Measurement of charge and light yields for ^{127}Xe L-shell electron captures in liquid xenon

D. J. Temples[1,*] J. McLaughlin[2] J. Bargemann[3] D. Baxter[1,†] A. Cottle[4,‡] C. E. Dahl[1,2]
W. H. Lippincott[3,‡] A. Monte[3,‡] and J. Phelan[5,†]

[1]*Fermi National Accelerator Lab (FNAL), Astrophysics Department, Batavia, Illinois 60510-5011, USA*
[2]*Northwestern University, Department of Physics and Astronomy, Evanston, Illinois 60208-3112, USA*
[3]*University of California, Santa Barbara, Department of Physics, Santa Barbara,
California 93106-9530, USA*
[4]*University of Oxford, Department of Physics, Oxford OX1 3RH, United Kingdom*
[5]*Massachusetts Institute of Technology, Department of Physics, Cambridge,
Massachusetts 02139-4307, USA*

Dark matter searches using dual-phase xenon time-projection chambers (LXe-TPCs) rely on their ability to reject background electron recoils (ERs) while searching for signal-like nuclear recoils (NRs). ER response is typically calibrated using β-decay sources, such as tritium, but these calibrations do not characterize events accompanied by an atomic vacancy, as in solar neutrino scatters off inner-shell electrons. Such events lead to emission of x rays and Auger electrons, resulting in higher electron-ion recombination and thus a more NR-like response than inferred from β-decay calibration. We present a cross-calibration of tritium β-decays and ^{127}Xe electron-capture decays (which produce inner-shell vacancies) in a small-scale LXe-TPC and give the most precise measurements to date of light and charge yields for the ^{127}Xe L-shell electron-capture in liquid xenon. We observe a 6.9σ (9.2σ) discrepancy in the L-shell capture response relative to tritium β decays, measured at a drift field of 363 ± 14 V/cm $(258 \pm 13$ V/cm), when compared to simulations tuned to reproduce the correct β-decay response. In dark matter searches, use of a background model that neglects this effect leads to overcoverage (higher limits) for background-only multi-kiloton-year exposures, but at a level much less than the 1-σ experiment-to-experiment variation of the 90% C.L. upper limit on the interaction rate of a 50 GeV/c^2 dark matter particle.

I. INTRODUCTION

For over a decade, dual-phase xenon time projection chambers (LXe-TPCs) have led the search for particle dark matter in the ~10 GeV–10 TeV mass range [1–4], which includes weakly interacting massive particles [5,6]. The success of the LXe-TPC technique stems from both its scalability and its ability to discriminate against events arising from ambient radioactivity using a combination of self-shielding, position reconstruction, and electron recoil (ER) vs nuclear recoil (NR) discrimination. The ER/NR discrimination is critical to eliminate low-energy, single-scatter ER backgrounds distributed uniformly in the detector, e.g., from the β decays of ^{85}Kr and ^{214}Pb. The three multiton LXe-TPC experiments that are now underway around the world [7–9] are sensitive to an additional low-energy ER background that must be addressed by ER/NR discrimination: neutrino-electron scattering by solar neutrinos [7,10]. The tacit (and not unreasonable) assumption made by these experiments to date is that the ERs produced by neutrino scatters will have the same, well-calibrated signature as β decays. The measurement described in this paper shows that this assumption of a universal ER response does not hold in LXe-TPCs, and a new model for the neutrino-electron scattering response is presented.

A. ER/NR discrimination and Auger cascades

Energy depositions in an LXe-TPC generate both prompt scintillation light (S1) and ionization. The ionization drifts under an applied electric field to a gas region where it produces a delayed scintillation signal (S2) via electroluminescence [11]. The ratio of ionization to scintillation (S2/S1) is, on average, larger for ERs than for NRs, for the same amount of S1 light. That ratio is determined by both the relative amounts of ionization and atomic excitation

---
*dtemples@fnal.gov
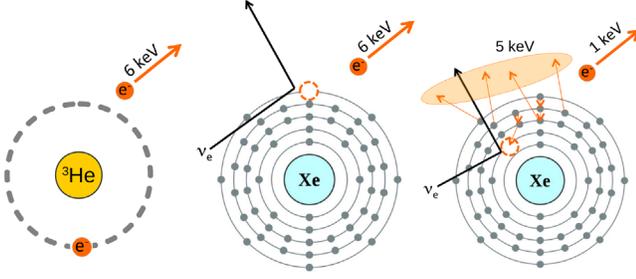†Formerly at 2.
‡Formerly at 1.

FIG. 1. A diagram depicting a β decay from tritium (left), a valence recoil from neutrino or Compton scattering (center), and an inner-shell recoil (right). Each of these interactions deposits a total of 6 keV via electron recoil(s), but the inner-shell recoil produces a distinct event topology from either the β decay or valence recoil.

initially present in the ER (or NR) track and by the fraction of ionization electrons that recombine with xenon ions at the interaction site, simultaneously reducing S2 and increasing S1 [12].

A neutrino scattering off a valence atomic electron will likely have the same signature as an equal-energy naked β decay (i.e., a β decay directly to the ground state of the daughter nucleus) because, in both cases, the energy deposited in the detector is given entirely to a single recoiling electron. The story changes when a neutrino scatters off an inner-shell electron, leaving a vacancy in the atomic structure of the struck xenon atom. In this case, the unstable atomic state relaxes by the emission of either an x ray or an Auger electron, with the potential to create additional atomic vacancies. The total energy deposited by these relaxation processes will equal the binding energy of the ionized inner-shell electron. Scatters off L-shell electrons are of particular interest to LXe-TPC dark matter searches as the xenon L-shell energy (4.8–5.4 keV) lies in the middle of the typical 1.5–15 keV$_{ee}$ (electron-equivalent) dark matter region of interest. In xenon, vacancies at the L-shell and beyond most frequently relax

via Auger emission of an electron from the next shell out, creating a cascade of Auger emission [13–15]. In this manner, a 6 keV energy deposition by a neutrino scatter on the L shell can be shared between up to eight electrons, all coming from the same xenon atom and all with energies below 2.5 keV. Figure 1 shows a schematic of the topology for a tritium β decay, a valence-shell scatter in xenon, and an inner-shell scatter.

Because an electron's stopping power increases as the electron energy falls, the Auger cascade event topology leads to a more compact energy deposition than a β decay of the same total energy. In technologies such as the bubble chamber, where ER/NR discrimination relies on differences in energy density, this has an orders-of-magnitude impact on discrimination [16], but how this may impact discrimination in an LXe-TPC is less obvious. The dominant effect driving ER/NR discrimination in LXe-TPCs is the primordial ionization to excitation ratio [17,18], but the increased ionization density in an Auger cascade could lead to increased recombination, reducing S2/S1 and causing the event to appear more NR-like. In this work, we show that there is a measurable decrease in the ionization-to-scintillation signature of events with inner-shell vacancies.

B. Measurement concept: ^{127}Xe electron capture decay

The effect of an inner-shell vacancy on recombination (and on the resulting S2/S1 distribution and discrimination power) can be measured by observing ^{127}Xe electron-capture decays in the LXe-TPC. In this decay process, an s-orbital electron is captured by the nucleus, and an electron neutrino is emitted. If the daughter nucleus is in the ground state, the only visible energy in this decay comes from the resulting atomic vacancy. In the case of ^{127}Xe, the daughter ^{127}I nucleus is left in an excited 375 keV (203 keV) state with 47% (53%) branching ratio.

In a measurement of ^{127}Xe decays by the LUX experiment, the gamma rays emitted by the excited ^{127}I nucleus were contained in the detector, creating multisite

TABLE I. ^{127}Xe decay scheme and gamma escape probabilities. Values in the middle columns give the probability that the decay captures on a given shell (denoted by the columns) and follows the given gamma cascade (denoted by row). The final column gives the probability that the only energy deposited within 2.0 cm of the interaction site comes from the atomic vacancy, i.e., that no internal conversion electrons are emitted and all emitted gammas travel at least 2.0 cm from the decay site without scattering. Roughly 2% of all ^{127}Xe decays result in an L-shell vacancy with no other visible energy deposition, primarily from the two bold entries in the table. Data taken from Refs. [14,20–22].

Gamma cascade	Intensity by shell (%)			P_{esc}
(keV)	K (33.2 keV)	L (5.2 keV)	M (1.1 keV)	(% @ 2.0 cm)
618.4	0.0	0.1	0.0	62.2
375.0	15.9	**2.6**	0.6	45.0
172.1 + 202.9	22.3	3.6	0.8	0.5
172.1 + 145.3 + 57.6	1.4	0.2	0.0	0.0
202.9	42.0	**6.0**	1.4	10.6
145.3 + 57.6	2.6	0.4	0.1	0.0

interactions where the S2, but not S1, generated by the relaxation of the atomic vacancy could be isolated. LUX found the charge yield (S2) for L-shell captures to be lower than expected based on calibrations with tritium β decays [19], but low rate and the lack of S1 limited the significance of the measurement. The objective of this work is to observe ^{127}Xe decays in a small LXe-TPC where the ^{127}I gammas can escape, collecting events where the sole energy deposition comes from the relaxation of the atomic vacancy (see Table I). The resulting, high-statistics measurements of the S1 and S2 generated by L-shell vacancies can then be directly compared to tritium β decays measured in the same detector.

II. EXPERIMENTAL SETUP

A. XELDA detector

The Xenon Electron-Recoil L-shell Discrimination Analyzer (XELDA) detector is a 6.33-cm-diameter LXe-TPC with a 1.27 cm tall active region, giving a 40 mL (177 g) LXe target (see Fig. 2). The XELDA design is based on the MiX detector at the University of Michigan [23]. Scintillation in the chamber is detected by five photomultiplier tubes (PMTs): a single, 3-in.-diameter Hamamatsu R11410 at the bottom of the detector to measure the prompt S1 scintillation and four 1-in.-square Hamamatsu R8520 PMTs at the top to measure S2 electroluminescence. The S2 hit pattern in the four top PMTs allows millimeter-resolution reconstruction of the (x, y) position of events in the middle of the detector and robust rejection of events occurring outside the central volume.

