error. In our previous work [8] we limited ourselves to only use information about the χ EFT convergence pattern up to the order at which we were sampling to estimate the variance \tilde{c}^2 of the truncation error, and we included the zeroth-order coefficients \tilde{c}_0 in the estimation. We then concluded that this procedure typically leads to an underestimation of the size of the truncation error. Here, we follow our own advice and infer the variance \tilde{c}^2 from all orders available to us, and exclude the rather uninformative and biasing zeroth-order coefficients. We use the same grid of observables, laboratory energies, and scattering angles, as in our previous work and arrive at $\tilde{c}_{np,pp} = 4.1$. This value is somewhat greater than the nominally expected natural scale of $\tilde{c} \approx 1$, yet nothing too alarming and certainly in line with what we have observed before. As discussed in Sec. II A, this is also the basis for us shifting the prior for the χ EFT truncation error for the effective range towards slightly greater values.

The importance of tuning the HMC parameters, along with various strategies for achieving efficient sampling, is detailed in Ref. [8]. In particular, we stress the importance of a suitable so-called mass matrix: a matrix that accounts for differences in scale between parameters in the sampled PDF. Previously we extracted a performant mass matrix by executing a short preliminary sampling and inverting the covariance matrix of the samples. Here, we instead optimize the LECs with respect to the posterior, i.e., we solve for

$$\tilde{\alpha}_* = \text{argmax}_{\tilde{\alpha}_{np,pp}} \text{pr}(\tilde{\alpha}_{np,pp}|D_{np,pp}, \tilde{c}_{np,pp}^2, I), \quad (16)$$

and estimate the parameter covariance matrix at that optimum. We use the BFGS optimization algorithm [45–48] which is a so-called quasi-Newton method that relies on first-order gradients to update approximations to the Hessian employed in the Newton algorithm. We employ automatic differentiation (AD) [38,49,50] to obtain the first-order gradients. In the end, we find that the approximate Hessian at the optimum, as found by the BFGS algorithm, makes for a suitable mass matrix in our application. In our case we have access to higher-order gradients via AD and could, in principle, obtain $\tilde{\alpha}_*$ using any expedient optimization method and compute an exact-to-machine-precision Hessian. However, we find that BFGS is sufficient for tuning the HMC sampler. Bypassing the need for the $\approx 20\,000$ function evaluations required during a preliminary HMC sampling is certainly a great improvement, as the BFGS algorithm typically terminates after at most a few hundred function evaluations. Better yet, the optimization-based method is in our experience more reliable and ultimately yields slightly higher performance in the subsequent sampling of the posterior.

TABLE I. MAP predictions of np and pp scattering lengths (a) and effective ranges (r) in units of fm. Empirical results are from Ref. [3].

Order	a_{np}	a_{pp}	r_{np}	r_{pp}
LO	−24.428	−24.056	1.766	1.768
NLO	−21.625	−18.476	2.497	2.585
NNLO	−23.732	−17.797	2.681	2.814
N3LO	−23.733	−16.544	2.693	2.867
Empirical	−23.740(20)	−17.3(4)	2.77(5)	2.85(4)

We gather $N \approx 10\,000$ samples in each of the ten chains we run at each chiral order. We find integrated autocorrelation times τ ($\tau_{\text{NLO}} = 0.71$, $\tau_{\text{NNLO}} = 0.93$, and $\tau_{\text{N3LO}} = 3.33$) that reach a plateau within very few samples, thus indicating converged HMC chains. The τ value should not be too great since it is inversely proportional to the effective sample size that measures the number of independent samples in the HMC chain. We note that the NLO and NNLO autocorrelation times are similar to our previous results [8], achieved with the more cumbersome method of preliminary samplings. We should reiterate the concerns that always surround the topic of convergence in MCMC. In real-world applications it is not possible for us to declare a chain “converged”. This is because we can never explore the entire parameter space in finite time, and we can miss non-negligible (or even dominating) probability regions. All we can do is to probe the MCMC chains for signs of nonconvergence. This issue becomes progressively more pressing as the dimensionality and domain of the parameter space grows, and is especially concerning at N3LO in our case. With that said, we have searched the parameter domains to the best of our abilities by initializing the BFGS optimization, and the HMC algorithm, at multiple locations and we did not find any signs of multimodality and nonconvergence.

In Table I we predict the scattering lengths and effective ranges for the np and pp isospin channels in the 1S_0 partial wave at the MAP points (16) of the LEC posteriors at different orders. Overall, the predictions are reasonably close to the experimental, or rather the empirical, values [3]. This builds confidence in our model inference. These point estimates give a first indication of our model’s performance in the low-energy limit. Overall, we observe convergence towards empirical results as we move to higher chiral orders. A notable exception is the N3LO prediction of the scattering length in the pp isospin channel. It differs from the empirical result by 1.9 standard deviations, compared to just 1.2 standard deviations for the NNLO prediction. Furthermore, the difference is in the opposite direction. This is possibly caused by a lack of informative data during the inference of the LECs.

2. Step 2: Evaluating the joint posterior

Equipped with an MCMC chain to represent $\text{pr}(\tilde{\alpha}_{np,pp}|D_{np,pp}, \tilde{c}_{np,pp}^2, I)$, we proceed to evaluate the posterior $\text{pr}(\tilde{\alpha}, \tilde{c}^2|D, \tilde{c}, I)$ in Eq. (15). Due to the expected weakness of IB effects in our χ EFT we can be fairly certain that the posterior probability mass of the latter PDF will not

TABLE II. MAP predictions of few-nucleon ground-state energies E (in MeV) and point-proton radii R (in fm) for nuclei with mass number $A = 2, 3, 4$ obtained using the Jacobi no-core shell-model [53] in a harmonic oscillator basis with frequency $22 \text{ MeV}/\hbar$ in model spaces with 251, 41, and 21 oscillator shells, respectively. Note that these predictions are not including a three-nucleon (NNN) interaction. For ${}^2\text{H}$ we also list the D -state probability (P_D) in % and the electric quadrupole moment (Q) in units of eb . The experimental and empirical values are the same as in Ref. [38].

