

Nuclear recoil response of liquid xenon and its impact on solar ^8B neutrino and dark matter searches

X. Xiang^{1,2,*}, R. J. Gaitskell,¹ R. Liu¹, J. Bang,¹ J. Xu,³ W. H. Lippincott⁴, J. Aalbers⁵,
J. E. Y. Dobson⁶, M. Szydagis⁷, G. R. C. Rischbieter⁷, N. Parveen⁷, D. Q. Huang^{1,8}, I. Olcina^{9,10},
R. J. James,⁶ and J. A. Nikoleyiczik¹¹

¹*Brown University, Department of Physics, Providence, Rhode Island 02912-9037, USA*

²*Brookhaven National Laboratory (BNL), Upton, New York 11973-5000, USA*

³*Lawrence Livermore National Laboratory (LLNL), Livermore, California 94550-9698, USA*

⁴*University of California, Santa Barbara, Department of Physics,
Santa Barbara, California 93106-9530, USA*

⁵*SLAC National Accelerator Laboratory, Menlo Park, California 94025-7015, USA*

⁶*University College London (UCL), Department of Physics and Astronomy,
London WC1E 6BT, United Kingdom*

⁷*University at Albany (SUNY), Department of Physics, Albany, New York 12222-1000, USA*

⁸*University of Michigan, Randall Laboratory of Physics, Ann Arbor, Michigan 48109-1040, USA*

⁹*University of California, Berkeley, Department of Physics, Berkeley, California 94720-7300, USA*

¹⁰*Lawrence Berkeley National Laboratory (LBNL), Berkeley, California 94720-8099, USA*

¹¹*University of Wisconsin-Madison, Department of Physics, Madison, Wisconsin 53706-1390, USA*



(Received 10 April 2023; accepted 7 July 2023; published 25 July 2023)

Knowledge of the ionization and scintillation responses of liquid xenon (LXe) to nuclear recoils is crucial for LXe-based dark matter experiments. Current calibrations carry large uncertainties in the low-energy region below ~ 3 keV_{nr} where signals from dark matter particles of < 10 GeV/ c^2 masses are expected. The coherent elastic neutrino-nucleus scattering (CE ν NS) by solar ^8B neutrinos also results in a continuum of nuclear recoil events below 3.0 keV_{nr} (99% of events), which further complicates low-mass dark matter searches in LXe experiments. In this paper, we describe a method to quantify the uncertainties of low-energy LXe responses using published calibration data, followed by case studies to evaluate the impact of yield uncertainties on ^8B searches and low-mass dark matter sensitivity in a typical ton-scale LXe experiment. We conclude that naively omitting yield uncertainties leads to overly optimistic limits by factor ~ 2 for a 6 GeV weakly interacting massive particle mass. Future nuclear recoil light yield calibrations could allow experiments to recover this sensitivity and also improve the accuracy of solar ^8B flux measurements.

DOI: [10.1103/PhysRevD.108.022007](https://doi.org/10.1103/PhysRevD.108.022007)

I. INTRODUCTION

Liquid xenon (LXe) dark matter experiments have greatly improved their sensitivities to rare nuclear recoil (NR) signals in the last decade. As of 2022, there are three multiton-scale LXe experiments in operation and searching for anticipated NR signals from weakly interacting massive particles (WIMPs)—LUX-ZEPLIN (LZ) [1], XENON-nT [2], and PandaX-4T [3]. Recently the XENON, LZ, and DARWIN collaborations have joined

force to study the possibility of building the next-generation LXe observatory for dark matter and neutrino physics [4].

The centerpiece of a typical LXe dark matter detector is a dual-phase xenon time projection chamber (TPC), where particle interactions in the liquid produce prompt scintillation light (S1) at ~ 175 nm and ionization electrons. The ionization electrons can either recombine with ions to produce additional scintillation light [5], or become liberated and drift away from the interaction site under the influence of an external electric field. Once the electrons are extracted into a gas xenon (GXe) region, they can produce secondary scintillation (S2) via electroluminescence. Combining S1 and S2 signals, LXe experiments have achieved low energy thresholds, accurate position reconstruction in 3D, and strong discrimination against electron recoil (ER) backgrounds.

*Corresponding author.
xxiang@bnl.gov

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 SCOAP³.

Coherent elastic neutrino-nuclei scattering (CE ν NS) is a neutral current interaction in which a neutrino of any flavor scatters off a nucleus as a whole, producing a nuclear recoil (NR). The process requires the momentum transfer to be significantly smaller than the inverse of the targeted nuclear size, restricting neutrino energies to below a few tens of MeV. The CE ν NS cross-section on a spin-0 nucleus with Z protons and N neutrons at rest without radiative corrections is well described by the Standard Model (neglecting a second order term) [6]:

$$\frac{d\sigma}{dE_r} = \frac{G_F^2 M_A}{\pi} \left[1 - \frac{E_r}{E_\nu} - \frac{M_A E_r}{2E_\nu^2} \right] \frac{Q_W^2}{4} F^2(q) \quad (1)$$

where E_r is the NR energy, E_ν is the neutrino energy, M_A is the target mass, G_F is the Fermi coupling constant, F is the nuclear form factor, q is the momentum transfer, and $Q_W = N - (1 - 4 \sin^2 \theta_W)Z$ is the weak nuclear charge with θ_W being the Weinberg's angle. In 2017, CE ν NS was first experimentally observed using a CsI[Na] crystal at a high energy neutrino beam by the COHERENT collaboration with 6.7-sigma significance [7].

Underground LXe TPC experiments are ideal to observe naturally occurring neutrinos via CE ν NS, thanks to the $\sim N^2$ -enhanced interaction cross section with a xenon target and their ability to separate CE ν NS from electron recoil (ER) backgrounds. With a sub-keV $_{\text{nr}}$ detection threshold, LXe detectors could observe CE ν NS for neutrino energies down to ~ 5 MeV, substantially below the neutrino energy used by the COHERENT experiment. Natural sources that can produce CE ν NS signals in underground LXe detectors include: solar ^8B neutrinos, solar *hep* neutrinos, the diffuse supernova neutrino (DSN) background, sub-GeV atmospheric neutrinos (*atm*), and the neutrinos from core-collapse of supernovae [$E_\nu \sim O(10 \text{ MeV})$]. The total flux, in units of $\text{cm}^{-2} \text{ s}^{-1}$, used in this analysis for ^8B , *hep*, DSN, and *atm*. neutrinos are: $(5.25 \pm 0.2) \times 10^6$ [8], $7.98(1 \pm 0.3) \times 10^3$, 86 ± 43 , and 10.5 ± 2.1 [9], respectively. Figure 1 shows energy spectra for the major sources of neutrinos as well as the maximum predicted Xe recoil energy from each.