The main TPC structure is fabricated from white polytetrafluoroethylene (PTFE, or Teflon) to increase light
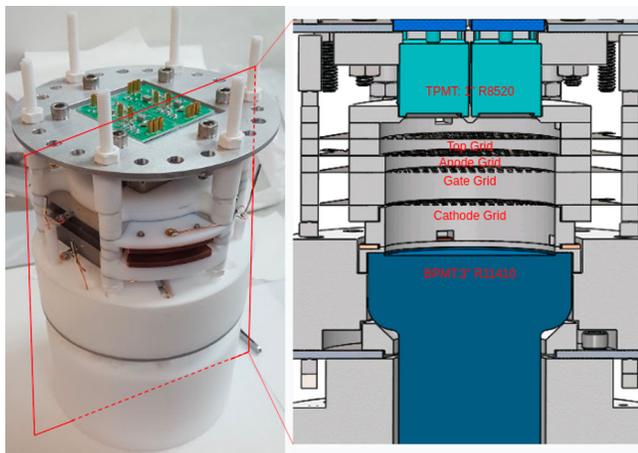

FIG. 2. Image (left) and solid model cross section (right) of the XELDA TPC. In normal operation, the liquid surface sits midway between the anode and gate grids. The liquid level is measured at three locations via the capacitance between pairs of copper plates—the bottom plate of one pair is visible in the center of the image on the left.

collection efficiency. Electric fields in the chamber are established by four electroformed, stainless steel, honeycomb-pattern grids. From bottom to top, these grids are the cathode, gate, anode, and top. The LXe-TPC target lies between the cathode and gate grids, and the liquid xenon surface lies between the gate and anode. Three parallel plate capacitors in the gate and anode grid planes measure the position of the liquid surface. Ionization electrons in the target volume drift upward past the gate grid to the liquid surface and into the gas phase, producing electroluminescence as they traverse the gap between the liquid surface and the anode grid. The anode and gate grids are aligned such that each hexagonal hole in the gate grid lies beneath a vertex in the anode grid, minimizing the spread in gas path lengths of drifting electrons. The top grid establishes a reverse-field region in the gas space above the anode, and a copper ring below the cathode collects ionization electrons produced in the reverse-field region between the cathode and bottom PMT.

The grids each have 2 mm pitch (distance between opposite sides of each hexagonal cell), and the grids in the liquid (gas) are 50 μm (120 μm) thick. The electroforming process gives the grid wires a hexagonal cross section, with wires roughly as wide as they are thick. The wire cross sections for each of the two grid types were measured under a microscope, and the as-fabricated grids were modeled in COMSOL to find the electric field leakage across each grid. Two electric field configurations are used in the data presented below, giving calculated TPC drift fields of 363 ± 14 V/cm and 258 ± 13 V/cm, with the dominant uncertainty coming from the position of the liquid surface. The gas extraction field is kept at 9.5 kV/cm in both field configurations. The electron drift times for interactions at the cathode (gate) grids are 8.5(1.0) μs.

B. Cryogenics and circulation system

The XELDA TPC hangs inside a 6-in.-diameter stainless steel vessel, which in turn is housed inside a vacuum cryostat repurposed from the SCENE experiment [24]. A xenon circulation system delivers liquid xenon to the bottom of the TPC vessel and removes gaseous xenon from the top, maintaining both the purity of the liquid xenon and the thermodynamic state of the chamber; see Fig. 3. Both "active" and "passive" circulation modes are needed to operate the chamber. During active circulation, an all-metal bellows pump (Senior Aerospace MB-111) draws xenon gas from the TPC and drives it through a heated zirconium getter (SAES MonoTorr PS4-MT3-R-1), back to a condenser held at 167.0 K by a cryocooler (Cryomech PT60) and proportional-integral-derivative-controlled heater. The condensed xenon drips down to the inlet at the bottom of the TPC. Xenon flow in this mode is throttled at 3 slpm by a mass flow controller at the pump inlet, and a heat exchanger between the gas streams leaving and returning to the cryostat keeps the heat load well within the PT60's cooling
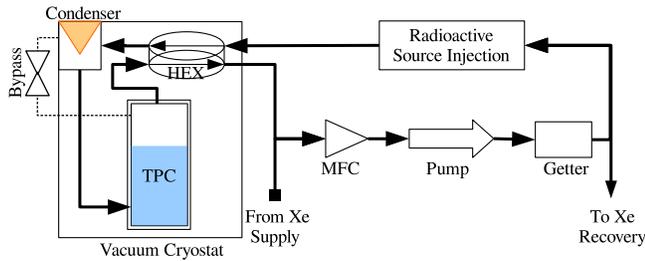
FIG. 3. Simplified piping and instrumentation diagram for the XELDA xenon circulation system. The system can operate either in an active circulation mode, where a pump drives xenon through the purification and source injection systems, or in a passive circulation mode, where gas from the TPC is drawn through a bypass to a cold head where it condenses and drips back to the inlet at the bottom of the TPC.

power. Circulation in this mode can increase electron lifetimes in the TPC from 1 to $> 10 \ \mu s$ in less than a week of circulation, limited primarily by inefficient mixing of xenon into the TPC.

In active circulation, the liquid level in the TPC fluctuates widely, driven by variations in the differential pressure between the TPC and condenser volumes. To achieve the stable liquid level needed for TPC operation, the pump is turned off, and a bypass valve connecting the TPC gas space directly to the condenser is opened. In this mode, the liquid xenon level in both the TPC and condenser drain line is determined entirely by the total xenon mass in the system and the xenon liquid temperature (density). Xenon is not efficiently purified in this mode (the TPC is able to maintain $> 10 \ \mu s$ electron lifetimes without active purification), but the TPC is still cooled by the incoming xenon as the condenser and TPC effectively form a closed-loop thermosyphon. In both active and passive circulation modes, the TPC temperature is regulated by a proportional-integral-derivative-controlled heater at the bottom of the TPC vessel.

C. Source production and injection

The measurements made with the XELDA detector require radioactive sources to be uniformly distributed within the TPC volume. These sources are injected directly into the gas stream during active circulation, downstream of the heated getter. Tritium was injected in the form of tritiated methane ($CH_3 T$) and was later removed from the system by resuming circulation through the heated getter. The ^{127}Xe sources were created at Fermilab by activating 25 and 50 mL high-pressure bottles of xenon in Fermilab's Neutron Irradiation Facility [25,26]. This activation also produces ^{129m}Xe (8.9 day half-life), ^{131m}Xe (11.8 day half-life), and ^{133}Xe (5.2 day half-life), but ^{127}Xe has the longest half-life of the activation products at 36.3 days and after a ~50-day-long cooldown period is the dominant radioactive isotope in the sample. We cannot remove the ^{127}Xe once
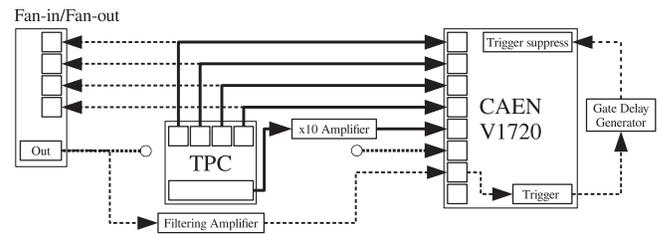

FIG. 4. Schematic showing the XELDA digitization and trigger electronics. Signal paths used in analysis are shown by solid lines, trigger paths are shown by dashed lines, and signals used for diagnostics only are shown by dotted lines.

injected, so this month-long half-life also dictates much of the run plan described below.

D. Electronics and trigger scheme

The XELDA data acquisition (DAQ) system deploys a CAEN V1720, eight-channel, 12 bit, 250 MHz digitizer. The bottom PMT signal passes through a ×10 amplifier immediately outside the vacuum jacket, and the amplified signal is digitized, ensuring a clear single-photoelectron signal and easy photon counting in small S1 pulses. This amplified signal saturates the digitizer for most S2 pulses in the region of interest, so S2 measurements are made using only the top four PMTs. Signals from the top PMTs are cloned by a nuclear instrumentation module fan-in/fan-out module, with one copy of each PMT signal digitized directly and a second copy going to trigger electronics. A schematic of the data acquisition electronics is shown in Fig. 4.

The XELDA DAQ operates primarily on an S2 trigger, with waveforms recorded for 15 μs before and after the trigger signal. This 30 μs window defines a single "event" in the TPC. To generate the trigger, a second fan in/fan out creates an analog sum of the top four PMTs, and that sum is fed to an Ortec 579 amplifier with 0.5 μs integration and differentiation times. The CAEN digitizer self-triggers when the amplified sum of the top PMT waveforms exceeds a fixed threshold for 300 ns. This trigger is highly efficient for S2 pulses down to five extracted electrons and has nonzero efficiency for triggering on single-electron pulses, but it is blind to S1s from events up to and beyond the K-shell energy. This trigger configuration allows the digitizer to ignore S1-only events from interactions in the reverse-field region below the cathode and in the xenon spaces outside the TPC field cage. Finally, a gate-delay generator enforces a 250 μs dead time after each trigger, preventing triggers on after-pulsing following large S2 pulses.