Chiral order	$E({}^2\text{H})$	$R({}^2\text{H})$	$P_D({}^2\text{H})$	$Q({}^2\text{H})$	$E({}^3\text{H})$	$R({}^3\text{H})$	$E({}^3\text{He})$	$R({}^3\text{He})$	$E({}^4\text{He})$	$R({}^4\text{He})$
LO	-0.855	2.854	5.347	0.411	-7.214	1.510	-6.327	1.598	-28.781	1.201
NLO	-1.810	2.105	3.052	0.282	-8.311	1.544	-7.559	1.700	-30.233	1.342
NNLO (NN only)	-2.165	1.982	3.212	0.267	-8.570	1.565	-7.791	1.745	-29.873	1.376
N3LO (NN only)	-2.268	1.978	3.476	0.278	-7.560	1.715	-6.829	1.918	-24.272	1.570
Experiment and empirical	-2.225	1.976(1)		0.270(11)	-8.428	1.587(41)	-7.718	1.766(5)	-28.30	1.455(6)

be very far from the one of the former. We take advantage of this expectation and proceed by using the principles of sampling/importance sampling [51], where we first define a sampling distribution

$$g(\vec{\alpha}, \vec{c}^2) \equiv \text{pr}(\vec{\alpha}_{np,pp} | D_{np,pp}, \vec{c}_{np,pp}^2, I) \times \text{pr}(\vec{c}^2 | \vec{c}, I) \times \pi(\alpha_{nn}), \quad (17)$$

with $\pi(\alpha_{nn})$ a simple, bounded uniform distribution. Since the likelihood for the single nn scattering length poses no significant computational challenge, for each chiral order we first explore the α_{nn} space to identify an interval of values encompassing the bulk of the probability mass. These intervals are (in units of 10^4 GeV^{-2}) $\tilde{C}_{150}^{nn} \in [-0.1495, -0.1445]$ at NLO, $\tilde{C}_{150}^{nn} \in [-0.1535, -0.1510]$ at NNLO, and $\tilde{C}_{150}^{nn} \in [-0.132, -0.125]$ at N3LO. We then sample the full posterior (15) using the following procedure:

- (1) Pick a sample $(\vec{\alpha}_i, \vec{c}_i^2)$ from the sampling distribution (17). In practice, we go through the entire MCMC chain of $\vec{\alpha}_{np,pp}$ sequentially since we know [8] that these are independent and random samples from $\text{pr}(\vec{\alpha}_{np,pp} | D_{np,pp}, \vec{c}_{np,pp}^2, I)$. For each $\vec{\alpha}_{np,pp}$ sample we draw a sample of \vec{c}^2 from its marginal posterior $\text{pr}(\vec{c}^2 | \vec{c}, I)$, see Eq. (12), and a sample of α_{nn} from $\pi(\alpha_{nn})$.
- (2) Evaluate the ratio $\omega_i = \text{pr}(\vec{\alpha}_i, \vec{c}_i^2 | D, \vec{c}, I) / g(\vec{\alpha}_i, \vec{c}_i^2)$. Due to the form of our sampling distribution (17) and the factorization of the joint posterior (15), this ratio is simply the product of the Gaussian likelihood for the scattering length (7) and the prior for α_{nn} (5).

Once we have the lists of samples $(\vec{\alpha}_i, \vec{c}_i^2)$ we compute the normalized (importance) weights $q_i = \omega_i / \sum_j \omega_j$. This provides us with a weighted chain of $(\vec{\alpha}, \vec{c}^2)$ values distributed as the target posterior $\text{pr}(\vec{\alpha}, \vec{c}^2 | D, \vec{c}, I)$ (15). The method described here is simple to apply but may not work well in all cases, e.g., if the marginal posterior of α_{nn} is relatively unknown and/or defined for a high-dimensional parameter domain. We have validated our results obtained using importance sampling by approximating $\text{pr}(\vec{\alpha}_{np,pp} | D_{np,pp}, \vec{c}_{np,pp}^2, I)$ with a multivariate normal distribution and directly sampling the full distribution (15) using HMC. Multivariate normal approximations to all posteriors are provided in the Supplemental Material [52].

We report bivariate and univariate marginals of the joint LEC posteriors in Appendix B. In Table II we make point-estimate predictions for selected few-nucleon ground-state observables using the MAP point of the joint posteriors. This is a first check of the model inference, and the results are reasonable. We note that the N3LO predictions deviate more from experiment than NNLO. However, we do not incorporate NNN interactions in this analysis and for that reason we also refrain from estimating χ EFT truncation errors for the results presented in Table II. On the other hand, we perform a full analysis of the NN effective range expansion in Sec. III.

Apart from the inclusion of \tilde{C}_{150}^{nn} and \vec{c} , the NLO and NNLO posteriors are overall similar to their counterparts in Ref. [8]. The differences that do occur arise prior to the inclusion of \tilde{C}_{150}^{nn} as a result of different characterizations of the χ EFT truncation error. The correlation patterns are nearly identical, and the majority of the marginal posteriors overlap at the 68% level. However, some parameters in higher partial waves have somewhat shifted values. For example, in Ref. [8] we reported $C_{3p2} = -0.162(1) \times 10^4 \text{ GeV}^{-4}$ at NLO, while Fig. 5 reveals that $C_{3p2} = -0.175(2) \times 10^4 \text{ GeV}^{-4}$.