Out of all natural neutrino sources, solar ^8B neutrinos produce the highest recoil rate in LXe, two orders of magnitude higher than the next most common, *hep* neutrinos, as shown in Fig. 2. Approximately 750 CE ν NS events can be expected from a tonne-year exposure in LXe. However, with a monotonically decreasing energy spectrum, 99% of the ^8B CE ν NS signals are expected to be below a $\sim 3 \text{ keV}_{\text{nr}}$. In this energy region, the detection efficiency in a LXe TPC is very low. A typical single-scatter interaction in LXe TPCs is required to contain an S1 and S2 pulse separated by the time required to drift the electrons to the gas phase. Usually an identifiable S1 signal requires the simultaneous detection of scintillation light by at least two (2-fold) or three (3-fold) PMT channels. A valid




FIG. 1. The neutrino flux that can induce observable CE ν NS events in a LXe detector. The gray area corresponds to an energy region that produces 0 observable quanta (photons and electrons) according to the NEST NR yield model, and is beyond the reach of a LXe detector. The diffused supernova neutrino background (DSNB) spectra are shown at various Fermi-Dirac temperatures in units of MeV [10]. Because the sub-GeV atmospheric neutrino remains unexplored by current experiments, the sub-GeV spectrum for dark matter experiments are based on FLUKA simulation. The low-energy cut off is caused by the lack of low-energy cosmic ray data and interaction uncertainties between cosmic rays and air nuclei as the simulation inputs [11,12]. The top x-axis is the maximum recoil energy from a neutrino back-scattering on a ^{131}Xe nucleus.

S2 signal typically requires a minimal number of 4 electrons extracted (N_{ee}) into the gas region [13]. With these practical energy thresholds taken into consideration, a 5.6 tonne fiducial mass detector like LZ [14], which requires at least a 3-fold (2-fold) S1 coincidence and a 5 N_{ee} S2 threshold, will have the total ^8B CE ν NS rate reduced to $\sim 1.8 \text{ per t} \times \text{yr}$ ($\sim 9 \text{ per t} \times \text{yr}$).

There are both opportunities and challenges that come with the appearance of CE ν NS in LXe detectors. On the one hand, the observation of CE ν NS from solar ^8B would demonstrate the low-energy nuclear recoil sensitivity of LXe experiments *in situ* and enable a new method to study natural neutrinos at the MeV-scale. For example, CE ν NS provides a unique opportunity to probe nonstandard neutrino interactions [15,16]. On the other hand, CE ν NS is an irreducible background in WIMP searches, because CE ν NS-induced recoil signatures are indistinguishable from those of WIMPs—single-scattering NRs uniformly distributed in the active Xe volume with no coincidence signals. In particular, the ^8B CE ν NS spectrum is nearly degenerate with that of a 6 GeV/ c^2 WIMP (Fig. 2), which weakens an experiment's ability to claim a discovery of WIMP interactions in this mass region. The degeneracy also amplifies any correlated uncertainty in calculating the WIMP sensitivity for masses around 6 GeV/ c^2 .




FIG. 2. The CE ν NS energy spectrum for all neutrino backgrounds that are relevant to for LXe dark matter experiments. The dashed and dotted spectra are 6 GeV (4×10^{-45} cm 2) and 100 GeV (3×10^{-49} cm 2) WIMPs whose spectrum shapes degenerate with those of ^8B and atmospheric neutrinos. The gray area corresponds to an energy region that produces 0 observable quanta (photons and electrons) according to the NEST NR yield model. The 1σ uncertainties associated with each neutrino source (transparent color bands) are the neutrino flux uncertainties.

In principle, a careful modulation analysis could resolve the phase difference between WIMP recoils (peak in June 1st [17]) and CE ν NS from solar neutrinos (peak at perihelion date around January 3rd), but this method requires plentiful statistics and is beyond the scope of this work. Directional information of the detected nuclear recoils could be another handle to discriminate between the two signals, but momentum reconstruction has not been demonstrated in any LXe experiment up to date.

As liquid xenon dark matter experiments inevitably march into the solar ^8B neutrino territory, it is crucial to quantify the common underlying nuclear recoil light yield (L_y) and charge yield (Q_y) at energies relevant for observing ^8B neutrinos. This paper focuses on demonstrating a method to quantify the low-energy L_y and Q_y uncertainties, exploring their scientific impacts for a generic LXe dark matter experiment in the presence of ^8B neutrinos. This paper is organized as follows: Sec. II reviews the published nuclear recoil calibration data on L_y and Q_y in LXe, and describes the method we adopt to quantify the low-energy yield uncertainties. Section III demonstrates the impact of the yield uncertainties on physics searches of ^8B neutrino and low-mass dark matter using a hypothetical LXe detector. Unless stated otherwise, the simulation work presented in this paper assumes a detector with the same performance as that predicted for LZ described in [14]. Section IV discusses possible strategies to mitigate this source of uncertainty. Finally, Sec. V summarizes the main results of this paper.

II. MODELING LOW-ENERGY YIELDS

The scintillation and ionization mechanisms of LXe are governed by complex micro-physics that cannot be accurately derived from first principles only. The Lindhard theory [18] has been used to estimate the magnitude and fluctuation of LXe's response to NRs. While this theory works well for recoil energies above $\mathcal{O}(10 \text{ keV}_{\text{nr}})$, deviation from the theory has been observed in a number of measurements in various target materials at low recoil energies. In this work we rely on existing nuclear recoil calibration data in the low energy region and develop a procedure to quantify the uncertainties of the yields.

A. Selection of published data

The noble element simulation technique (NEST) is a widely used package to model the L_y and Q_y responses of xenon to energy depositions within the noble liquid community [19]. Multiple experiments have reported L_y and Q_y calibration results in the energy region of interest for this analysis (below $3.7 \text{ keV}_{\text{nr}}$) [20–24], which are shown alongside the default NEST yield model v2.1 in Fig 3. These L_y and Q_y results are obtained with neutrons elastically scattering off xenon nuclei. Because current LXe dark matter detectors all operate at drift fields in the range of 100 to 500 V/cm, we leave out the ZEPLIN-III results which were measured at 3.4 kV/cm. The LUX Run3 Q_y data is included in this analysis, but it is found to make a minimal impact on the final results due to its relatively large uncertainty values.

Among the remaining calibration data, the Q_y values measured at the Lawrence Livermore National Laboratory (LLNL) using a drift field of 220 V/cm [24] and those measured in LUX Run4 using a field of 400 V/cm [21,22] both reached the sub-keV energy scale (down to $\sim 0.3 \text{ keV}_{\text{nr}}$). These two results are consistent with each other with the LLNL data reporting smaller systematic uncertainty values. The data points shown in Fig. 3 have vertical error bars calculated as the sum of the statistical and the systematical uncertainties in quadrature. Since the Q_y data trend exhibits a steep slope below $1 \text{ keV}_{\text{nr}}$, an additional Q_y uncertainty is added to the lowest points at $\sim 0.3 \text{ keV}_{\text{nr}}$ and $\sim 0.5 \text{ keV}_{\text{nr}}$ to account for recoil energy uncertainties (δE_r): $\delta Q_y = s \delta E_r$, where $s = 0.6 \text{ electrons/keV}_{\text{nr}}^2$ is an estimate of the slope.