III. DATA COLLECTION, REDUCTION, AND CROSS-CALIBRATION

A. Summary of data collected

The XELDA detector acquired data over a period from August 2018 to October 2019. Data were taken with two
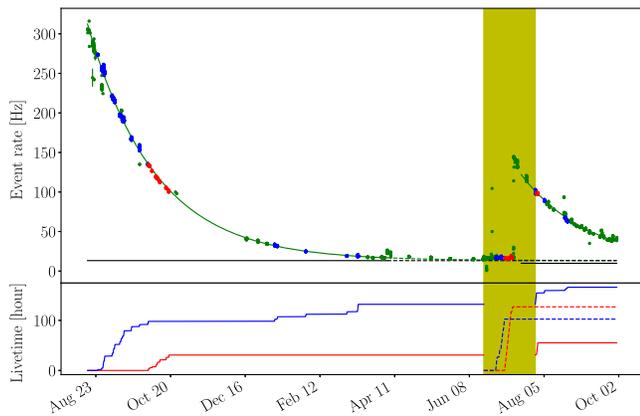
FIG. 5. Trigger rates (above) and accumulated live times (below) for data taken by XELDA from August 2018 through October 2019. Above: blue and red points in the rate plot indicate data used in this paper at 363 and 258 V/cm, respectively. Green points indicate datasets not used in analysis but shown here to illustrate the decay of ^{127}Xe in the detector. These runs were rejected due to effects including poor electron lifetime ($< 10 \mu s$), unstable liquid level, lack of clear single- and double-electron peaks in the S2 spectrum, or abnormal trigger rate. Lines indicate fits to the data rate taking a flat background plus decaying exponential with fixed 36.3 day half-life. The shaded band indicates the period where tritium is present in the detector. In the bottom panel, solid (dashed) lines indicate xenon-only (tritium) live times.

electric field configurations (drift fields of 363 and 258 V/cm) and two source configurations: "xenon-only" data, with trigger rates from ^{127}Xe of 10–300 Hz, and "tritium" data, which include some residual ^{127}Xe activity as well as a ~ 2 Hz trigger rate from tritium decay. Figure 5 shows the rates, calendar durations, and accumulated live times for each of the four data categories. There are three source injections relevant to this analysis: an initial ^{127}Xe injection, a tritium injection once the ^{127}Xe rate had decayed, and a final ^{127}Xe injection after tritium data taking was complete.

Data were collected using a customized version of the DAQMAN software [27], which also provided the framework for data reduction described in the next section. We do not process events where the signal from any top PMT exceeded the digitization range, removing all ^{127}Xe events where the associated gammas deposited significant energy in the TPC forward-field regions, as well as ^{127}Xe K-shell capture events near the walls of the detector where the S2 signal falls predominantly in a single channel. ^{127}Xe K-shell captures in the fiducial region of the detector (defined below) and all outer shell ^{127}Xe capture events are unaffected by this data preselection.

B. Pulse identification and classification

Offline waveform processing breaks each event into a sequence of S1-like and S2-like scintillation pulses. This is

accomplished by scaling each single-PMT waveform to units of vacuum ultraviolet-photons-detected (phd) per time, suppressing baseline noise in each waveform, summing the five waveforms, identifying pulses within the summed waveform, calculating a set of reduced quantities (RQs) for each pulse, and finally using these RQs to classify each pulse as "S1," "S2," or "other." The methods used here were drawn in part from past work on similar scale TPCs including Refs. [24,28]. This section gives key technical details for each of these steps.

The single-photo-electron (phe) response of each PMT is measured at the start of each data-taking session using a blue light-emitting diode (LED) coupled by a fiber optic to the TPC, tracking both excursions and general trends in phe size (or gain) in each PMT. Unlike the LED's 470 nm photons, 175 nm xenon scintillation photons frequently liberate two photoelectrons, so a fixed-value scale factor is applied to the LED data to give the response per vacuum ultraviolet phd. This correction is based on measurements by the LUX-ZEPLIN (LZ) Collaboration of typical double-photoelectron emission rates in the R8520 and R11410 PMT models used in XELDA [29], which we take to be 20.5% and 22.5%, respectively.

The first 3 μs of each PMTs waveform is pulse free for reconstructable events and sets the initial baseline for that waveform. The baseline for the remainder of the waveform is built from a rolling 160 ns average of the waveform, until encountering a pulse that differs from the rolling baseline by seven times the rms found in the initial 3 μs window. The rolling baseline calculation resumes when the waveform again falls within 1× the rms of the prepulse baseline, and the baseline during the excursion is taken to be the linear interpolation between the baselines found on either side of the pulse. Once the baseline is found for the entire 30 μs waveform, the baseline is subtracted, and all points not within 120 ns of an excursion beyond a fixed threshold are suppressed. The chosen suppression threshold keeps > 92% of single-photoelectron pulses.

The scaled, baseline-suppressed signals from all five PMTs are summed, and pulses are found in the summed waveform using a combination of 300 ns and 1 μs top-hat filters. Time windows containing S2-like pulses are found first by subtracting from the 1-μs-filtered signal the sliding-window maximum of the 300-ns-filtered signal, i.e., looking for 1 μs windows that contain significant pulses even after the largest 300 ns pulse within the window is excluded. After the edges of such S2-like pulses are found, S1-like pulses are identified in the remaining regions of the waveform using the 300-ns-filtered signal only.

Reduced quantities are calculated for each pulse found, including the pulse area (phd) by channel, the start time of the pulse (time to reach 5% of total pulse area in the summed waveform), the prompt fractions of the pulse (pfXXX refers to the fraction of pulse area, in the summed channel, reached within XXX ns of the pulse start time), and

the width of the pulse (`fwYYZZ` refers to the time it takes the summed waveform to grow from `YY`% to `ZZ`% of the total pulse area). Based on comprehensive hand scanning, a pulse is classified as S1 if `pf200` > 83% and `fw1050` < 150 ns and classified as S2 if `pf200` < 83% and `fw1050` is between 60 ns and 1.4 μs. All other pulses are classified as other.

C. Event reconstruction

Event reconstruction in XELDA is optimized for the accurate reconstruction of single-scatter events in the fiducial volume of the TPC (defined below) and for the removal of events that either are not single scatters or fall outside the fiducial volume. To start this process, the largest S2 pulse (determined from the total phd in the four top PMTs) is designated as the "main" S2, and the largest S1 pulse (determined from phd in the bottom PMT) prior to that S2 is designated as the main S1. Events with no S2 pulses or with no S1 pulses prior to the main S2 are discarded.

With the main S1 and S2 pulses identified, three-dimensional (3D) position reconstruction of the event is possible. The drift time separating S1 and S2 (based on the pulse start times defined in the previous section) gives the z position of the event, while the hit pattern of the S2 in the top four PMTs gives the (x, y) position. The (x, y) position is found through a maximum likelihood optimization given an (x, y)-dependent light response function (LRF) for each PMT. The LRFs are parametrized as in Ref. [30], with LRFs for the four PMTs differing only by 90° rotations and an overall scaling to account for tube-to-tube variations in quantum efficiency (QE). Both LRF parameters and relative QEs are found using L-shell capture events, selected based on S2 size and assumed to uniformly populate the TPC volume. Relative QEs are fixed by events at the centroid of the (x, y) distribution, assumed to have equal geometric light collection efficiency for all four top PMTs. LRF parameters are found iteratively, using reconstructed positions to generate a new LRF until the LRF parameters stabilize. Finally, two (adjacent) top PMT tubes show nonlinear response for ionization signals beginning around the K-shell capture peak. For the purposes of (x, y) reconstruction only, a quadratic correction determined from events at the centroid of the K-shell (x, y) distribution is applied to the measured signal size in the two saturating tubes. The accuracy and resolution of the final (x, y) reconstruction is evident in the reconstructed honeycomb pattern of the gate grid using K-shell events, shown in Fig. 6.

Two pulse-area corrections are applied during event reconstruction, both affecting S2 pulses only. An exponential electron lifetime correction, based on the observed drift-time dependence of uncorrected K-shell S2s, scales all S2 pulses to the zero-drift-time value. S2 signals in the top PMTs are also corrected for the relative QEs found in (x, y) reconstruction, with one tube receiving a unity correction.
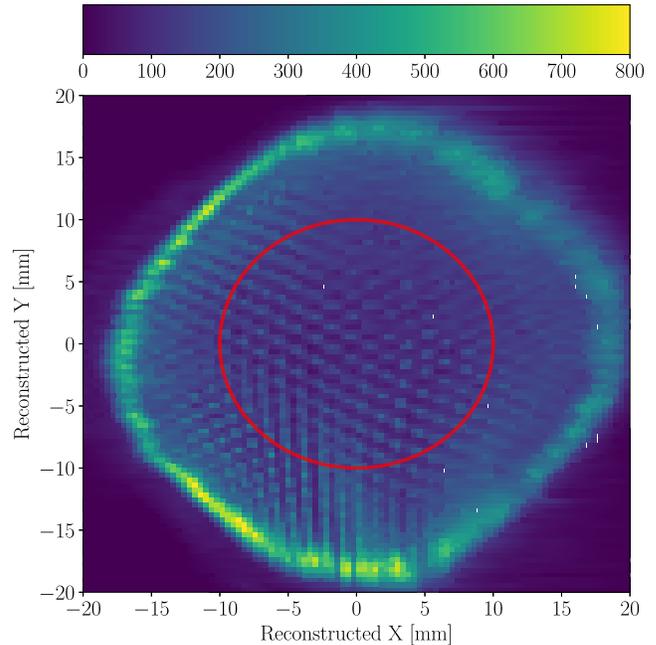

FIG. 6. Reconstruction of (x, y) position based on the S2 hit pattern in the top PMTs for K-shell events. The honeycomb pattern imprinted by the focusing of electrons as they drift past the gate grid is clearly visible inside the fiducial volume. The red circle indicates the radial extent of the fiducial region of the TPC (10 mm).

All references below to S2$_{\text{top}}$ and S2$_{\text{CH12}}$ include these corrections. No additional corrections (e.g., position-dependent corrections) were found to improve resolution in either S1$_{\text{bottom}}$ or S2$_{\text{top}}$ for events in the fiducial volume. PMT nonlinearity does affect S2$_{\text{top}}$ for K-shell capture events, but this is handled during the cross-calibration described in Sec. III D rather than as an explicit correction.