The posteriors presented here are confined to rather small volumes in parameter space and, as expected, the marginal posteriors for the \tilde{C}_{150} LEC acting in the different isospin channels are similar to each other. The marginal posterior for \tilde{C}_{150}^{nn} is roughly twice as wide as those for \tilde{C}_{150}^{np} and \tilde{C}_{150}^{pp} . This can largely be attributed to the vastly more abundant np and pp scattering data. The posteriors pick up correlations between \tilde{C}_{150}^{nn} and \tilde{C}_{150}^{np} , \tilde{C}_{150}^{pp} . At N3LO, \tilde{C}_{150}^{nn} shows stronger correlations with other LECs than at NLO and NNLO. The fourth-order contact LECs, D , are sensitive to high momentum observables and therefore less constrained by the data. We find that several of the D LECs return the prior if we do not condition the posterior on data in the $T_{\text{lab}} \in [40.5, 290] \text{ MeV}$ region. Furthermore, the posteriors indicate that some D LECs are notably unnatural; for instance, we have $D_{150} = -29.1(3) \times 10^4 \text{ GeV}^{-6}$, which is of particular interest in this case as it acts in the same partial wave as—and is strongly correlated with—our explicitly isospin breaking LECs. The inference of this LEC, then, is influential on our predictions of 1S_0 scattering lengths and may explain the stark difference between the N3LO predictions for a_{np} and a_{pp} in Table I. There are precedents for large values of D_{150} ; see, e.g., Ref. [54]. Unlike earlier works, we also find rather unnaturally sized




FIG. 2. Prior and posterior PDFs for the πN LECs c_1 , c_3 , c_4 , in units of GeV^{-1} . The posteriors were obtained using HMC. The inner (outer) ellipses enclose approximately 39% (86%) of the probability mass. Black lines indicate the prior as provided by the Roy-Steiner analysis [6]. Blue lines indicate the posterior conditioned on low-energy scattering data only, i.e., with T_{lab} up to 40.5 MeV. Red lines indicate the posterior conditioned on scattering data up to $T_{\text{lab}} = 290$ MeV.

MAP values for $D_{3S1} = -32.2(3) \times 10^4 \text{ GeV}^{-6}$ and $D_{3P1} = -20.4(1) \times 10^4 \text{ GeV}^{-6}$, where the uncertainties denote 68% credible intervals. At N3LO, the various \tilde{C} and C parameters are akin to their values inferred at NNLO, with the exception of C_{3P1} which is both several times larger and has the opposite sign, i.e., $C_{3P1} = -0.956(5) \times 10^4 \text{ GeV}^{-4}$ at NNLO and $C_{3P1} = 6.13(1) \times 10^4 \text{ GeV}^{-4}$ at N3LO.

We have also investigated how the posteriors are affected if we do not exclude the unexpectedly large NNLO-N3LO shift in a_{th}^{pp} from the estimation of the truncation error (see Sec. II A). We find that the width of the marginal \tilde{C}_{1S0}^{nm} posterior roughly doubles at NLO, widens by a small but noticeable amount at NNLO, and is virtually unaffected at N3LO. From this we draw the conclusion that the theoretical uncertainty dominates at NLO, the theoretical and experimental uncertainties are roughly equal at NNLO, and the experimental uncertainty dominates at N3LO.

We find nonoverlapping priors and posteriors for the πN LECs at both NNLO and N3LO. This tension was previously seen at NNLO in Ref. [8]. It turns out that if the posteriors are conditioned on only a low-energy data set, $T_{\text{lab}} \in [0, 40.5]$ MeV, we return our priors from a Roy-Steiner analysis. This is summarized for the NNLO case in Fig. 2. Although there are significant deviations between the πN prior and posterior when we condition on high-energy data, the discrepancies are rather small on a naturalness scale. As such, the statistical model we have set up to relate low-energy data and χ EFT at NNLO appears to preserve the long-range physics rather well. Moreover, the credible intervals of the πN posterior cannot be straightforwardly compared with the variance of the prior determined in a Roy-Steiner analysis [5], for which it is difficult to estimate the truncation uncertainty. At N3LO we

similarly find that conditioning on the energy-truncated data set returns the prior for the πN LECs. The prior and posterior at N3LO (conditioned on the full data set) are shown side by side in Appendix B. It is peculiar that the marginal and univariate posterior for the πN LEC c_4 sits on top of the prior.

III. POSTERIOR PREDICTIVE DISTRIBUTIONS OF 1S_0 ISOSPIN BREAKING

Equipped with samples from the joint posterior PDFs of the NN LECs up to N3LO we proceed to sampling the posterior predictive distributions (PPDs) for the S -wave scattering length a and effective range r defined by the effective range expansion (ERE)

$$p \cot \delta(p) = -\frac{1}{a} + \frac{1}{2} r p^2 + \mathcal{O}(p^4). \quad (18)$$

Here, p is the relative momentum and $\delta(p)$ is the scattering phase shift. The PPD is the PDF of unobserved values for a and r conditioned on np and pp scattering data as well as the empirical nm scattering length.