For L_y , only the two calibration datasets from LUX are selected and their uncertainties are treated as independent. The LLNL experiment did not use reflectors to enhance light collection and thus did not report a simultaneous L_y measurement. Similar to the treatment of Q_y , the L_y error bars in Fig. 3 include statistical and systematical uncertainties summed in quadrature. Because the L_y values are relatively insensitive to NR energy, we do not include energy uncertainties in the L_y error bars.




FIG. 3. A summary of all NR L_y and Q_y measurements at various electric fields published in the past decade. The default NEST v2.1 curves come from a global fit taking into account all available data extending up to $100 \text{ keV}_{\text{nr}}$. Bottom: the detected energy spectra for various low-mass WIMPs computed using NEST v2.1 yield and the detector parameters in [14]. The solid spectra are for a detector threshold with 3-fold S1 coincidence and ≥ 5 electron S2 requirement. The dashed spectra are for a 2-fold 5 electron detector threshold.

The response of LXe to a nuclear recoil depends on the electric field strength applied at the interaction site. The original data in Fig. 3 were collected from experiments under different drift fields. These datasets are corrected to the same field using NEST. For the L_y fitting analysis, a downward scaling is applied to the LUX Run3 L_y values to align them with the LUX Run 4 L_y values at 400 V/cm . The scaling is relatively small, with a maximum shift of 7% occurring at $1 \text{ keV}_{\text{nr}}$. Similarly, we applied a constant downward scaling factor of 0.969 to the LUX Run4 Q_y values to align them with Livermore's Q_y value taken at 220 V/cm . This data scaling approach is equivalent to adjusting the drift field parameter in the NEST model.

B. NEST-based yield model

When a nuclear recoil occurs in liquid xenon, a fraction of the recoil energy, represented by f_n , is converted into observable electronic excitation and ionization, while the remainder dissipates as heat (e.g., atomic motion). [25]. We refer f_n as the quenching factor. The value of f_n depends

on the recoil energy and exhibits a non-linear functional form.

Within the NEST model, L_y and Q_y are defined as the average numbers of scintillation photons (N_{ph}) and ionization electrons (N_e) created per keV of deposited energy, respectively. Lindhard's theory provides a framework for NEST to predict the total number of quanta for a nuclear recoil (NR), and the Thomas-Imel box model offers a method to partition the signal into scintillation and ionization. With additional quenching applied to scintillation to account for the Penning effect—where two excitons can interact to produce a single photon—the NEST models can reasonably predict yields between $\sim 3 \text{ keV}_{\text{nr}}$ and $\sim 100 \text{ keV}_{\text{nr}}$ [5].

Despite substantial efforts to model the quenching effect for NR in liquid xenon, a consistent theory that can reliably predict f_n at both high and low energies remains elusive. Recent measurements below $3 \text{ keV}_{\text{nr}}$ indicate that the NEST model does not align with ionization quenching measurements for low-energy recoils [24]. To align with ionization data, a sigmoid function was integrated into the NEST model, thus empirically introducing a notable drop-off feature, as depicted in Fig. 3. This sigmoid function also compels a drop-off for L_y below $\sim 0.4 \text{ keV}_{\text{nr}}$, despite the absence of L_y data to substantiate this. P. Sorensen argued that a minimal energy threshold must exist below which the recoil energy is insufficient to create a single quanta [26], a perspective that provides a rationale for introducing the drop-off feature in the NEST models.

The NEST model parameters are obtained through a simultaneous fit to multiple calibration datasets across a wide energy range. Because the measured NR yields below $\sim 3 \text{ keV}_{\text{nr}}$ carry much larger uncertainties than those at higher energies, the NEST global fit over a wide energy range is less constrained by the low energy data points and is thus not ideal for this analysis. In addition, there are 11 free parameters in the NEST model, which pose a technical challenge for the error bands of the yields to be properly evaluated. Incorporating such a large number of variables is also computationally expensive in our profile likelihood ratio (PLR) framework, which is commonly used by LXe experiments to calculate experimental sensitivity.

We make a simplification to the NEST yield model to capture the uncertainties in the low energy region. We denote L_y^0 and Q_y^0 as the default light and charge yield values from NEST version 2.1, and obtain the parametrized L_y' and Q_y' by coupling L_y^0 and Q_y^0 to linear terms:

$$\begin{aligned} L_y' &= L_y^0 (a^{L_y} + b^{L_y} (E_r - E_0^{L_y})) \\ Q_y' &= Q_y^0 (a^{Q_y} + b^{Q_y} (E_r - E_0^{Q_y})) \end{aligned} \quad (2)$$

where E_r is the NR recoil energy in the unit of keV_{nr} and L_y' is the parametrized light yield, and Q_y' is the parametrized charge yield as a function of recoil energy.

By construction, each yield curve has two independent free parameters (a, b). The E_0 parameter serves as a constant offset for the ease of sampling (a, b) points in later steps. We choose $E_0^{Ly} = 2 \text{ keV}_{nr}$ and $E_0^{Qy} = 1 \text{ keV}_{nr}$, but the exact values of E_0 do not affect the yield values because the parametrization is linear. Henceforth, Eq. (2) is referred to as the “(a, b) model.”

The nuisance parameters (a, b) allow us to explore different variations of NEST yields; parameter a scales the yields up and down while parameter b introduces an additional degree of freedom by modifying the slopes. This model reduces to the default NEST model when $(a, b) = (1, 0)$. We comment that because the NEST model is a global fit to all data, the optimal (a, b) obtained using the selected low-energy data do not necessarily equal to (1, 0) exactly, but we anticipate them to not deviate strongly from (1, 0).

While the (a, b) model allows us to explore the variations of the entire yield curves, we expect the yields to converge to the NEST predictions beyond 3.4 keV_{nr} because the yields are well-constrained by high energy measurements. To avoid the overconstraint by high energies, we introduce a sigmoid (or Fermi-Dirac) function F_D coupling to the (a, b) terms as the following:

$$\begin{aligned} Ly' &= Ly^0[(a^{Ly} + b^{Ly}(E_r - E_0^{Ly}))F_D + (1 - F_D)] \\ Qy' &= Qy^0[(a^{Qy} + b^{Qy}(E_r - E_0^{Qy}))F_D + (1 - F_D)] \end{aligned} \quad (3)$$

The sigmoid is a function of recoil energy E_r , and it is defined as:

$$F_D = \frac{1}{1 + e^{(E_r - \mu)/\epsilon}} \quad (4)$$

where μ specifies the location of the transition and ϵ specifies the slope of the transition. When $E_r \ll \mu$, Eq. (3) is reduced to Eq. (2). When $E_r \gg \mu$, it returns to the default NEST yield model. The values of μ , and ϵ are summarized in Table I. The choice of $\epsilon = 1$ allows a smooth transition toward higher energy.

To obtain the optimal (a, b) model yields for low energy NRs and the corresponding uncertainties, we use the Pearson χ^2 as the cost function:

$$\chi^2 = \sum_i^n \frac{(y_i - f(x_i))^2}{\sigma_i^2} \quad (5)$$

TABLE I. Summary of parameters used in the (a, b) model.