The final selection of "golden" single-scatter events is based on the following set of cuts. The fiducial volume is set first by requiring the drift time of the event to fall between 2.5–7.5 μs, avoiding regions with nonuniform electric field near the cathode and gate grids. Next, the reconstructed radial position of the event must be < 10 mm from the central axis of the detector, eliminating events near the walls of the detector where drifting charge may be lost and (x, y) reconstruction is challenging. Events must fall within the region of interest in (S1$_{\text{bottom}}$, S2$_{\text{top}}$) space, defined as 2.0–500 phd in S1$_{\text{bottom}}$ and 500–70000 phd in S2$_{\text{top}}$. No pulses may fall in the first 3 μs of the event, and the rms of each PMT waveform in that 3 μs must be < 1.5 analog-to-digital converter counts, eliminating events with poorly characterized baselines. Multiple-scatter events in the TPC are eliminated by requiring events to contain no pulse larger than 10% of the main S2. Events with misidentified S1s (due to a particular sporadic electronic pickup problem) are eliminated by requiring S1$_{\text{top}}$/S1$_{\text{bottom}}$ < 3. For normal events, this ratio is < 0.1. Finally, various

anomalous event topologies are eliminated by requiring that no more than 20% of the total event area be contained in pulses outside the main S1 and S2, no more than 5% of the total event area be contained in pulses classified as other, and no pulses fall between the main S1 and S2 except for single-electron S2s (identified by pulse size and time profile) and S1s smaller than 10% of the main S1.

D. Cross-calibration of datasets

The long calendar duration of the XELDA data-taking campaign makes it essential to cross-calibrate data taken at different times and in different detector conditions. In particular, changes in liquid level due to fluctuations in total xenon payload affect electroluminescence production (S2) and potentially the light collection efficiency in the bottom PMT (S1). These and other phenomena affecting S1 and S2 yields are addressed by using the ^{127}Xe K-shell capture peak as a *common candle* in all XELDA data. At 33.2 keV, this peak sits above the tritium β-decay end point but is low enough in energy to avoid saturation effects in most (three of five) of the PMTs. We emphasize that we do not use the K-shell events to study any aspect of ER discrimination but instead use it solely as a common point of reference that is easily visible in both types of data that we collected.

The tritium data used in this analysis was taken in a single continuous period (see Fig. 5), and we see no evidence of variation in the K-shell peak position or in the tritium β continuum, in either the 363 or 258 V/cm tritium datasets. By contrast, both gradual and discrete changes in K-shell S1 and S2 were seen in the xenon-only data, so run-by-run (roughly day-by-day) S1$_{bottom}$ and S2$_{top}$ correction factors are applied to xenon-only data, aligning each run's K-shell peak with that observed in the tritium data at the corresponding drift field.

The dominant effect seen in the S1$_{bottom}$ amplitude is a gradual reduction in K-shell signal size, dropping 10% by the end of the first three months of data taking. This loss of signal appears to be a reduction in the bottom PMT's QE, consistent with aging studies performed on the same tube model at Brown [31] given the > 100 C of charge collected at the PMT anode during the initial high-rate period. This effect plateaus in later data, as the event rate falls and the PMT aging process slows. Correction factors for this effect are derived from a smoothing spline applied piecewise in calendar time to run-by-run Gaussian fit means of the K-shell S1 peak. The smoothing spline avoids overfitting to statistical fluctuations in the peak position, with breaks between splines inserted by hand to accommodate extended periods when the PMT was unbiased, slowing the aging process. The average absolute residual between the spline and individual S1$_{bottom}$ fit means in the ^{127}Xe data is 0.25%.

Correcting for variation in the S2$_{top}$ signal is more complicated because of the nonlinear response at K-shell energies in two of the top PMTs. Fortunately, we see no evidence for differences in S2 variation across individual tubes (consistent with the hypothesis that liquid-level fluctuations are the primary cause of S2 drift), so we use the signal in the unsaturated top PMT channels 1 and 2, denoted S2$_{CH12}$, to make the cross-calibration.

We first scale all xenon-only data to a representative xenon run at each drift field, then find the correction factor between the combined xenon-only K-shell peaks and the tritium K-shell peaks. Run-by-run xenon-only S2 corrections are made by selecting events whose S1 falls within the K-shell region of interest and then finding the S2 scale factor that best aligns the S2$_{CH12}$ distribution with a reference distribution, determined by maximizing the p value returned by a Kolmogorov-Smirnov (KS) test. The reference xenon-only S2$_{CH12}$ distribution is given by a high-statistics run taken prior to tritium injection, and the maximum correction required in the pretritium xenon-only runs is 8%. Xenon-only runs after tritium was removed exhibit much smaller S2s, requiring 20%–33% corrections, consistent with an expected decrease in S2 yield from the increased liquid level following source injections (see below).

Once the xenon-only runs are made self-consistent, the relative S2 scale between the xenon-only and tritium datasets is found by again maximizing the p value returned by a KS test. Because S2$_{CH12}$ exhibits a strong position dependence (unlike S2$_{top}$), the S2$_{CH12}$ K-shell spectrum in the fiducial volume is rather broad (see Fig. 7), and the KS test is influenced by an unmodeled gamma background that appears beneath the relatively low-rate K-shell peak in the tritium data. Therefore, a floating uniform background component is added to the xenon-only S2$_{CH12}$ distribution in order to match the observed tritium distribution. Optimized p values of 0.692 (0.623) are found for the 363 V/cm (258 V/cm) datasets, giving a relative scaling between the xenon and tritium data of 0.8235 (0.7945), similar to the drop in S2 seen in going from pretritium xenon-only data to post-tritium xenon-only data.

The changes in ionization gain found above are correlated with discrete changes in the liquid level in the TPC, which result from source injections and xenon recovery events. In the pretritium xenon data, the liquid height as measured by the capacitive level sensors was $23 \pm 5\%$ of the distance from the gate to anode. The liquid surface in the tritium and post-tritium xenon data was at $45 \pm 10\%$ of the gate-anode gap. This reduction in the height of the electroluminescence region (with the corresponding increase in the extraction field) gives an estimated ~15% reduction S2 gain [17], confirming the expectation that liquid-level shifts are the dominant contributor to the observed changes in S2 gain.

Liquid-level shifts also affect the TPC drift field, indicating that the pretritium xenon data had a ~20 V/cm lower drift field (at both field settings) than
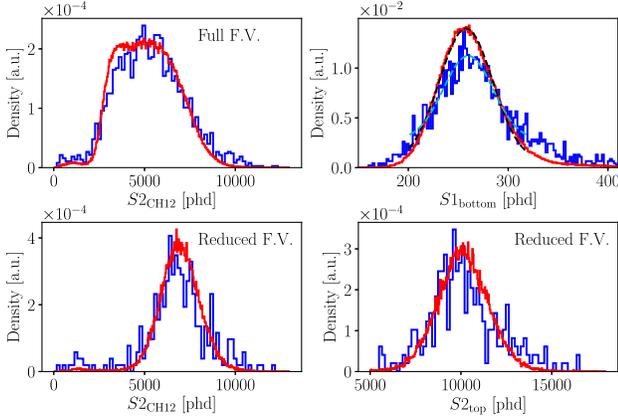
FIG. 7. Comparison of the K-shell common candles at 363 V/cm, reconstructed after cross-calibration between xenon-only and tritium datasets is complete. In each figure, the red histogram shows the xenon-only distribution, and the blue histogram shows the tritium distribution. Dashed curves show Gaussian fits to the $S1_{bottom}$ distributions. The top figures use the full fiducial volume, while the bottom two use a reduced volume that is offset from the center of the detector to avoid PMT saturation effects. The vertical axis is scaled such that all the histograms integrate to unity.

the corresponding tritium and post-tritium data. The K-shell common candle alignment described above implicitly corrects for this field shift, with the caveat that the L-shell peak exhibits slightly less field-dependence than is seen in the K shell. This is evident in the ratio of the K- to L-shell $S1_{bottom}$, which drops from 7.7 at 258 V/cm to 7.5 at 363 V/cm—that is, the K-shell events show more "field quenching" in S1 than L-shell events. By relying on the K-shell common candle, we overcorrect for the effect of any changes in the drift field on the L-shell signal, causing us to potentially underestimate the significance of the inner-shell vacancy effect.

Whether from to variation in ionization/scintillation gain or drift field, uncertainties in the S1 and S2 scale factors between the xenon and tritium datasets dominate the systematic uncertainty in the analysis that follows. The S2 uncertainty is estimated by bootstrapping from the $S2_{CH12}$ distribution and building a Neymann construction from the ratio of the global maximum p value described above to the p value obtained using simulation truth information. This construction finds that the true scale factor's p value is greater than 0.6 times the maximum p value in 68% of trials. Applying this condition to the scale factors tested in our data and marginalizing over the unknown background component gives a relative uncertainty of 1.2% (0.85%) for the 363 V/cm (258 V/cm) xenon-to-tritium S2 scale factor. A similar method is employed for the $S1_{bottom}$ channel, giving relative uncertainty of 0.84% (1.1%) in the S1 scale factor.

Figure 7 shows the xenon-only and tritium $S1_{bottom}$ and $S2_{CH12}$ K-shell distributions at 363 V/cm after the above

scaling is complete. As a cross-check, we also compare the full $S2_{top}$ distributions for events falling in a reduced fiducial region that selects events where most S2 light goes to the nonsaturating PMTs. In all cases, a consistent K-shell common candle is observed between the cross-calibrated xenon-only and tritium datasets.