We quantify the PPD in all three isospin channels $t_z = pp, np, nn$ of the 1S_0 partial wave. The full PPDs are given by the finite set of samples

$$\{a_{\text{th}}^{t_z}(\vec{\alpha}) + \delta a_{\text{th}}\} \quad (19)$$

and

$$\{r_{\text{th}}^{t_z}(\vec{\alpha}) + \delta r_{\text{th}}\}, \quad (20)$$

with $(\vec{\alpha}, \vec{c}^2) \sim \text{pr}(\vec{\alpha}, \vec{c}^2 | D, \vec{c}, I)$. Furthermore, the conjugacy of the \mathcal{IG} prior for the χ EFT truncation error variance enables a pointwise and closed-form evaluation of the associated χ EFT truncation errors δa_{th} and δr_{th} ; see Ref. [43] and Appendix A. As a result of our sampling procedure (see Sec. II B) the sets in Eqs. (19) and (20) are composed of approximately 10^5 weighted samples. In the following we present results based on an inference of \vec{c} where we excluded the NNLO-N3LO shift of the a_{th}^{pp} (see Table I). It turns out that this does not impact any predictions beyond NLO; see Appendix D for comparison.

The results at NLO, NNLO, and N3LO are shown in Fig. 3. Clearly, our model is not fine-tuned to reproduce any of the ERE parameters in Fig. 3 except a_{nn} , which therefore serves as another model check, but all PPDs agree with the empirical results within their uncertainties and the different distributions at different orders overlap reasonably with each other. The exception here is the N3LO prediction of a_{pp} , that we also omitted from the estimation of the χ EFT truncation error. In contrast, for a_{np} , both NNLO and N3LO agree perfectly with the rather precise empirical value. For the empirical values of the pp ERE parameters, the error bands emerge predominantly from the model dependence in the analysis of scattering data and removal of electromagnetic effects. The PPDs for the effective ranges show a great deal of congruity. This implies that the NN contact C_{1S_0} that enters at NLO with a quadratic momentum dependence is sensibly inferred. This contact is missing at LO and the predictions of effective ranges at that order are thus rather poor, as seen in Table I. Overall, one must




FIG. 3. PPDs of scattering lengths and effective ranges at NLO (blue), NNLO (purple), and N3LO (red). Empirical results are shown as black lines, with corresponding 1σ (2σ) uncertainties as a dark (light) gray area. The empirical results are without electromagnetic effects and gathered from Ref. [3], except for a_{nn} which we take from Ref. [26]. Note that the histograms have been scaled such that they have the same peak height.

go beyond NNLO to achieve predictions that are more precise than the corresponding empirical uncertainties.

To quantify the strength of IB effects, we convert the PPDs for the ERE parameters in the 1S_0 partial wave to standard measures of the CIB and charge symmetry breaking (CSB), where the latter amounts to a π rotation around the y axis in isospace, i.e.,

$$\begin{aligned} \Delta a_{\text{CIB}} &= \frac{1}{2}(a_{\text{th}}^{pp} + a_{\text{th}}^{nn}) - a_{\text{th}}^{np}, & \Delta a_{\text{CSB}} &= a_{\text{th}}^{pp} - a_{\text{th}}^{nn}, \\ \Delta r_{\text{CIB}} &= \frac{1}{2}(r_{\text{th}}^{pp} + r_{\text{th}}^{nn}) - r_{\text{th}}^{np}, & \Delta r_{\text{CSB}} &= r_{\text{th}}^{pp} - r_{\text{th}}^{nn}. \end{aligned} \quad (21)$$

The PPDs for the CIB and CSB effects, including χ EFT errors, are summarized in Fig. 4. We first note that it is only at NNLO and beyond that we can detect, with confidence, the overall magnitudes of IB effects in the 1S_0 partial wave. Only CIB in the scattering length can be said to be different from zero with any confidence at NLO. We also find that our PPDs for CSB and CIB agree, within uncertainties, with existing empirical data and a range of point estimates using well-known chiral potentials at NNLO and N3LO. These point estimates are all rather close to each other and fall




FIG. 4. CSB and CIB in the 1S_0 scattering lengths and effective ranges; see (21). Colors and shadings as in Fig. 3. Empirical results (a_{nn} from Ref. [26]; all others from Ref. [3]) are shown in black with gray error bands. Calculations using three widely used potentials are also shown: NNLO_{sat} [31], Idaho-N3LO[54], and Δ -NNLO_{GO}(394) [55]. Note that the three latter potentials agree almost exactly regarding $\Delta r_{\text{CSB,CIB}}$ and are thus difficult to visually distinguish. Also note that the histograms have been scaled such that they have the same peak height.

within the empirical uncertainties at all orders. The CSB and CIB in the effective range is somewhat underestimated and overestimated, respectively, compared to the empirical values. The outlier N3LO result for a_{pp} of course propagates to the results for CIB and CSB and induces a comparatively large isospin breaking at this order. Yet, the results are in line with other chiral potentials.

IV. SUMMARY AND OUTLOOK

In this paper, we sampled high-dimensional posteriors up to N3LO in χ EFT and studied the effects of IB in the NN sector. We used Bayesian inference, conditioned on np and pp scattering data as well as the empirical value a_{exp}^{nm} for the 1S_0 scattering length in the nm channel, to infer posterior PDFs for the LECs at LO, NLO, NNLO, and N3LO. We split the inference in two steps. First we employed HMC and conditioned the LEC posteriors on np and pp scattering data and accounted for uncorrelated χ EFT truncation errors. A new approach to tuning the mass matrix, based on posterior optimization, enabled us to extract a converged 31-dimensional posterior also at N3LO. We find such advancements pivotal for enabling robust MCMC sampling of the N3LO posterior. In the next step of the inference, we included a_{exp}^{nm} and employed importance sampling to marginalize in the nm contact LEC. For the variance of the truncation error in the ERE parameters in the 1S_0 partial wave we employed a conjugate \mathcal{IG} prior. In the end, we find that the resulting LEC posteriors at NLO and NNLO match our previous [8] posteriors in the np and pp isospin sectors and that the LEC $C_{1S_0}^{nm}$ is roughly twice as broad but exhibits a correlation pattern similar to $\tilde{C}_{1S_0}^{np}$ and $\tilde{C}_{1S_0}^{pp}$.