	Constants			Fit Results				
	E_0	ϵ	μ	a_{\min}	b_{\min}	σ_a	σ_b	ρ_{ab}
Qy	1.0	1.0	3.4	0.92	0.05	0.060	0.052	-0.585
Ly	2.0	1.0	3.4	1.04	-0.05	0.086	0.072	-0.354

where x_i are data points from NR calibrations, σ_i are the overall uncertainty values, and f is the parametrized model. After combining systematic and statistical uncertainties in quadrature, we observe rather symmetric error bars (Fig. 3). The minimum χ^2_{\min} is obtained by performing a simple linear regression fitting of the parametrized model to the best published data. The (a_{\min}, b_{\min}) parameters found at the χ^2_{\min} are considered as the median scenario for a yield curve.

The error propagation is handled by the Minuit package [27] to compute L_y and Q_y errors across energy:

$$\sigma_f = \sqrt{\left(\frac{\partial f}{\partial a}\right)^2 \sigma_a^2 + \left(\frac{\partial f}{\partial b}\right)^2 \sigma_b^2 + 2\rho_{ab}\sigma_a\sigma_b \frac{\partial f}{\partial a} \frac{\partial f}{\partial b}} \quad (6)$$

where the values of Hesse errors σ_a , σ_b , and correlation ρ_{ab} are summarized in the Table I. The resulting uncertainty bands are shown in Fig. 4.

The Brazil error bands in Fig. 4 appear to converge to a narrow region due to the underlying sigmoid function in the




FIG. 4. The Brazilian flag uncertainty bands for the yield curves. Top: The bands represent the 1 and 2 sigma uncertainties for Q_y propagated from the (a^{Qy}, b^{Qy}) model. The Livermore data at 220 V/cm has the smallest uncertainties for Q_y . The * superscript indicates the Q_y data points are adjusted to 220 V/cm before the fit. The gray shadow marks the region where a sharp fall-off in NEST is enforced. Bottom: same as the top but for L_y . The LUX Run 4 data at equivalent 400 V/cm has the smallest uncertainties for L_y . The * superscript indicates that the LUX Run 3 L_y data points are adjusted to 400 V/cm before the fit.

NEST model. The linear parametrization method employed in this analysis does not counteract this convergence. Various hypotheses can be proposed for L_y values below 0.4 keV_{nr} in the absence of measurements, but these might not necessarily align with physical reality. We have verified that assuming a flat L_y below 0.4 keV_{nr} does not significantly impact the sensitivity presented in Sec. III, so we will not explore the wide range of assumptions below 0.4 keV_{nr} in this work.

C. Simplified yield model

We also consider a simpler version of the (a, b) parametrization with one free parameter in each yield function:

$$\begin{aligned} Ly' &= Ly^0 + a^{Ly} \\ Qy' &= Qy^0 + a^{Qy} \end{aligned} \quad (7)$$

This is similar to the XENON-1T yield model, which uses one free “interpolation parameter” to uniformly shift the Q_y curve and a free coupling parameter to scale the L_y curve [13]. We refer to this parametrization as the a model for short, and use the same treatment as the (a, b) model to smoothly transition to the default NEST values at higher energies. The uncertainty bands for this simple model are shown in the Appendix.

III. PHYSICS CASE STUDIES

We perform three case studies to evaluate the impacts of the yield uncertainties on physics searches. All three studies assume a hypothetical LXe TPC detector with the same operation condition and background levels as outlined in sensitivity studies of the LUX-ZEPLIN (LZ) experiment [14]. To highlight the effects of the yield uncertainties, we only consider physical backgrounds in this work although instrumental backgrounds including pathological electron emission and accidental backgrounds are reported in majority of currently active xenon TPCs experiments.

A. Detection efficiency

Low-energy NRs generally produce a small number of photons and electrons in the xenon target, which can be challenging to detect in a xenon TPC experiment. An inaccurate quantification of detection efficiency uncertainties would compromise an experimental sensitivity to signals that produce weak NRs. We calculate the signal efficiency at a specific recoil energy as the fraction of uniformly distributed monoenergetic NR events that survive the detector threshold cuts.

The yield uncertainties are propagated into the detection efficiency by sampling the yield models following a χ^2 distribution, as explained below. First, with the yield model parametrization explained in Sec. II B, we randomly pick




FIG. 5. Top: the median NR detection efficiency (dashed) as a function of true recoil energy for two sets of S1 + S2 detector thresholds. The two shaded color bands are the 68% and 95% uncertainties from L_y . On the same plot, we show differential rate spectra for solar ^8B CE ν NS surviving the detection thresholds, whose normalization is indicated by the y-axis on the right. The shaded blue and red bands are derived from the 1σ uncertainty in detection efficiencies. The region between two vertical dotted lines indicate where the central 68% of the events are. Bottom: the fractional uncertainties, defined as the difference from the $\pm 1\sigma$ probability to the median probability divided by the median probability, as a function of energy.

(a, b) values uniformly distributed around (1, 0) and calculate the corresponding χ^2 value [Eq. (5)]. Then we accept the yield model at each (a, b) point with a probability determined by its distance to the minimum, $\delta\chi^2 = \chi^2 - \chi^2_{\min}$, using the χ^2 PDF with 2 degrees of freedom. Next, NEST simulations are performed for each L_y and Q_y model and produce an estimate of the detection efficiency at each NR energy, based on which the spread of NR detection efficiency is computed. The uncertainty bands in Fig. 5 represent the central 68% (heavily shaded) and 95% (lightly shaded) quantiles of the efficiency curves. We stop the calculation at 0.4 keV_{nr} due to the prohibitive computation time and the impractically low detection efficiency.

Figure 5 also shows the ^8B energy spectra convolved with the efficiency curves to illustrate the relevant energy region for ^8B . When the S1 coincidence threshold requirement is relaxed from the 3-fold to 2-fold, Fig. 5 top shows a drastically improved efficiency that results in an overall increase of ^8B event rate by a factor of ~ 5 , while Fig. 5-Bottom shows a significantly reduced relative uncertainty in efficiency. Consequently, lowering the S1 coincidence threshold in a LXe detector is advantageous to probe the low-energy NR phenomenon in the absence of additional background events.

B. Constraint of ^8B flux

A precise measurement of the ^8B CE ν NS rate could yield evidence of new physics complementing results from oscillation and spallation source experiments. This study evaluates an experiment's capability to measure the ^8B neutrino flux using Monte Carlo-simulated data in the absence of WIMPs. The standard CE ν NS interaction with the Helm nuclear form factor was adopted, and seasonal flux modulations are not considered in this analysis. The detector signals are generated using the default L_y and Q_y models in NEST v2.1 but the statistical analysis assume that the yield values are uncertain with spreads estimated in Sec. II.