IV. RESULTS

Figures 8 and 9 show the two-dimensional (2D) [S1, $\log_{10}(S2/S1)$] and (S1, S2) distributions of the 363 V/cm xenon-only and tritium datasets after all cuts and corrections described in the previous section are complete. Visually, it is clear that the L-shell peak location is offset from the centroid of the tritium ER band, but it is conceivable that this offset is due to the falling tritium
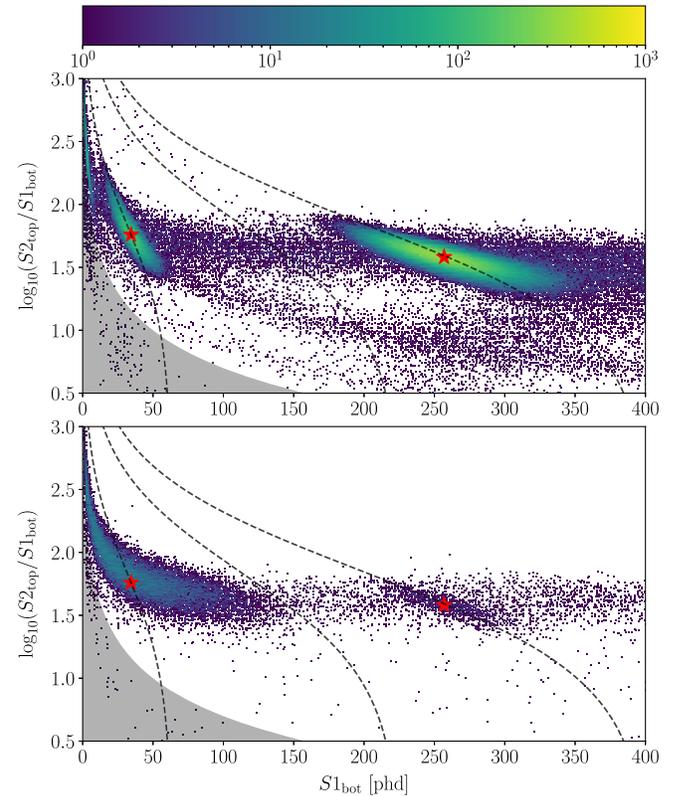

FIG. 8. Distribution of $\log_{10}(S2_{top}/S1_{bot})$ vs $S1_{bot}$ for the 363 V/cm Xe-only data (top) and ^{127}Xe plus tritium data (bottom). The dashed curves denote constant energies at 33.2 keV (K-shell), 18.6 keV (tritium Q value), and 5.2 keV (L shell). The color represents the absolute counts per bin, and the scale is logarithmic. The red stars indicate the centroids of the K- and L-shell distributions, as determined by the xenon-only dataset. Those points are reproduced in the tritium data to help guide the eye. The L-shell centroid is clearly offset below the center of the tritium band. The gray regions indicate the analysis threshold, and events in this region are excluded from the analysis.
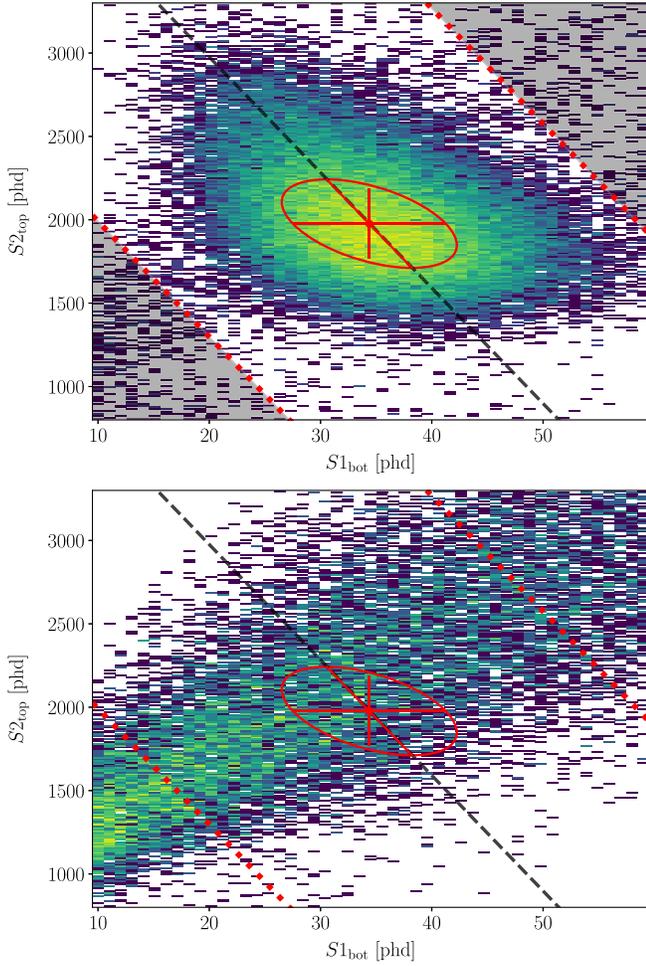
FIG. 9. (Top) the ^{127}Xe L-shell electron capture peak, as observed in the 363 V/cm data. The red dotted lines denote the limits of the energy region of interest (the shaded regions are not used in L-shell analysis), while the black dashed line denotes the 5.2 keV energy of this peak. The red point is the median in both $S1_{bottom}$ and $S2_{top}$ for events falling within the energy region of interest (3.20–7.20 keV), while the error bars shown denote the widths (σ) of the distribution decomposed into independent physical fluctuations: S1 fluctuations (horizontal), S2 fluctuations (vertical), and recombination fluctuations (diagonal). (Bottom) the ^{3}H β-decay continuum in the K-shell energy region of interest, as observed in the 363 V/cm data, with the same lines of constant energy. The 1σ widths of the L shell are drawn on top, again to illustrate the offset.

energy spectrum, since the tritium recombination fraction increases (S2/S1 falls) as energy increases.

A. Simulation and absolute gain calibration

To determine whether the observed offset indicates a true shift in S2/S1 response, we would ideally compare L-shell events to an equal-energy β decay. Since we do not have a monoenergetic β source, we use the Noble Element Simulation Technique (NEST) [17,32] software package to

bridge the gap. We first tune the NEST detector parameters to find agreement between our tritium data and a tritium simulation in NEST and then use NEST to generate monoenergetic 5.2 keV electron recoils.

To correctly simulate the statistical variation in signal size, NEST requires the gains of the S1 and S2 signals, defined as

$$S1 = g_1 n_\gamma, \qquad S2 = g_2 n_e, \qquad (1)$$

where n_γ and n_e are the numbers of scintillation photons and ionization electrons extracted from the interaction site. We use the corrected $S1_{bottom}$ and $S2_{top}$ as S1 and S2 in Eq. (1), deriving a single g_1 (g_2) at each drift field representing the photons detected per photon (electron) emitted in the tritium dataset.

The gains g_1 and g_2 can also be used to construct a linear, drift-field-independent ER energy scale, given by [11]

$$E_{ee} = \left(\frac{S1}{g_1} + \frac{S2}{g_2}\right) \times 13.7 \text{ keV}, \qquad (2)$$

where the subscript on E_{ee} indicates the reconstructed electron-recoil equivalent energy.

We find g_1 and g_2 independently at each field by matching the observed tritium distribution to that simulated by NEST, performing a binned maximum likelihood fit (with Poisson statistics) on the 2D [S1, log(S2/S1)] distribution. The S1 window containing the L-shell peak (15–60 phd) is excluded from the fit to avoid contamination by L-shell events in the tritium dataset. We also use Eq. (2) to constrain g_1 and g_2 by requiring the L-shell centroid to reconstruct to 5.2 keV, leaving a single free parameter for the maximum likelihood fit. The best-fit g_1 and g_2 with the dominant systematic uncertainties are given in Table II, and the values obtained at each drift field agree to within the uncertainties. Figure 10 shows the reconstructed energy spectra based on these g_1 and g_2 values for tritium decays and K-, L-, and M-shell electron capture decays. A comparison of $\log(S2/S1)$ between the tritium data and the best-fit simulated tritium profile at 363 V/cm is shown in Fig. 11.

All subsequent references in this work to quantities simulated using NEST use these best-fit XELDA detector parameters. With these values in hand, we can also

TABLE II. Observed g_1 and g_2 values at the two drift fields considered in this analysis.

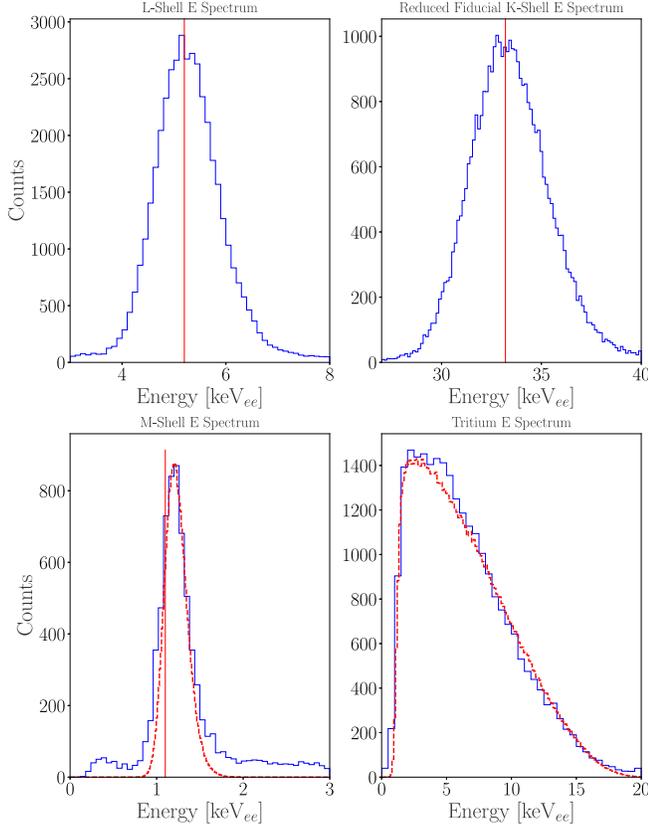| Drift field	g_1	g_2
(V/cm)	(phd/photon)	(phd/electron)
363	0.166 ± 0.001	11.5 ± 0.001
258	0.163 ± 0.002	11.5 ± 0.002

FIG. 10. Reconstructed energy distributions for the ^{127}Xe electron capture peaks and tritium continuum (bottom-right) from 363 V/cm data (blue). Vertical red lines denote the true energy of the K- (top right) and L-capture (top left) peaks (red). The NEST expectation for the tritium energy spectrum (red, bottom right) in XELDA, including our two-phd S1 threshold, is shown for comparison to the observed data. The NEST expectation for a 1.2 keV β, also with a two-phd S1 threshold applied, is overlaid on the M-shell peak (bottom left). The K-shell peak is shown for a reduced fiducial volume where saturation effects are not present.

transform to "physical" axes, by dividing out the detector-specific gains. In this case, the discrimination parameter becomes $\log(n_e/n_\gamma)$ rather than $\log(S2/S1)$.