Next, we sampled the PPDs for the ERE parameters and found that our results are consistent with existing point estimates using well-known chiral potentials, with the exception of a_{pp} at N3LO. This outlier might be traced back to insufficient information in the NN scattering data set we conditioned the first part of the inference on. When accounting for χ EFT truncation errors, we find that one must go to NNLO to confidently detect IB effects.

One way to improve a Bayesian analysis of the strong interaction in the low-energy region of relevance to the ERE might be to mix pionless EFT [4,56] and χ EFT. Indeed, χ EFT harbors an intrinsic uncertainty due to the low-energy scale set by m_π . Since the exact form of the χ EFT expansion parameter Q for scattering amplitudes is obscured by the nonperturbative resummation in the Lippmann-Schwinger equation its uncertainty $\text{pr}(Q|I)$ should be accounted for. It could then be interesting to apply Bayesian mixture models to combine pionless EFT and χ EFT predictions.

When we condition the inference on low-energy NN scattering data with $T_{\text{lab}} \in [0, 40.5]$ MeV we return the Roy-Steiner priors for the πN LECs at NNLO and N3LO. This is in accordance with χ EFT being a low-energy theory with long-ranged physics governed by the πN interaction. When conditioning the inference on all NN data up to the pion-production threshold at 290 MeV, the marginal πN LEC posteriors are significantly shifted with respect to the priors. Still, the discrepancies are small on a naturalness scale.

Ab initio studies [55,57] indicate that the inclusion of the Δ isobar yields more realistic predictions for bulk properties of nuclei and nuclear matter. However, a Roy-Steiner prior for the πN LECs with explicit Δ isobars is less precise [6]. As such, it will also become more important to employ additional nuclear structure and reaction data in the inference and employ error models that account for correlations between the χ EFT expansion coefficients. Although it rapidly becomes computationally challenging to evaluate likelihoods encompassing nuclear data for increasing mass numbers, the development of fast and accurate emulators [58–61] appears to provide sufficient leverage. In particular, the present work serves as an example how to sequentially incorporate multiple sets of low-energy nuclear structure and reaction data in a Bayesian framework to perform inference and quantify the theoretical precision.

ACKNOWLEDGMENTS

This work was supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (Grant Agreement No. 758027), the Swedish Research Council (Grants No. 2017-04234 and No. 2021-04507). The computations were enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) partially funded by the Swedish Research Council through Grant Agreement No. 2018-05973

APPENDIX A: THE APPEARANCE OF THE STUDENT'S t DISTRIBUTION

We assume that the n_c expansion coefficients $\vec{c} = (c_1, \dots, c_{n_c})$ in (8) are independent and identically distributed, and drawn from a normal distribution with variance \bar{c}^2 :

$$\text{pr}(c_i|\bar{c}^2) = \mathcal{N}(0, \bar{c}^2). \quad (\text{A1})$$

We then place an inverse-gamma (\mathcal{IG}) prior on \bar{c}^2 , and consequently get an (\mathcal{IG}) posterior, (12), as well due to conjugacy with respect to the normal distribution:

$$\text{pr}(\bar{c}^2|\vec{c}) = \mathcal{IG}(\alpha, \beta), \quad (\text{A2})$$

where the parameters α (shape) and β (scale) have been updated, via (13), by the exposure to \vec{c} . We will now show that these assumptions result in a PPD for an unseen coefficient c_+ given by a Student's t distribution, i.e.,

$$\text{pr}(c_+|\vec{c}) = t_\nu(0, \tau^2), \quad (\text{A3})$$

where ν and τ are the degrees of freedom and scale, respectively, of the Student's t distribution using a method of derivation that can be applied in many other situations where conjugate priors are employed.

In general, the PPD $\text{pr}(c_+|\vec{c})$ for a new datum c_+ given a parameter (or parameters) \bar{c}^2 and observed data \vec{c} can be expressed as

$$\text{pr}(c_+|\vec{c}) = \frac{\text{pr}(c_+|\bar{c}^2, \vec{c}) \times \text{pr}(\bar{c}^2|\vec{c})}{\text{pr}(\bar{c}^2|c_+, \vec{c})}, \quad (\text{A4})$$

which is easily verified using the product rule of probabilities. We know that $\text{pr}(\bar{c}^2|\vec{c}) = \mathcal{IG}(\alpha, \beta)$, and $\text{pr}(\bar{c}^2|c_+, \vec{c})$ is obvi-




FIG. 5. The joint posterior $\text{pr}(\bar{\alpha}, \bar{c}^2 | \text{NLO}, \bar{c}, D)$ conditioned on NN scattering data with $T_{\text{lab}} \in [0, 290]$ MeV, the 1S_0 nm scattering length, and a conjugate prior for the χ EFT truncation error of this quantity. The \tilde{C} LEC values are given in units of 10^4 GeV^{-2} , and C in units of 10^4 GeV^{-4} .

ously also an inverse-gamma distribution with further updated parameters (α_+, β_+)

$$\begin{aligned} \alpha_+ &= \alpha + \frac{1}{2} = \alpha_0 + \frac{n_c}{2} + \frac{1}{2}, \\ \beta_+ &= \beta + \frac{c_+^2}{2} = \beta_0 + \frac{\bar{c}^2}{2} + \frac{c_+^2}{2}. \end{aligned} \quad (\text{A5})$$