Figure 6 (top) shows the S2 distribution of the simulated dataset, along with background PDFs generated using the NEST v2.1 package. The simulation assumes an S1 threshold of 3-fold coincidence and an S2 threshold of 5 extracted electrons. A typical WIMP analysis in a LXe detector is performed in (S1, S2) observable space. In this analysis,

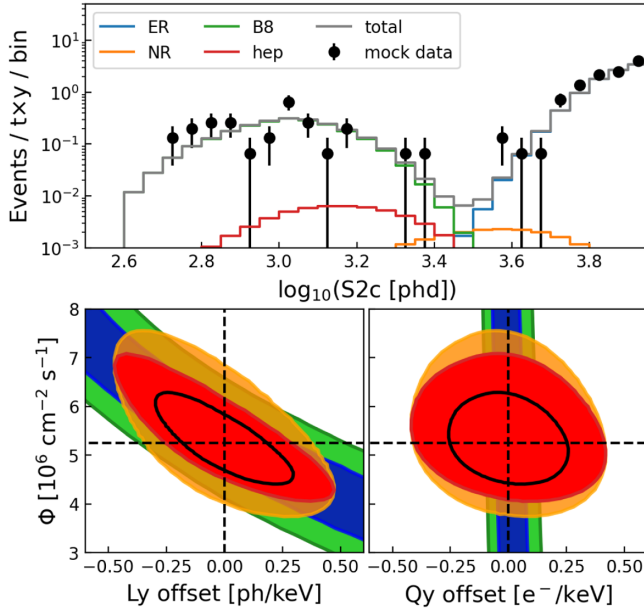


FIG. 6. Top: the mock data generated for a $15.3 \text{ t} \times \text{yr}$ exposure using the nominal ^8B flux ($\Phi = 5.25 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$) [28] and the default NEST L_y and Q_y curves. The detector thresholds for S1 is 3-fold coincidence and for S2 is $5 N_{ee}$. The error bars are Poissonian. The ER component combines all background labeled as “ER” in Table IV of [14]. The NR combines all “NR” labeled background except for solar *hep* neutrino. The bin width is $0.05 \log(\text{phd})$. Bottom: projections of the 90% confidence volumes in θ_{L_y} and Φ parameter space. Green (blue) stripe shows the confidence interval for a $15.3 \text{ t} \times \text{yr}$ ($50 \text{ t} \times \text{yr}$) exposure without any constraints of L_y and Q_y . The orange (red) contour shows the interval for a $15.3 \text{ t} \times \text{yr}$ ($50 \text{ t} \times \text{yr}$) exposure with L_y and Q_y constrained to the calibration data using the model in Sec. II C. The solid black line is a $50 \text{ t} \times \text{yr}$ exposure where we artificially reduce L_y and Q_y uncertainties to half.

however, beyond the trigger requirement, including S1 in the analysis does not lead to a significant improvement in the result due to the small S1 signal amplitudes for both signals and background in this energy region. Therefore, this study projects the 2-dimensional data distribution onto the one-dimensional S2 observable space before performing any inference analysis.

A significant excess of events due to ^8B signals can be observed clearly over known sources of physical backgrounds, as shown in Fig. 6 (top). As discussed above (Sec. III), all LXe TPCs observe some levels of instrumental backgrounds, often called accidental backgrounds, arising from accidental coincidence of spurious S1s (e.g., PMT dark count coincidence) and S2s (e.g., grid electron emission). The characteristics of these backgrounds vary from experiment to experiment, and are still being actively studied [13,29,30], making a generic model unreliable. For one example, the XENON1T experiment reported accidental backgrounds at $0.47 \text{ events/t} \times \text{yr}$ rate in a 1.3 t fiducial mass but only $0.08 \text{ events/t} \times \text{yr}$ in a 0.65 t core mass [31]. Here we restrict our study to the known physical backgrounds to highlight the effect of the low-energy yield uncertainties, leaving out the instrumental backgrounds.

To construct confidence intervals, we perform an extended unbinned maximum likelihood fits of the mock data using the Minuit2 package. The model has three parameters of interest: the ^8B flux Φ , a L_y parameter θ_{L_y} , and a Q_y parameter θ_{Q_y} , where θ_{L_y} and θ_{Q_y} provide offsets to shift the default NEST yield curves up and down as prescribed in Sec. II C. Although both θ_{L_y} and θ_{Q_y} can affect the amplitude and shape of NR PDFs, significant changes in shape only occur under extreme variations of L_y and Q_y outside the yield boundaries constrained by published data. Therefore we only allow the amplitude to vary in this work.

The full log-likelihood function consists of two parts—an extended event likelihood term ($\mathcal{L}^{\text{event}}$) and a profile term ($\mathcal{L}^{\text{profile}}$) to constrain the nuisance parameters (θ) by their associated uncertainties:

$$\log \mathcal{L}^{\text{total}}(\Phi, \theta) = \log \mathcal{L}^{\text{event}}(\Phi, \theta) + \log \mathcal{L}^{\text{profile}}(\theta) \quad (8)$$

where Φ is the ^8B flux, and the two nuisance parameters are θ_{L_y} and θ_{Q_y} . For a given dataset \mathcal{D}_{obs} , the event likelihood is defined as the following:

$$\mathcal{L}^{\text{event}}(\Phi, \theta | \mathcal{D}_{\text{obs}}) = \text{Pois}(n_{\text{obs}}; \mu_{\text{tot}}) \prod_{i=1}^{n_{\text{obs}}} \left[\mu_s(\Phi, \theta) f_s(x_i) + \sum_b \mu_b(\theta) f_b(x_i) \right] \quad (9)$$

where the first term is a Poisson distribution, n_{obs} is the number of observed events in this dataset, μ_{tot} is the total expected number of events, μ_s is the expected signal

number, μ_b is the expected background number, and $f_s(\mathbf{x}_i)$ and $f_b(\mathbf{x}_i)$ are the normalized signal and background PDFs in the observable space \mathbf{x} (i.e., logS2). By definition, $\mu_{\text{tot}} = \mu_s + \sum_b \mu_b$. The signal number μ_s is a function of θ_{L_y} and θ_{Q_y} , while all expected background number μ_b are fixed to their respective nominal values.

The profile term are used to constrain the yield nuisances as the following:

$$\mathcal{L}^{\text{profile}}(\theta) = \mathcal{G}(\theta_{L_y}; 0, \sigma_{L_y}) \times \mathcal{G}(\theta_{Q_y}; 0, \sigma_{Q_y}) \quad (10)$$

where \mathcal{G} are the Gaussian whose width σ_{L_y} (σ_{Q_y}) are derived from fitting external L_y (Q_y) calibration data.

Figure 6 (bottom) shows the estimated 90% confidence intervals for the obtained ^8B flux. The green and blue stripes represent the results when the L_y and Q_y constraints are removed; in this case a measurement of the ^8B flux is impossible due to the strong correlation between the flux and the yield values. When constraints in L_y and Q_y are imposed using NR calibration data, the yield-flux degeneracy is broken and measurements of the ^8B flux can be achieved for projected exposures of 15.3 t \times yr (orange) and 50 t \times yr (red). The black contour is for the hypothetical scenario where the L_y and Q_y uncertainties [21,22,24] are reduced by half, illustrating how improved calibrations can lead to better confined confidence intervals for the ^8B flux in an experiment.

C. WIMP sensitivity

In this section, we evaluate the projected WIMP sensitivity (i.e., the exclusion limit to reject elastic spin-independent WIMP-nucleon scattering at a 90% confidence level) for the hypothetical LXe experiment in the presence of the ^8B background with uncertain experimental L_y and Q_y yield values.