We next find the light and charge yields of L-shell electron-capture events and compare to the yields found for monoenergetic β decays simulated in NEST. We define these quantities as

$$\mathcal{Q}_y = \left\langle \frac{S2}{g_2 E_{ee}} \right\rangle, \qquad \mathcal{L}_y = \left\langle \frac{S1}{g_1 E_{ee}} \right\rangle, \qquad (3)$$

where $\langle \cdots \rangle$ indicates an average taken over a fixed window in reconstructed energy, 3.2 keV $< E_{ee} <$ 7.2 keV. We also define the ratio

$$q = \frac{\mathcal{Q}_y^{\mathrm{L}}}{\mathcal{Q}_y^{^3\mathrm{H}}}, \qquad (4)$$

where the denominator gives the tritium charge yield averaged over the same reconstructed energy window. The advantage of the ratio q is that the g_1 and g_2 factors enter only via their effect on the energy window and on signal resolution in the NEST simulation, significantly reducing systematic uncertainty.

Table III shows the charge yields, light yields, and q ratios observed in XELDA data and NEST simulation at both drift fields. Uncertainties on the NEST values are driven by uncertainty on the drift field strength. Systematic uncertainties on the values derived from XELDA data include uncertainties from the cross-calibration of the xenon-only and tritium datasets as well as, for $\mathcal{Q}_y$ and $\mathcal{L}_y$, uncertainties on the gains g_1 and g_2. The q ratios show a 6.9 (9.2)-σ discrepancy at 363(258) V/cm between the L-shell response and the response expected for a β decay.

B. Recombination model for L-shell capture

Our hypothesis is that the reduced charge yield seen in L-shell capture events is due to increased recombination at the event site, and we model this effect by modifying the recombination physics in the NEST package. We float the mean recombination fraction for the simulated monoenergetic "β," leaving the recombination fluctuation model unchanged, and perform a binned maximum-likelihood fit to the observed $\log(n_e/n_\gamma)$ distribution for L-shell events. The unmodified and best-fit $\log(n_e/n_\gamma)$ distributions are shown with XELDA data in Fig. 12. The unmodified and best-fit mean recombination fractions (r) are listed in Table III, along with the simulated charge and light yields using the modified recombination model.

V. DISCUSSION

A. Discrimination of L-shell capture events

A key aspect to the success of LXe-TPCs for dark matter is the efficient rejection of ER events using the S2/S1 (or n_e/n_γ) ratio, and a figure of merit for the discrimination power is the fraction of ER events that fall below the NR median. We use NEST to predict the NR median in XELDA, shown by the vertical dashed green line in Fig. 13, and to extrapolate the low-n_e/n_γ tails of the tritium and ^{127}Xe L-shell spectra, shown by the dashed curves. The extrapolation is necessary as the observed tails in the XELDA tritium and L-shell data are dominated by low-energy external ER backgrounds that are unmodeled in our NEST simulations and that show anomalously poor discrimination due to detector effects arising in any small chamber, particularly wall effects and multisite events that produce simultaneous energy depositions in the TPC and reverse-field region(s) where the ionization signal is suppressed, also known as $\gamma - X$ events. Our tuned NEST model provides a convenient way to extrapolate spectra to regions that we cannot measure directly due to these backgrounds. The extrapolated tritium spectrum and NR median
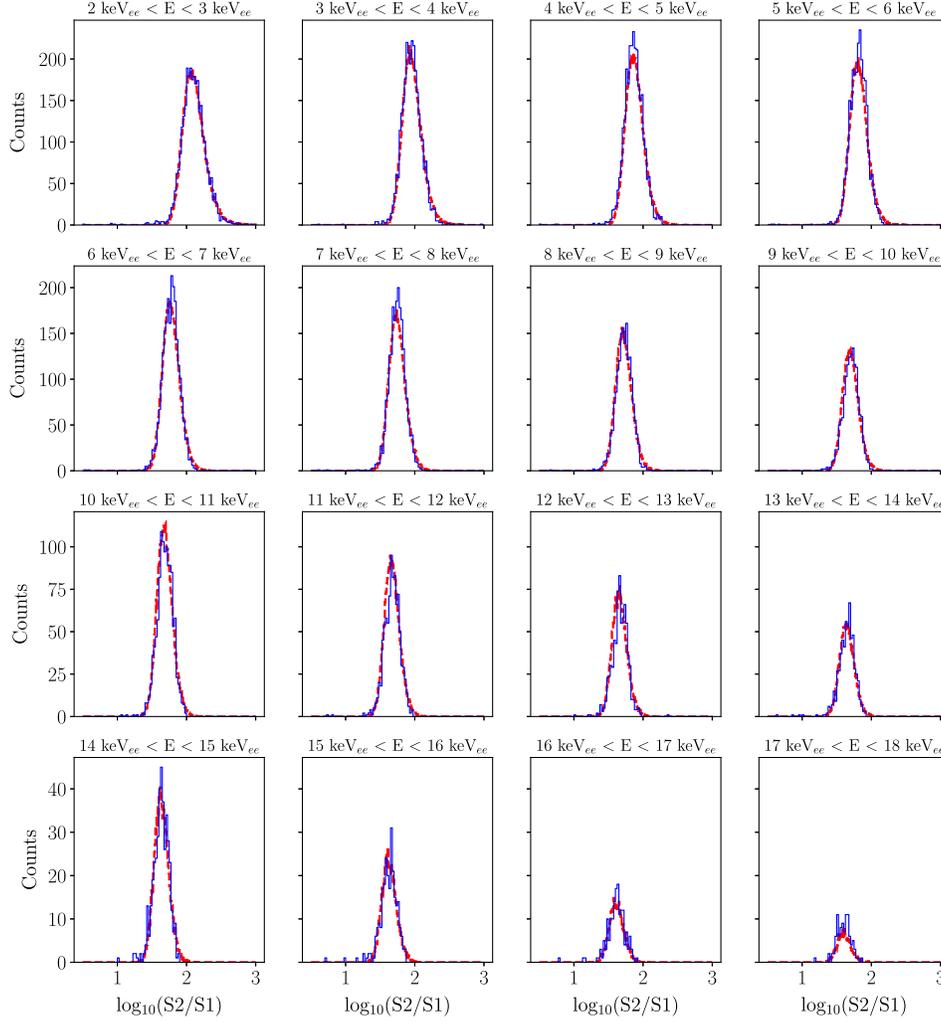
FIG. 11. Distributions of log(S2/S1) for the tritium β-decay continuum, in 1 keV energy bins, comparing the observed tritium data (blue) to expected distributions from a NEST simulation (red) at 363 V/cm. Threshold effects become relevant for signals below 2 keV, so those bins are excluded from this analysis. Above this energy, the observed and simulated distributions are in good agreement. The slight excess in data in the 3–8 keV bins is caused by residual ^{127}Xe L-shell activity in the tritium dataset.

TABLE III. The values of $\mathcal{Q}_y$, $\mathcal{L}_y$, and q for the ^{127}Xe L-shell electron capture as measured by XELDA at both 363 and 258 V/cm. Also shown are the values for a 5.2 keV monoenergetic β and a 5.2 keV β with modified recombination, simulated by NEST with the corresponding recombination probability r as indicated. Uncertainties in the underlying NEST tritium model are not included in the systematic uncertainties quoted above and should be considered in the application of $\mathcal{Q}_y$ and $\mathcal{L}_y$ as described in Sec. V D. The unitless ratio q is robust against mismodeling in NEST—reanalysis with different NEST tritium models gives a 0.3% shift in q for ^{127}Xe $- L$ and 0.06% shifts in q for the simulated β's, subdominant to the experimental systematic uncertainty.

|	Source	$\mathcal{Q}_y$ (e/keV$_{ee}$)	$\mathcal{L}_y$ (γ/keV$_{ee}$)	q
XELDA	^{127}Xe $- L$, 363 V/cm	33.63 ± 0.03 (stat) ±0.33 (sys)	39.36 ± 0.33 (stat) ±0.36 (sys)	0.917 ± 0.001 (stat) ±0.009 (sys)
NEST	β, $r = 0.4789$	36.42 ± 0.14	36.44 ± 0.13	0.9766 ± 0.0006
NEST-mod	β, $r = 0.5196$	32.98 ± 0.02	40.02 ± 0.02	0.901 ± 0.004
XELDA	^{127}Xe $- L$, 258 V/cm	32.87 ± 0.07 (stat) ±0.37 (sys)	40.12 ± 0.07 (stat) ±0.37 (sys)	0.909 ± 0.003 (stat) ±0.007 (sys)
NEST	β, $r = 0.4984$	35.10 ± 0.23	37.90 ± 0.23	0.9753 ± 0.0005
NEST-mod	β, $r = 0.5319$	32.16 ± 0.03	40.83 ± 0.03	0.911 ± 0.006
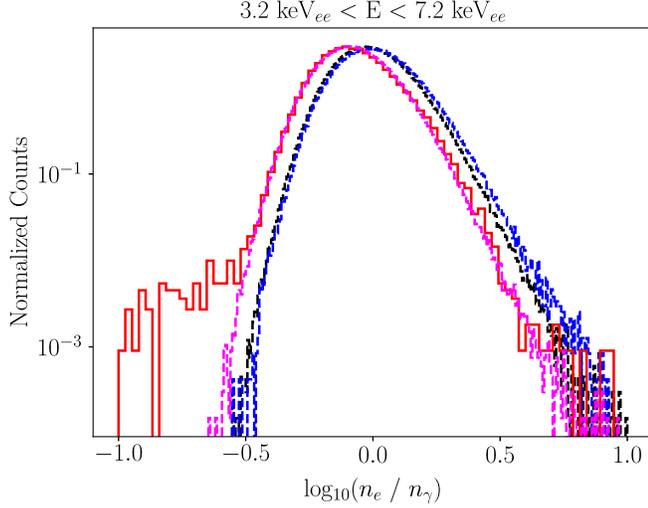
FIG. 12. Distributions of discrimination parameters for events within the 3.2–7.2 keV$_{ee}$ region of interest at 363 V/cm. Curves include the ^{127}Xe L shell as measured by XELDA (red), the NEST expectations for the tritium continuum (blue), a monoenergetic β of 5.2 keV (black), and a monoenergetic 5.2 keV beta with our modified recombination fraction (magenta). The L-shell data, as well as our modified recombination beta, shows a clear downward shift in n_e/n_γ when compared to NEST expectations.
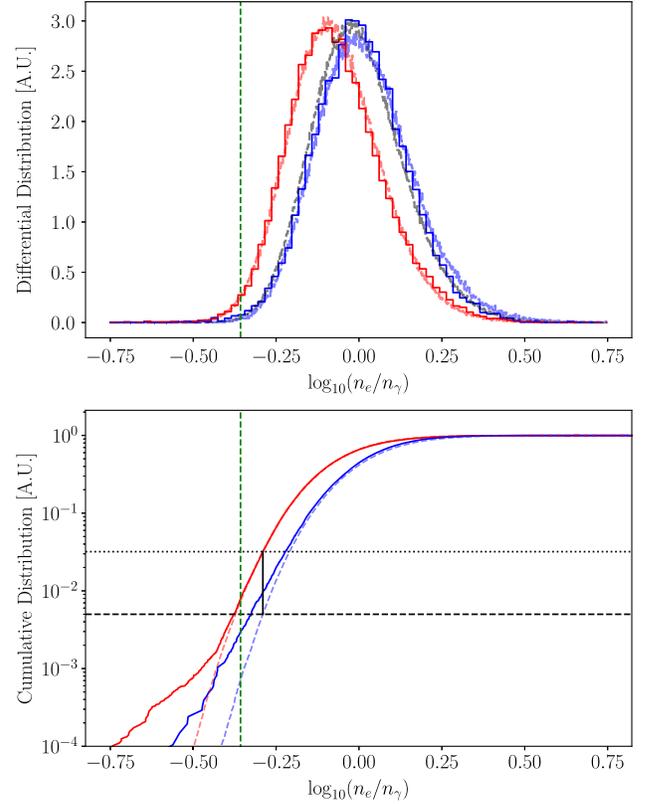