Since \bar{c}^2 is given in $\text{pr}(c_+ | \bar{c}^2, \bar{c})$, this is just a normal distribution as in (A1). Hence

$$\text{pr}(c_+ | \bar{c}) = \frac{\mathcal{N}(0, \bar{c}^2) \times \text{IG}(\alpha, \beta)}{\text{IG}(\alpha_+, \beta_+)} \quad (\text{A6})$$




FIG. 6. The joint posterior $\text{pr}(\bar{\alpha}, \bar{c}^2 | \text{NNLO}, \bar{c}, D)$ conditioned on NN scattering data with $T_{\text{lab}} \in [0, 290]$ MeV, the 1S_0 nn scattering length, and a conjugate prior for the χ EFT truncation error of this quantity. The \tilde{C} LEC values are given in units of 10^4 GeV^{-2} , C in units of 10^4 GeV^{-4} , and c in units of GeV^{-1} .

and we obtain for the PPD

$$\begin{aligned} \text{pr}(c_+ | \bar{c}) &= \frac{1}{\sqrt{2\pi} \bar{c}^2} \exp\left(-\frac{c_+^2}{2\bar{c}^2}\right) \\ &\times \frac{\beta^\alpha}{\Gamma(\alpha)} \times \bar{c}^{2(-\alpha-1)} \exp\left(-\frac{\beta}{\bar{c}^2}\right) \\ &\times \frac{(\beta_+)^{\alpha_+}}{\Gamma(\alpha_+)} \times \bar{c}^{2(-\alpha_+-1)} \exp\left(-\frac{\beta_+}{\bar{c}^2}\right). \end{aligned} \quad (\text{A7})$$

This is can be written as

$$\text{pr}(c_+ | \bar{c}) = \frac{1}{\sqrt{2\pi}} \frac{\beta^\alpha \Gamma(\alpha_+)}{\beta_+^{\alpha_+} \Gamma(\alpha)}. \quad (\text{A8})$$

Using (A5) we can rewrite this further as

$$\text{pr}(c_+ | \bar{c}) = \frac{1}{\sqrt{2\pi} \beta} \frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)} \times \left(1 + \frac{c_+^2}{2\beta}\right)^{-\alpha - \frac{1}{2}}. \quad (\text{A9})$$




FIG. 7. A subset of the joint posterior $\text{pr}(\vec{\alpha}, \vec{c}^2 | \text{N3LO}, \vec{c}, D)$ conditioned on NN scattering data with $T_{\text{lab}} \in [0, 290]$ MeV, the 1S_0 nn scattering length, and a conjugate prior for the χ EFT truncation error of this quantity. The \tilde{C} LEC values are given in units of 10^4 GeV^{-2} , C in units of 10^4 GeV^{-4} , and D in units of 10^4 GeV^{-6} .

This is the PDF for a Student's t distribution with $\nu = 2\alpha$ degrees of freedom and scale $\tau = \sqrt{\beta/\alpha}$, i.e.,

$$\text{pr}(c_+ | \vec{c}) = t_\nu(0, \tau^2) = t_{2\alpha}\left(0, \frac{\beta}{\alpha}\right). \quad (\text{A10})$$

One can analogously show that the truncation error is given by

$$\text{pr}(\delta a_{\text{th}} | \vec{a}) = t_{2\alpha}\left(0, a_{\text{ref}}^2 \frac{Q^{2(k+1)} \beta}{1 - Q^2 \alpha}\right), \quad (\text{A11})$$

where \vec{a} is a set of observations of a that can be transformed into expansion coefficients \vec{c} , and α and β are given by (13).

APPENDIX B: LEC POSTERiors

Figures 5 and 6 show corner plots, i.e., marginal uni- and bi-variate PDFs, for the LEC posterior at NLO and NNLO, respectively. A subset of the marginal posteriors at N3LO is shown in Fig. 7. The corner plot of the 31-dimensional N3LO posterior is too large to print and we thus provide it as Supplemental Material to this paper [52]. In Fig. 8 we show a comparison of the πN prior and posterior at N3LO.

For convenience we list the LECs inferred at each order here, sorted as they appear in the corner plots. The inferred LECs at NLO are

$$\vec{\alpha} = (\tilde{C}_{150}^{nn}, \tilde{C}_{150}^{np}, \tilde{C}_{150}^{pp}, \tilde{C}_{3S1}, C_{150}, C_{3P0}, \\ \times C_{1P1}, C_{3P1}, C_{3S1}, C_{3S1-3D1}, C_{3P2}). \quad (\text{B1})$$




FIG. 8. Prior and posterior PDFs for the c and d πN LECs, in units of GeV^{-1} (c) and GeV^{-2} (d). The posterior (red) was obtained using HMC. The inner (outer) ellipses enclose approximately 39% (86%) of the probability mass. Black lines indicate the prior as provided by the Roy-Steiner analysis [6].