The sensitivity is calculated using the profile likelihood ratio (PLR) software developed by the LZ collaboration. This analysis follows the convention recommended by [9] to use a two-sided Frequentist test statistics, a strictly positive signal strength estimator, and avoids the asymptotic approximation. The WIMP interaction rate is calculated using the Helm nuclear form factor and the standard halo model with the astrophysical parameter values [32–36] summarized in [9].

The event model is constructed from two-dimensional PDFs in (S1, logS2) for each background and signal source. We use the same background models as explained in [14] but merge similar background PDFs together to increase the computation speed so we can explore multiple scenarios with different yield assumptions. In particular, all ER background PDFs are combined after weighing each component by its expected event rate; the atmospheric neutrino (ν_{atm}) and diffused supernova neutrino (ν_{DSN}) are absorbed into the environmental neutron background PDF.

This simplification is not expected to affect the final outcome because these backgrounds have minimal contamination in the signal region: < 0.1 events of ER backgrounds may leak into the 90% event contour of a 10 GeV/ c^2 WIMPs in 15.3 t \times yr of exposure; ν_{atm} and ν_{DSN} , when combined, only contribute 0.024 events/t \times yr in the energy range from .25 to 6.0 keV $_{\text{nr}}$. Similar to the previous case study, instrumental backgrounds are not considered here.

The construction of the likelihood function follows Eqs. (8) and (9), except that the flux Φ is replaced by the WIMP-nucleon cross section $\sigma_{\chi-N}$, and the background rates $\mu_b(\theta)$ are no longer constants. Since the PDF shapes of low-mass WIMPs (≤ 10 GeV) and ^8B backgrounds are insensitive to modest yield variations near the measured values, we define the L_y nuisance parameters as relative scaling factors ($\theta_s^{L_y}, \theta_b^{L_y}$). These parameters are coupled to the signal and background rates as $\mu_s = \mu_s^0 \theta_s^{L_y}$ and $\mu_b = \mu_b^0 \theta_b^{\Phi} \theta_b^{L_y}$, where θ_b^{Φ} nuisance account for ^8B flux uncertainty, μ_s^0 and μ_b^0 are the nominal WIMP and ^8B rates in the absence of uncertainty. This treatment was verified to not affect the outcome while simplifying the PLR computation.

A profile term is defined to capture the strong correlation between the WIMP signal and the ^8B background as a result of low-energy yield uncertainties:

$$\mathcal{L}^{\text{profile}}(\theta) = \prod_b \mathcal{G}_b(\theta_b; a_b, \sigma_b) \times \mathcal{C}(\theta_s^{L_y}, \theta_{8\text{B}}^{L_y}) \quad (11)$$

where each \mathcal{G}_b is a Gaussian constraint for a background rate indexed as b , and \mathcal{C} is a 2-dimensional analytic constraint function describing the correlation between low-mass WIMPs at a specific mass value and ^8B neutrinos as a result of the L_y variation. The correlation function is obtained from NEST simulations ahead of the PLR computation using a procedure similar to that described in Sec. III A. For each set of (a, b) parameters sampled according to their probability distribution function, we simulate ^8B and WIMP events in (S1, logS2) space using NEST. Each NEST run returns a ^8B rate and a WIMP rate, which are recorded in a 2-dimensional histogram. Then we obtain a distribution that describes the correlation between the two scaling factors $\theta_s^{L_y}$ and $\theta_b^{L_y}$ arising from the yield uncertainties. Finally, we fit an empirical function \mathcal{C} to describe the distribution smoothly. The following functional form works well in our case:

$$\mathcal{C}(\theta_s^{L_y}, \theta_{8\text{B}}^{L_y}) = \mathcal{G}(\theta_{8\text{B}}^{L_y}) \mathcal{G}(\theta_s^{L_y} | m(\theta_{8\text{B}}^{L_y}), w(\theta_{8\text{B}}^{L_y})) \quad (12)$$

where the second Gaussian is a conditional PDF whose mean m quadratically depends on $\theta_{8\text{B}}^{L_y}$ and width w linearly depends on $\theta_{8\text{B}}^{L_y}$. The fitting is easier done in two steps—first




FIG. 7. Top: The projected median sensitivity (90% confidence limit) for a $15.3 \text{ t} \times \text{yr}$ of LXe exposure as a function of WIMP mass. The solid curves assume the standard 3-fold $5 N_{ee}$ detector threshold for the LZ-like detector, and dashed curves assume a relaxed 2-fold threshold $5 N_{ee}$ detector threshold. The black curves are the baseline scenario where neither L_y nor Q_y uncertainty is included. Green curves include Q_y uncertainty only, orange curves include L_y uncertainty derived from the single parameter model (Sec. II C), red curves include L_y uncertainty derived from (a, b) models (Sec. II B). Bottom: the ratio, computed by taking the colored curves from the top plot, divided by the black curves (baseline).

we fit the x-projection of the histogram with $\mathcal{G}(\theta_{8B}^{L_y})$, and then we proceed to fit the entire 2D histogram with the 2 parameters fixed. The profiling term for Q_y follows the same construction.

The evaluated median WIMP sensitivities for a projected $15.3 \text{ t} \times \text{yr}$ exposure are shown in Fig. 7. The black curves represent the ideal scenario where both the L_y and Q_y yields are precisely known, and the colored curves include different yield uncertainties, which produce worsened limits as expected. Three observations can be made from Fig. 7. First, the yield uncertainties impact WIMP masses around 5.5–6 GeV the most due to the degeneracy between ^8B CE ν NS and WIMP spectra for those masses. For WIMP masses around 10 GeV/ c^2 or higher, the effect of the yield uncertainties diminishes and can be safely neglected, because the WIMP spectra significantly diverges from that of ^8B . Second, the WIMP sensitivity obtained with a 2-fold S1 coincidence threshold is more susceptible to the yield uncertainties than that with a 3-fold threshold, even though the 2-fold threshold generally leads to improved limits. This may be explained as the larger yield uncertainties at lower NR energies. Third, the impact of L_y on WIMP sensitivities dominates over that of Q_y due to the larger L_y uncertainty in available calibration data, which




FIG. 8. The projected 6 GeV/ c^2 WIMP sensitivity (the median 90% exclusion limit) for a 5.6-ton LZ-like detector as a function of live days. The solid curves are the detector threshold of 3-fold S1 coincidence and $5 N_{ee}$ S2 requirement, and the dashed curves are a relaxed 2-fold $5 N_{ee}$ detector threshold. The black curves are the naive scenario where yield uncertainty is not included. The shaded bands are the 68% overall uncertainties from PLR calculation.

highlights the need to improve L_y accuracy in future NR calibrations.