FIG. 13. Differential (top) and cumulative (bottom) distributions of $\log_{10}(n_e/n_\gamma)$ for 363 V/cm L-shell (solid red) and tritium (solid blue) events in the 3.2–7.2 keV region of interest. The dashed curves show the NEST expectations for a 5.2 keV β (black, top only), the tritium continuum in this region of interest (blue), and the modified-recombination 5.2 keV β (red). Also shown is the expected median of the NR distribution (green dashed). The horizontal dashed and dotted lines in the lower plot show the nominal 0.5% inefficiency predicted by NEST and 1.65% inefficiency from the L-shell data, respectively, at the cut choice indicated by the vertical black line. Data diverge from NEST expectations at low acceptance values because the NEST simulation does not include background γ-X events, wall leakage, nor other effects that impact discrimination in small detectors.

give a discrimination power of 7×10^{-4} in XELDA (for the 3.2–7.2 keV window), comparable to that achieved by the LUX Collaboration [12], and an order of magnitude lower than the leakage observed for L-shell events. We note that the modified recombination β model accurately reproduces the observed L-shell leakage for values of $\log_{10}(n_e/n_\gamma)$ above the NR median. All of these cases are illustrated for the 363 V/cm data in the bottom panel of Fig. 13.

For a point of comparison where the effect of the unmodeled backgrounds is reduced, we use the value of $\log_{10}(n_e/n_\gamma)$ at which NEST would predict an ER leakage fraction of 0.005 for tritium data. At this value of $\log_{10}(n_e/n_\gamma)$, the L-shell leakage fraction for 363 V/cm data is 3.2 times higher than that observed for tritium. Alternatively, a direct comparison to the predicted tritium leakage from NEST without the unmodeled backgrounds finds that the observed L-shell rejection inefficiency is a factor of 6.4(6.2) times higher than expected at 363(258) V/cm. In either case, the L-shell leakage is significantly worse than what would be expected based on a tritium calibration alone.

For another look at the change in discrimination power, Fig. 14 shows the observed leakage in both tritium and L-shell data as a function of the predicted tritium leakage fraction from NEST, for both drift fields. The value of the $\log_{10}(n_e/n_\gamma)$ cut that leads to the predicted NEST tritium leakage rate is shown on the secondary (upper) x axis. As in Fig. 13, unmodeled backgrounds induce deviations

from NEST predictions in the tails of the observed distributions, but for cut inefficiencies greater than $\sim 1\%$, the data are robust against the effects of these tails. We can predict L-shell leakage below 1% using the modified recombination model, which shows a roughly constant increase in leakage relative to tritium regardless of cut position. In all scenarios, in both data and simulation, L-shell events show significantly greater leakage than is seen for tritium.

B. Extending the model to L-shell scatters

To relate the observed L-shell electron capture signal to scatters off L-shell electrons, we must model the contribution to the signal of the ejected electron itself. We expect
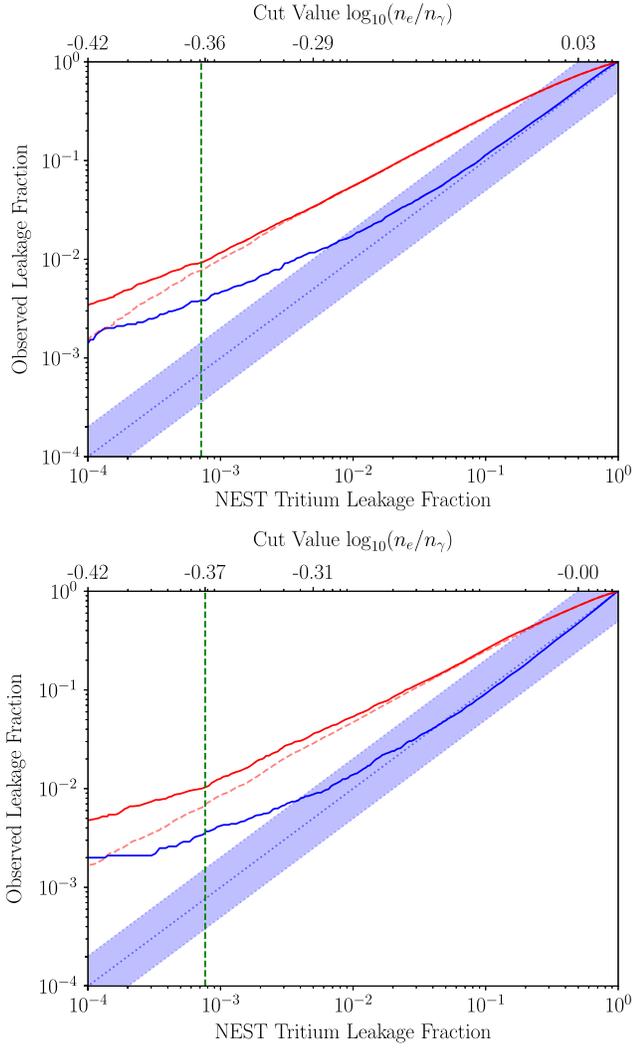
FIG. 14. The observed leakage fractions of tritium (blue) and ^{127}Xe L shell (red) for an energy region of 3.2–7.2 keV as a function of NEST's predicted tritium leakage fraction for the 363 (top) and 258 V/cm (bottom) datasets. The blue band shows a factor of 2 around $y = x$ (perfect agreement with NEST), while the red dashed line indicates NEST's simulation of the modified-recombination β model. The NR median from NEST is indicated by the vertical green dashed line. The secondary x axis is the value of the cut in $\log_{10}(n_e/n_\gamma)$ that gives rise to the expected and observed leakage fractions.

an L-shell scatter event to have the same charge yield as the ^{127}Xe L-shell decay when $E_{\text{recoil}} = 5.2$ keV, i.e., when the ejected electron carries no energy. As the energy of the recoil increases, however, more of that energy will be deposited by the ejected electron, and the event will approach the same charge yield as a valence scatter. We model this by taking an energy-weighted average of the L-shell recombination mean (r_L, given in Table III) and the recombination mean of a valence shell recoil of energy E (r_E, from Ref. [17] or [33]). In this way, the modified inner-shell ER recombination is

$$r'(E) = \frac{E_L r_L + (E - E_L) r_E}{E}, \quad (5)$$

where E is the total deposited energy and E_L is the binding energy of an L-shell electron (5.2 keV). The recombination mean obtained from this average may be less than r_E, in which case the valence recombination mean should be used instead, giving

$$r(E) = \begin{cases} r'(E) & \text{for } r_E \leq r' \\ r_E & \text{for } r_E > r' \end{cases}. \quad (6)$$

This model can be applied to the expected solar neutrino ER spectrum [7,34], where our corrected recombination factor is applied to 8/52 of events between L- and K-shell energies (or to the eight L-shell electrons out of 52 xenon electrons with binding energies in the region of interest). All other events are generated using the nominal recombination fraction in NEST. In the LZ region of interest (1.5–15 keV$_{\text{ee}}$), this model predicts a 4.8% increase in the number of leaked ER events from solar neutrinos over what would be expected from applying the β model to the neutrino ER background. The excess events are almost entirely in the 5.2–8.0 keV$_{\text{ee}}$ region, where this model predicts a leakage increase of 7.9%. As NR acceptance goes down, this relative increase in overall leakage stays fairly constant, down to at least the 35th percentile of the NR band.

C. Impact on dark matter sensitivity

To investigate the impact of this effect on detector sensitivity and dark matter limits, we consider both two-sided profile likelihood ratio (PLR) tests and optimized cut-and-count analyses in three different scenarios, for exposures up to 3 kton-yr. All three scenarios include backgrounds from neutrino sources only, including ER and NR events from electron- and nuclear-scattering of solar, atmospheric, and supernova neutrinos [7,34,35]. The first two scenarios compare dark matter sensitivity with and without the increased L-shell ER leakage, using Eq. (6) and the standard ER recombination model, respectively, to both simulate data and create the background model input to the limit-setting analysis. The third scenario considers the effect of mismodeling the L-shell background, simulating data based on the modified ER recombination model but using the standard ER recombination expectations to build the background model for limit setting.