The inferred LECs at NNLO are

$$\vec{\alpha} = (c_1, c_3, c_4, \tilde{C}_{150}^{nn}, \tilde{C}_{150}^{np}, \tilde{C}_{150}^{pp}, \tilde{C}_{3S1}, C_{150}, \\ \times C_{3P0}, C_{1P1}, C_{3P1}, C_{3S1}, C_{3S1-3D1}, C_{3P2}). \quad (\text{B2})$$

The inferred LECs at N3LO are

$$\vec{\alpha} = (c_1, c_2, c_3, c_4, d_1 + d_2, d_3, d_5, d_{14} - d_{15}, \\ \times \tilde{C}_{150}^{nn}, \tilde{C}_{150}^{np}, \tilde{C}_{150}^{pp}, C_{150}, D_{150}, \tilde{C}_{3S1}, \\ \times C_{3S1}, C_{3S1-3D1}, D_{3S1}, D_{3S1-3D1}, D_{3D1}, \\ \times C_{3P0}, D_{3P0}, D_{3P1}, C_{3P1}, C_{3P2}, D_{3P2}, \\ \times D_{3P2-3F2}, C_{1P1}, D_{1P1}, D_{1D2}, D_{3D2}, D_{3D3}). \quad (\text{B3})$$

APPENDIX C: MULTIVARIATE NORMAL APPROXIMATIONS OF THE LEC POSTERiors

The LEC posteriors presented in this paper are stored as arrays of samples from which subsequent results are produced. The posteriors are, however, approximately Gaussian, and we therefore provide multivariate normal approximations for convenience. These approximations are available in text format in the supplemental material [52]. The filenames reflect the chiral order, and the contents are structured as follows:

- (i) The first line shows the ordering of the LECs.
- (ii) Then follows a vector $\vec{\mu}$ of corresponding mean values.
- (iii) Finally, a covariance matrix $\vec{\Sigma}$ encodes the uncertainties.




FIG. 9. PPDs of scattering lengths and effective ranges at NLO (blue), NNLO (purple), and N3LO (red) when including all np and pp scattering lengths up to N3LO to infer \bar{c} . Empirical results are shown as black lines, with corresponding 1σ (2σ) uncertainties as a dark (light) gray area. The empirical results are without electromagnetic effects and gathered from Ref. [3], except for a_{nn} which we take from Ref. [26]. Note that the histograms have been scaled such that they have the same peak height.

Put together, these quantities define a multivariate normal distribution $\mathcal{N}(\vec{\mu}, \vec{\Sigma})$ that approximates the LEC posterior.

APPENDIX D: POSTERIOR PREDICTIVE DISTRIBUTIONS: INCORPORATING THE c_4^{pp} DATA

For completeness, in Fig. 9 we show the PPDs for the effective range expansion with an EFT truncation error informed also by the somewhat irregular NNLO-N3LO

shift in a_{pp} ; see Table I. This irregularity is however modulated by our prior expectations on \bar{c} and the convergence of the EFT expansion. In the end, although the value for \bar{c} increases some, we find that the experimental and theoretical errors are roughly equally sized at NNLO while the theoretical (experimental) error dominates at NLO (N3LO), as before. As a result, mainly the NLO predictions for the scattering length is modified compared to the result in Fig. 3.

-
- [1] S. Weinberg, *Phys. Lett. B* **251**, 288 (1990).
[2] E. Epelbaum, H. W. Hammer, and U. G. Meißner, *Rev. Mod. Phys.* **81**, 1773 (2009).
[3] R. Machleidt and D. R. Entem, *Phys. Rep.* **503**, 1 (2011).
[4] H. W. Hammer, S. König, and U. van Kolck, *Rev. Mod. Phys.* **92**, 025004 (2020).
[5] M. Hoferichter, J. Ruiz de Elvira, B. Kubis, and U. G. Meißner, *Phys. Rep.* **625**, 1 (2016).
[6] D. Siemens, J. Ruiz de Elvira, E. Epelbaum, M. Hoferichter, H. Krebs, B. Kubis, and U. G. Meißner, *Phys. Lett. B* **770**, 27 (2017).
[7] S. Wesolowski, I. Svensson, A. Ekström, C. Forssén, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, *Phys. Rev. C* **104**, 064001 (2021).
[8] I. Svensson, A. Ekström, and C. Forssén, *Phys. Rev. C* **105**, 014004 (2022).
[9] R. J. Furnstahl, N. Klco, D. R. Phillips, and S. Wesolowski, *Phys. Rev. C* **92**, 024005 (2015).
[10] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, *Phys. Lett. B* **195**, 216 (1987).
[11] G. A. Miller, B. M. K. Nefkens, and I. Slaus, *Phys. Rep.* **194**, 1 (1990).
[12] J. L. Friar, U. van Kolck, M. C. M. Rentmeester, and R. G. E. Timmermans, *Phys. Rev. C* **70**, 044001 (2004).
[13] E. Epelbaum, U. G. Meißner, and J. E. Palomar, *Phys. Rev. C* **71**, 024001 (2005).
[14] G. Hagen *et al.*, *Nat. Phys.* **12**, 186 (2016).
[15] R. F. Garcia Ruiz *et al.*, *Nat. Phys.* **12**, 594 (2016).
[16] B. A. Brown, *Phys. Rev. Lett.* **119**, 122502 (2017).
[17] A. Koszorús *et al.*, *Nat. Phys.* **17**, 439 (2021); **17**, 539(E) (2021).
[18] X. Roca-Maza, G. Colò, and H. Sagawa, *Phys. Rev. Lett.* **120**, 202501 (2018).
[19] B. Hu, W. Jiang, T. Miyagi, Z. Sun, A. Ekström, C. Forssén, G. Hagen, J. D. Holt, T. Papenbrock, S. R. Stroberg, and I. Vernon, *Nat. Phys.* **18**, 1196 (2022).
[20] S. J. Novario, D. Lonardonì, S. Gandolfi, and G. Hagen, *arXiv:2111.12775* [nucl-th].
[21] T. Naito, G. Colò, H. Liang, X. Roca-Maza, and H. Sagawa, *Phys. Rev. C* **105**, L021304 (2022).
[22] T. Naito, X. Roca-Maza, G. Colò, H. Liang, and H. Sagawa, *Phys. Rev. C* **106**, L061306 (2022).
[23] P. G. Reinhard and W. Nazarewicz, *Phys. Rev. C* **105**, L021301 (2022).