The effect of yield uncertainties on projected WIMP sensitivities as a function of accumulated exposure is shown in Fig. 8, where the color scheme is the same as used in Fig. 7. In the ideal scenario of zero yield uncertainties, the projected WIMP sensitivity continues to improve without any signs of saturation as the exposure increases. However, when L_y uncertainties are introduced, the projected sensitivity of the hypothetical experiment to 6 GeV/ c^2 WIMP begins to saturate around an exposure of 30.7 tonne \cdot year. This effect is not as significant at low exposure values because the expected ^8B rate is small and the sensitivity is mostly limited by Poisson fluctuation. The transition between the two regimes occurs at ~ 700 tonne \cdot day. LZ is currently the leading LXe dark matter experiment with a reported exposure of 330 tonne \cdot day [1], and has not yet reached the saturation point. As LZ and other ton-scale LXe experiments continue to accumulate statistics, they may find it difficult to constrain the ^8B background rate at $< 20\%$ accuracy without further reducing the L_y systematic uncertainty. Eventually this issue may cause them to approach the neutrino floor near the 6 GeV/ c^2 mass sooner than previously expected.

IV. DISCUSSION

Traditionally LXe dark matter experiments mostly focus on the WIMP mass range of 10 GeV/ c^2 to 1 TeV/ c^2 , but

recent progresses have enabled them to start probing WIMP masses below $10 \text{ GeV}/c^2$ [37–39]. In this mass region, a LXe experiment’s WIMP sensitivity is negatively impacted by the presence of ^8B CE ν NS interactions, which is further amplified by the large uncertainties in measured NR yields as demonstrated in Sec. III.

LXe direct detection experiments hold a unique position as they can detect both CE ν NS from solar ^8B and elastic neutrino-electron scattering from solar pp . Leveraging LXe’s excellent ER/NR discrimination, LXe experiments can simultaneously measure pp and ^8B events, thus providing a crucial data point that complements results from oscillation and spallation source experiments. complement to the results from oscillation and spallation source experiments. However, the low-energy yield uncertainties could undermine an LXe experiment’s capacity to accurately measure the ^8B CE ν NS, obscure new physics beyond the Standard Model.

Figure 9 illustrates the estimated uncertainty of a ^8B measurement assuming different levels of L_y calibration accuracy. In this study, we artificially reduce the overall L_y uncertainties reported in [21,22] by a chosen factor and then use the method explained in Sec. III A to estimate the ^8B rate uncertainty. Assuming a detector threshold of 3-fold S1 and 5-electron S2 and a projected $15.3 \text{ t} \times \text{yr}$ full exposure for LZ, the L_y uncertainty needs to be reduced to 93% of its current value in order to match the level of Poisson fluctuations, as illustrated by the horizontal dashed blue line. If a 2-fold S1 threshold is used, a stronger reduction factor of 0.67 will be required, as illustrated by the horizontal dashed orange line. Since the statistical uncertainty is proportional to the inverse square root of exposure, future LXe experiments would need even stricter calibration requirements to take full advantage of the larger

exposures, as illustrated by the horizontal dotted lines in Fig. 9 for a hypothetical $50 \text{ t} \times \text{yr}$ exposure experiment.

Elastic neutron scattering is a demonstrated method to produce nuclear recoils of desired energy distributions, and is the approach used in the most precise calibrations in LXe discussed in Sec. II A. To improve the L_y calibration accuracy in the keV energy region, low-energy mono-energetic neutrons, such as those produced by (p,n) interactions or through down-scattering of MeV DD fusion neutrons on a deuterium- or hydrogen-rich reflector [40], will be needed. To control the systematic uncertainties in such calibrations, neutron timing derived from the neutron beam or by instrumenting the down-scattering neutron reflector volumes can provide critical background rejection power. Photo-neutron sources such as YBe and SbBe have also been proposed for calibrations of dark matter detectors in the low-energy region. A well-designed calibration experiment with carefully controlled systematics may provide the much-needed improvement in L_y to recover the losses of light mass WIMP and ^8B sensitivities in large LXe experiments.

V. CONCLUSION

We developed NEST-based models to parametrize the low-energy NR ionization and scintillation yields. We propagate the yield uncertainties into the WIMP sensitivity in the mass range of 4 to $10 \text{ GeV}/c^2$ in the presence of solar ^8B neutrino CE ν NS backgrounds. In this work, we primarily used calibration data from LUX Run3 [23] and Run4 [21,22] to constrain the L_y uncertainty below $3.4 \text{ keV}_{\text{nr}}$, and data from Livermore [24] and LUX Run4 [21,22] to constrain the Q_y uncertainty below $3.4 \text{ keV}_{\text{nr}}$. We found that the dominant systematical uncertainty in terms of the expected ^8B CE ν NS rate in LZ comes from the L_y , with a fractional uncertainty of $\sim 20\%$ for 3-fold, 5 N_{ee} detector threshold, and $\sim 14\%$ for 2-fold 5 N_{ee} threshold. The low-energy Q_y systematic uncertainty in terms of predicted ^8B rate is $\sim 6\%$ for a $N_{ee} > 5$ threshold and can be absorbed into the L_y nuisance parameter during the PLR calculation.

The yield uncertainties introduce correlations between ^8B background and low-mass WIMP signals, which constructively amplify the impact on the WIMP sensitivity. In addition, we obtained a more conservative projections when we parametrized L_y with the (a, b) model (Sec. II B) than simply shifting the L_y up and down (Sec. II C), suggesting more than one degree of freedom is needed to characterize the L_y uncertainty.

The impact of yield uncertainties on WIMP sensitivity is the strongest at $5.5\text{--}6.0 \text{ GeV}/c^2$, but diminishes rather quickly toward 10 GeV . The effect of yield uncertainties is subdominant in smaller exposures ($< 1.5 \text{ t} \times \text{yr}$) due to Poisson fluctuations. However, the long-term impact becomes more severe as the exposure increase, ultimately

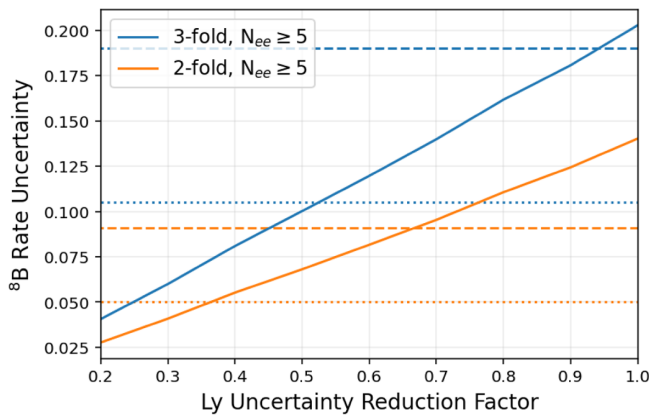


FIG. 9. Calibration requirement of the L_y uncertainties for the ^8B rate at two different thresholds. The dashed lines are the corresponding levels of Poisson uncertainties for ^8B neutrino in a projected $15.3 \text{ t} \times \text{yr}$ exposure. The dotted lines are for a $50 \text{ t} \times \text{yr}$ exposure.

reaching saturation for 6 GeV/c² WIMP around 20 t × yr. Hence, it is crucial for the next generation LXe dark matter experiment to carry out better low-energy NR calibrations, especially for constraining the uncertainty in L_y .