Both the PLR and cut-and-count analyses show a negligible difference in sensitivity between the first two scenarios, regardless of exposure. This is not surprising—while the increase in leakage for L-shell capture events on their own is striking, relatively few neutrino events both scatter on the L-shell and fall in the narrow energy window where the effect is significant. The small ($\sim$5%) increase in overall leakage, when correctly modeled, has virtually no

impact on sensitivity, consistent with the analysis performed in Ref. [36].

Using an inconsistent background model when setting a limit has a slightly larger effect. In this third scenario, the underestimation of ER leakage leads to a higher limit (overcoverage) than would be obtained using the correct model. This effect appears in both the PLR and cut-and-count analyses, and the magnitude of the effect grows with exposure—but remains much smaller than the 1-σ experiment-to-experiment variation in the 90% C.L. upper limit at 50 GeV/c^2 dark matter mass for the largest exposure simulated (3 kton-yr).

Though this work focuses on the inner-shell effect in the context of the neutrino background, ^{127}Xe itself is a transient background in LXe-TPC dark matter searches. This background can be modeled directly using the values for observed recombination presented in Sec. IV.

D. Future work

The modified recombination model presented above is sufficient for modeling the relatively small neutrino-electron background expected in the current generation of LXe-TPC dark matter searches and for planning next-generation searches where neutrino-electron scattering will be the dominant ER background to the weakly interacting massive particle signal. If (or when) these next-generation experiments go forward, further measurements of this effect should be performed to ensure correct statistical coverage and to address questions left unanswered in this work. We have not measured the tails of the S2/S1 distribution associated with L-shell vacancies—instead, we assume the same recombination fluctuations that are seen in β decays. We have not measured this effect when accompanied by an initially ejected electron—instead, we introduce an *ad hoc* model with reasonable limiting behavior. We have not measured this effect on the M shell, and we have explored only a small range in drift field—drift-field dependence should be expected, since recombination at higher fields is (in some models) more sensitive to small-scale changes in ionization density. All of these measurements are feasible, many with the ^{127}Xe technique described here.

Recently, Ref. [37] published a measurement of the average xenon work function of $W = 11.5$ eV, in agreement with a recent measurement from the EXO-200 Collaboration [38]. This value is 16% lower than the value of 13.7 eV adopted here and in the nominal NEST model and used widely throughout the field over the last decade. For our analysis, a shift in the value of W would be absorbed by an equal rescaling of the gain factors g_1 and g_2 and would impact our tritium and xenon-only datasets identically. The values of the absolute quanta yields ($\mathcal{Q}_y$ and $\mathcal{L}_y$) reported in this work would shift by the same amount, but the construction of the ratio q ensures the significance of the L-shell effect is unchanged. Similarly, any miscalibration of g_1 and g_2 stemming from a particular tritium recombination model will propagate into the absolute light and charge yields directly, but the mean q is robust against these effects. For example, changing from NEST release 2.2.1 patch 1 (used in this analysis) to the current development version (2.3-beta) results in a 5% decrease (increase) in L shell $\mathcal{Q}_y$ ($\mathcal{L}_y$) due to shifts in the inferred g_2 (g_1), but q shifts by only 0.3%.

[1] D. S. Akerib *et al.* (LUX Collaboration), Results from a Search for Dark Matter in the Complete LUX Exposure, Phys. Rev. Lett. **118**, 021303 (2017).

[2] X. Cui *et al.* (PandaX-II Collaboration), Dark Matter Results from 54-Ton-Day Exposure of PandaX-II Experiment, Phys. Rev. Lett. **119**, 181302 (2017).

[3] E. Aprile *et al.* (XENON Collaboration), Dark Matter Search Results from a One Ton-Year Exposure of XENON1T, Phys. Rev. Lett. **121**, 111302 (2018).

[4] E. Aprile *et al.* (XENON Collaboration), Constraining the Spin-Dependent WIMP-Nucleon Cross Sections with XENON1T, Phys. Rev. Lett. **122**, 141301 (2019).

[5] G. Jungman, M. Kamionkowski, and K. Griest, Super-symmetric dark matter, Phys. Rep. **267,** 195 (1996).

[6] L. Roszkowski, E. M. Sessolo, and S. Trojanowski, WIMP dark matter candidates and searches—current status and future prospects, Rep. Prog. Phys. **81,** 066201 (2018).

[7] D. S. Akerib *et al.* (LUX-ZEPLIN Collaboration), Projected WIMP sensitivity of the LUX-ZEPLIN dark matter experiment, Phys. Rev. D **101,** 052002 (2020).

[8] H. Zhang *et al.* (PandaX Collaboration), Dark matter direct search sensitivity of the PandaX-4T experiment, Sci. China Phys. Mech. Astron. **62,** 31011 (2019).

[9] E. Aprile *et al.* (XENON Collaboration), Projected WIMP sensitivity of the XENONnT dark matter experiment, J. Cosmol. Astropart. Phys. 11 (2020) 031.

[10] J.-W. Chen, H.-C. Chi, C.-P. Liu, and C.-P. Wu, Low-energy electronic recoil in xenon detectors by solar neutrinos, Phys. Lett. B **774,** 656 (2017).

[11] V. Chepel and H. Araújo, Liquid noble gas detectors for low energy particle physics, J. Instrum. **8,** R04001 (2013).

[12] D. S. Akerib *et al.* (LUX Collaboration), Discrimination of electronic recoils from nuclear recoils in two-phase xenon time projection chambers, Phys. Rev. D **102,** 112002 (2020).

[13] S. Nagy, *Radiochemistry and Nuclear Chemistry*, edited by S. Nagy (EOLSS Publishers, Oxford, 2009), Vol. 1, ISBN 9781848265769.

[14] International Atomic Energy Agency—Nuclear Data Services, *Livechart of nuclides* (International Atomic Energy Agency - Nuclear Data Services, 2021).

[15] A. Ringbom, Calculation of electron–photon coincidence decay of ^{131m}Xe and ^{133m}Xe including atomic relaxation, Appl. Radiat. Isot. **70,** 1499 (2012).

[16] C. Amole *et al.* (PICO Collaboration), Data-driven modeling of electron recoil nucleation in PICO C3F8 bubble chambers, Phys. Rev. D **100,** 082006 (2019).

[17] J. Brodsky *et al.*, *NESTCollaboration/nest: Geant4 Integration Fixes and Updates* (CERN, Meyrin, 2019).

[18] C. E. Dahl, The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter, Ph.D. thesis, Princeton University, 2009.

[19] D. S. Akerib *et al.* (LUX Collaboration), Ultralow energy calibration of LUX detector using ^{127}Xe electron capture, Phys. Rev. D **96,** 112011 (2017).

[20] S. Seltzer, *XCOM-Photon Cross Sections Database, NIST Standard Reference Database 8* (National Institute of Standards and Technology, 1987), https://dx.doi.org/10.18434/T48G6X.

[21] P. Linstrom, *NIST Chemistry WebBook, NIST Standard Reference Database 69* (National Institute of Standards and Technology, 1997), https://dx.doi.org/10.18434/T4D303.

[22] M. Rodrigues, M.-C. Lépy, P. Cassette, X. Mougeot, and M. Bé, Standardization of xenon-127 and measurement of photon emission intensities, Appl. Radiat. Isot. **87,** 342 (2014).

[23] S. Stephenson *et al.*, MiX: a position sensitive dual-phase liquid xenon detector, J. Instrum. **10,** P10040 (2015).

[24] H. Cao *et al.* (SCENE Collaboration), Measurement of scintillation and ionization yield and scintillation pulse shape from nuclear recoils in liquid argon, Phys. Rev. D **91,** 092007 (2015).

[25] FNAL, Neutron therapy facility, https://www-bd.fnal.gov/ntf/research_ni/index.html.

[26] H. I. Amols, J. F. Dicello, M. Awschalom, L. Coulson, S. W. Johnsen, and R. B. Theus, Physical characterization of neutron beams produced by protons and deuterons of various energies bombarding beryllium and lithium targets of several thicknesses, Med. Phys. **4,** 486 (1977).

[27] B. Loer, bloer/daqman: Citation release (version v0.4) (2019), https://doi.org/10.5281/zenodo.3347152.

[28] J. Angle *et al.* (XENON10 Collaboration), Search for Light Dark Matter in XENON10 Data, Phys. Rev. Lett. **107,** 051301 (2011).

[29] C. Faham, V. M. Gehman, A. Currie, A. Dobi, P. Sorensen, and R. J. Gaitskell, Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes, J. Instrum. **10,** P09010 (2015).

[30] V. N. Solovov *et al.* (ZEPLIN Collaboration), Position reconstruction in a dual phase xenon scintillation detector, IEEE Trans. Nucl. Sci. **59,** 3286 (2012).

[31] A. Vaitkus and C. Lynch (private communication).

[32] M. Szydagis, N. Barry, K. Kazkaz, J. Mock, D. Stolp, M. Sweany, M. Tripathi, S. Uvarov, N. Walsh, and M. Woods, NEST: a comprehensive model for scintillation yield in liquid xenon, J. Instrum. **6,** P10002 (2011).

[33] D. S. Akerib *et al.* (LUX Collaboration), Tritium calibration of the LUX dark matter experiment, Phys. Rev. D **93,** 072009 (2016).

[34] N. Vinyoles, Aldo M. Serenelli, Francesco L. Villante, S. Basu, J. Bergström, M. C. Gonzalez-Garcia, M. Maltoni, C. Peña-Garay, and Ningqiang Song, A new generation of standard solar models, Astrophys. J. **835,** 202 (2017).

[35] F. Ruppin, J. Billard, E. Figueroa-Feliciano, and L. Strigari, Complementarity of dark matter detectors in light of the neutrino background, Phys. Rev. D **90,** 083510 (2014).

[36] A. Gaspert, P. Giampa, and D. E. Morrissey, Neutrino backgrounds in future liquid noble element dark matter direct detection experiments, arXiv:2108.03248.

[37] L. Baudis, P. Sanchez-Lucas, and K. Thieme, A measurement of the mean electronic excitation energy of liquid xenon, arXiv:2109.07151.

[38] G. Anton *et al.*, Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment, Phys. Rev. C **101,** 065501 (2020).