- [24] R. N. Pérez, J. E. Amaro, and E. Ruiz Arriola, *Phys. Rev. C* **88**, 024002 (2013).
- [25] R. N. Pérez, J. E. Amaro, and E. R. Arriola, *Phys. Rev. C* **88**, 064002 (2013).
- [26] A. Gårdestig, *J. Phys. G: Nucl. Part. Phys.* **36**, 053001 (2009).
- [27] C. C. Chang, A. N. Nicholson, E. Rinaldi, E. Berkowitz, N. Garron, D. A. Brantley, H. Monge-Camacho, C. J. Monahan, C. Bouchard, M. A. Clark, B. Joó, T. Kurth, K. Orginos, P. Vranas, and A. Walker-Loud, *Nature (London)* **558**, 91 (2018).
- [28] U. van Kolck, *Few Body Syst. Suppl.* **9**, 444 (1995).
- [29] P. Reinert, H. Krebs, and E. Epelbaum, *Phys. Rev. Lett.* **126**, 092501 (2021).
- [30] M. Piarulli, L. Girlanda, R. Schiavilla, R. N. Pérez, J. E. Amaro, and E. Ruiz Arriola, *Phys. Rev. C* **91**, 024003 (2015).
- [31] A. Ekström, G. R. Jansen, K. A. Wendt, G. Hagen, T. Papenbrock, B. D. Carlsson, C. Forssén, M. Hjorth-Jensen, P. Navrátil, and W. Nazarewicz, *Phys. Rev. C* **91**, 051301(R) (2015).
- [32] P. Reinert, H. Krebs, and E. Epelbaum, *Eur. Phys. J. A* **54**, 86 (2018).
- [33] E. Epelbaum, H. Krebs, and U. G. Meißner, *Phys. Rev. Lett.* **115**, 122301 (2015).
- [34] D. R. Entem, R. Machleidt, and Y. Nosyk, *Phys. Rev. C* **96**, 024004 (2017).
- [35] E. Epelbaum, W. Gloeckle, and U. G. Meißner, *Eur. Phys. J. A* **19**, 125 (2004).
- [36] E. Epelbaum, W. Gloeckle, and U. G. Meißner, *Eur. Phys. J. A* **19**, 401 (2004).
- [37] S. Wesolowski, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, *J. Phys. G: Nucl. Part. Phys.* **46**, 045102 (2019).
- [38] 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).
- [39] R. Machleidt and I. Slaus, *J. Phys. G: Nucl. Part. Phys.* **27**, R69 (2001).
- [40] M. Göbel, T. Aumann, C. A. Bertulani, T. Frederico, H. W. Hammer, and D. R. Phillips, *Phys. Rev. C* **104**, 024001 (2021).
- [41] Q. Chen, C. R. Howell, T. S. Carman, W. R. Gibbs, B. F. Gibson, A. Hussein *et al.*, *Phys. Rev. C* **77**, 054002 (2008).
- [42] E. Epelbaum, H. Krebs, and U. G. Meißner, *Eur. Phys. J. A* **51**, 53 (2015).
- [43] J. A. Melendez, R. J. Furnstahl, D. R. Phillips, M. T. Pratala, and S. Wesolowski, *Phys. Rev. C* **100**, 044001 (2019).
- [44] Student, *Biometrika* **6**, 1 (1908).
- [45] C. G. Broyden, *IMA J. Appl. Math.* **6**, 76 (1970).
- [46] R. Fletcher, *Comput. J.* **13**, 317 (1970).
- [47] D. Goldfarb, *Math. Comput.* **24**, 23 (1970).
- [48] D. F. Shanno, *Math. Comput.* **24**, 647 (1970).
- [49] A. Griewank, *Acta Numer.* **12**, 321 (2003).
- [50] I. Charpentier and J. Utke, *Optim. Methods Software* **24**, 1 (2009).
- [51] A. F. M. Smith and A. E. Gelfand, *Am. Stat.* **46**, 84 (1992).
- [52] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevC.107.014001> for the full bivariate marginal of the LEC posterior at N3LO as well as multivariate normal approximations to the posterior at orders NLO, NNLO, N3LO.
- [53] P. Navrátil, G. P. Kamuntavicius, and B. R. Barrett, *Phys. Rev. C* **61**, 044001 (2000).
- [54] D. R. Entem and R. Machleidt, *Phys. Rev. C* **68**, 041001(R) (2003).
- [55] W. G. Jiang, A. Ekström, C. Forssén, G. Hagen, G. R. Jansen, and T. Papenbrock, *Phys. Rev. C* **102**, 054301 (2020).
- [56] P. F. Bedaque and U. van Kolck, *Annu. Rev. Nucl. Part. Sci.* **52**, 339 (2002).
- [57] A. Ekström, G. Hagen, T. D. Morris, T. Papenbrock, and P. D. Schwartz, *Phys. Rev. C* **97**, 024332 (2018).
- [58] S. König, A. Ekström, K. Hebeler, D. Lee, and A. Schwenk, *Phys. Lett. B* **810**, 135814 (2020).
- [59] A. Ekström and G. Hagen, *Phys. Rev. Lett.* **123**, 252501 (2019).
- [60] J. A. Melendez, C. Drischler, A. J. Garcia, R. J. Furnstahl, and X. Zhang, *Phys. Lett. B* **821**, 136608 (2021).
- [61] X. Zhang and R. J. Furnstahl, *Phys. Rev. C* **105**, 064004 (2022).