ACKNOWLEDGMENTS

This work used two software packages developed by the LZ Collaboration—LZNESTUTILS, a NEST interfaces package, and LZSTATS, a profile likelihood ratio statistical package. This work used computing resources from the Center for Computation and Visualization provided by Brown University, and from the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract Nos. DE-AC02-05CH11231 and DE-SC0010010. The authors would like to thank many NEST and LZ collaboration members for insightful discussion and feedback.

APPENDIX: SIMPLIFIED YIELD MODEL UNCERTAINTIES

One can derive the error envelope by varying one parameter per yield. For instance, Sec. II C illustrate an alternative method by shifting the yield curves up and down by a constant. This method generates a relatively conservative error bands as shown in Fig. 10 in comparison to the (a, b) parametrization. The fitting procedure and error propagation are identical to the description in the main paragraph. There are implicit shape variations that are not characterized by simply having one parameter shifting the NEST yield curve up and down. However, the shape variation only becomes significant when a LXe experiment is searching physics at very large exposures. It is convenient for the current generation LXe experiments to stick with one parameter variation, since they are likely to report low-energy physics results at a lower exposure.




FIG. 10. The Brazilian flag 1 and 2 sigma error bands for Q_y and L_y propagated from the single parameter model, similar to Fig. 10.

Unless L_y calibration is significantly improved, it is recommended for the next generation experiment to incorporate two parameters in characterizing L_y uncertainties when calculating the physics results. Furthermore, the one parameter method is likely sufficient for an LXe experiment that is overwhelmed by other sources of uncertainties in the same low-energy region, such as the instrumental background rising from random coincidence of phony S1s and spurious S2s.

-
- [1] J. Aalbers *et al.* (LZ Collaboration), [arXiv:2207.03764](#).
 - [2] E. Aprile *et al.* (XENON Collaboration), *J. Cosmol. Astropart. Phys.* **11** (2020) 031.
 - [3] Y. Meng *et al.* (PandaX-4T Collaboration), *Phys. Rev. Lett.* **127**, 261802 (2021).
 - [4] J. Aalbers *et al.*, *J. Phys. G* **50**, 013001 (2023).
 - [5] B. Lenardo, K. Kazkaz, A. Manalaysay, J. Mock, M. Szydagis, and M. Tripathi, *IEEE Trans. Nucl. Sci.* **62**, 3387 (2015).
 - [6] D. Z. Freedman, *Phys. Rev. D* **9**, 1389 (1974).
 - [7] D. Akimov *et al.* (COHERENT Collaboration), *Science* **357**, 1123 (2017).
 - [8] B. Aharmim *et al.* (SNO Collaboration), *Phys. Rev. C* **88**, 025501 (2013).
 - [9] D. Baxter *et al.*, *Eur. Phys. J. C* **81**, 907 (2021).
 - [10] L. E. Strigari, *New J. Phys.* **11**, 105011 (2009).
 - [11] M. Honda, T. Kajita, K. Kasahara, and S. Midorikawa, *Phys. Rev. D* **83**, 123001 (2011).
 - [12] E. Richard *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. D* **94**, 052001 (2016).
 - [13] E. Aprile *et al.* (XENON Collaboration), *Phys. Rev. Lett.* **126**, 091301 (2021).
 - [14] D. S. Akerib *et al.* (LUX-ZEPLIN Collaboration), *Phys. Rev. D* **101**, 052002 (2020).

- [15] D. Akimov *et al.* (COHERENT Collaboration), *Phys. Rev. Lett.* **126**, 012002 (2021).
- [16] S. J. Brice *et al.*, *Phys. Rev. D* **89**, 072004 (2014).
- [17] K. Freese, M. Lisanti, and C. Savage, *Rev. Mod. Phys.* **85**, 1561 (2013).
- [18] J. Lindhard, M. Scharff, and H. E. Schiøtt, *Kgl. Danske Videnskab. Selskab. Mat. Fys. Medd.* **33**, 14 (1963), <https://www.osti.gov/biblio/4701226>.
- [19] M. Szydagis, G. A. Block, C. Farquhar, A. J. Flesher, E. S. Kozlova, C. Levy, E. A. Mangus, M. Mooney, J. Mueller, G. R. C. Rischbieter, and A. K. Schwartz, *Instruments* **5**, 13 (2021).
- [20] E. Aprile, M. Anthony, Q. Lin, Z. Greene, P. De Perio, F. Gao, J. Howlett, G. Plante, Y. Zhang, and T. Zhu, *Phys. Rev. D* **98**, 112003 (2018).
- [21] D. Huang, Ultra-low energy calibration of the LUX and LZ dark matter detectors, Ph.D. thesis, Brown University, 2020, 10.26300/zvs6-fx07.
- [22] D. S. Akerib *et al.* (LUX Collaboration), [arXiv:2210.05859](https://arxiv.org/abs/2210.05859).
- [23] D. S. Akerib *et al.* (LUX Collaboration), [arXiv:1608.05381](https://arxiv.org/abs/1608.05381).
- [24] B. G. Lenardo, J. Xu, S. Pereverzev, O. A. Akindele, D. Naim, J. Kingston, A. Bernstein, K. Kazkaz, M. Tripathi, C. Awe, L. Li, J. Runge, S. Hedges, P. An, and P. S. Barbeau, *Phys. Rev. Lett.* **123**, 231106 (2019).
- [25] P. Sorensen and C. E. Dahl, *Phys. Rev. D* **83**, 063501 (2011).
- [26] P. Sorensen, *Phys. Rev. D* **91**, 083509 (2015).
- [27] F. James and M. Roos, *Comput. Phys. Commun.* **10**, 343 (1975).
- [28] M. Anderson *et al.* (SNO+ Collaboration), *Phys. Rev. D* **99**, 012012 (2019).
- [29] D. S. Akerib *et al.* (LUX Collaboration), *Phys. Rev. D* **104**, 012011 (2021).
- [30] D. S. Akerib *et al.*, [arXiv:2101.08753](https://arxiv.org/abs/2101.08753).
- [31] E. Aprile *et al.* (XENON Collaboration), *Phys. Rev. Lett.* **121**, 111302 (2018).
- [32] J. D. Lewin and P. F. Smith, *Astropart. Phys.* **6**, 87 (1996).
- [33] C. McCabe, *J. Cosmol. Astropart. Phys.* **02** (2014) 027.
- [34] M. C. Smith *et al.*, *Mon. Not. R. Astron. Soc.* **379**, 755 (2007).
- [35] R. Schoenrich, J. Binney, and W. Dehnen, *Mon. Not. R. Astron. Soc.* **403**, 1829 (2010).
- [36] J. Bland-Hawthorn and O. Gerhard, *Annu. Rev. Astron. Astrophys.* **54**, 529 (2016).
- [37] D. S. Akerib *et al.* (LUX Collaboration), *Phys. Rev. D* **101**, 042001 (2020).
- [38] D. S. Akerib *et al.* (LUX Collaboration), *Phys. Rev. Lett.* **122**, 131301 (2019).
- [39] E. Aprile *et al.* (XENON Collaboration), *Phys. Rev. Lett.* **123**, 251801 (2019).
- [40] J. R. Verbus *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **851**, 68 (2017).