Search for Neutrinoless Double Beta Decay Using the Full MAJORANA DEMONSTRATOR Dataset

Alexandru Hostiuc

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2023

Reading Committee:

Jason Detwiler, Chair

Alvaro Chavarria

Silas Beane

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Search for Neutrinoless Double Beta Decay Using the Full MAJORANA DEMONSTRATOR Dataset

Alexandru Hostiuc

Chair of the Supervisory Committee:
Professor Jason Detwiler
Physics

The MAJORANA Collaboration is searching for the neutrinoless double-β decay of ^{76}Ge. The MAJORANA DEMONSTRATOR consists of two modular arrays of high-purity Ge detectors operated in vacuum cryostats housed in a low-background shield at the Sanford Underground Research Facility in Lead, South Dakota. The arrays have operated with up to 48.5 kg of detectors (30.9 kg enriched to $\sim$88% in ^{76}Ge). The DEMONSTRATOR has completed operation of its enriched detectors, and a limit has been produced with 64.5 ± 0.9 kg-yr of enriched exposure. This thesis describes the development and implementation of the DCR (Delayed Charge Recovery) estimator, used for surface alpha background rejection. The properties of the alpha spectrum described herein show the absolutely critical importance of surface alpha background rejection, and DCR in particular, in the performance of the MAJORANA DEMONSTRATOR. The thesis also describes a parallel, independent statistical analysis of the full MAJORANA DEMONSTRATOR ^{76}Ge dataset, yielding a $0\nu\beta\beta$ decay half-life lower limit of 8.4×10^{25} years, and an upper limit on $m_{\beta\beta}$ of $112 - 267$ meV, depending on the nuclear matrix element used.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

GLOSSARY

$0\nu\beta\beta$: Neutrinoless Double-Beta Decay

$2\nu\beta\beta$: Two Neutrino Double-Beta Decay

$Q_{\beta\beta}$: Q-value of Neutrinoless Double-Beta Decay, i.e. the total kinetic energy released in a beta decay process

SM: The Standard Model of Particle Physics

BSM: Beyond the Standard Model

ROI: Region-of-Interest

MJD: The Majorana Demonstrator

GERDA: GERmanium Detector Array

LEGEND: Large Enriched Germanium Experiment for Neutrinoless Double-Beta Decay

HPGE: High-Purity Germanium

PPC: P-type Point Contact (detector)

ICPC: Inverted Coaxial Point Contact (detector)

BEGE: Broad Energy Germanium (detector)

DCR: Delayed Charge Recovery

AVSE: Multi-site Discrimination Parameter (Maximum Current Amplitude vs Energy)

LQ: Late Charge Parameter

PSA: Pulse-Shape Analysis

DS#: Majorana Demonstrator Data Set #

SURF: Sanford Underground Research Facility

LNGS: Laboratori Nazionali del Gran Sasso (INFN Gran Sasso National Laboratory)

FWHM: Full Width Half Maximum

LMFE: Low-Mass Front End

TRAPENFCAL: Calibrated Energy

BI: Background Index

FEP: Full-Energy Peak

SEP: Single Escape Peak

DEP: Double Escape Peak

CC: Compton Continuum

MS: Multi-site

SS: Single-site

MC: Monte-Carlo (Simulation)

SIGGEN: siggen is a computer code developed by David Radford of Oak Ridge National Laboratory that simulates the generation of waveforms in germanium detectors given a crystal geometry, an energy deposition and location, and some other configuration parameters.

WF: waveform, i.e. the trace that is recorded by the Majorana data acquisition system when an energy deposition is detected

GAT: Germanium Analysis Toolkit, a software package developed by the Majorana collaboration for processing the output of the Majorana Demonstrator

ACKNOWLEDGMENTS

I would like to say thank you to everyone, collaborators, family, friends, who supported me, during the sometimes arduous adventure of obtaining a Ph.D. in Physics.

Certainly, I would like to thank Jason Detwiler for being a wonderful advisor with whom I could collaborate during this difficult endeavor. Thank you for understanding my skill set and interests and helping expand my horizons while working for the MAJORANA DEMONSTRATOR Collaboration. I am grateful for this support, without which I probably wouldn't have made it this far. I would like to thank Nick Ruof for being an excellent friend in and out of the office, and working together to finish the official MAJORANA analysis. The moments of levity and laughter are just as appreciated as the help and support when I needed clarifications or a second opinion in work matters. I would also like to thank Walter Pettus, Ian Guinn and Micah Buuck, and chronologically later Clint Wiseman, for their support and help when I was starting out work for the MAJORANA DEMONSTRATOR. Thank you to David Peterson, Tim van Wechel and the rest of the CENPA Staff who have helped along the way. Thanks to Catherine Provost for her support and encouragement. Thanks also to the new generation of CENPA grad students, both in the LEGEND group (CJ Nave, Grace Song, Sam Borden) and outside it (Jimmy Sinnis, Harry Ni, Winston DeGraw), for the almost daily post-COVID lunches as well as fun events like Tuesday Trivia nights (which also often helped financially!) and Friday Table Tennis Tournaments to take the edge off. Thanks to Pak Kau Lim for being a great friend and talking things through when necessary.

Last but not least, thanks to important non-UW people in my life. Thanks to my mom Carmen and (posthumously) dad Alex for always supporting my educational goals and pushing me in the right direction whenever I needed. Thanks to my grandma, Viorica, for

helping raise me and for her love and support since birth. A posthumous thanks to my grandfather Victor for working hard so that my family has a future and financial support. Thanks to Monica Mocanu for helping me with my applications to Ph.D. programs, so as to gain admission to the UW. Thanks to Mircea Fianu, my 5th through 8th grade math teacher, for helping me develop my skills into a love of math and science, and thanks to Bill Griffiths, Tom Medlock and Jeff Sigler, my high school science teachers, for helping me progress on this journey in high school. Thank you also to Professors Jim Pilcher, David Miller and Stephan Meyer for the research experience I gained at The University of Chicago.

DEDICATION

To scientific progress

Chapter 1
INTRODUCTION

Neutrinoless double beta decay experiments are important because they can tell us, at a fundamental level, about the nature of the Universe, the nature of things, and could even help explain why we are able to exist. In addition, if neutrinoless double beta decay is observed, regardless of how it occurs, it will have been the first time matter was ever created in the laboratory by humans. This process is not strictly a conversion from energy into mass, but rather the creation of two new electrons, violating lepton number conservation. At the very least, this makes the process interesting and worth studying.

1.1 Radioactivity and Beta Decays

In the early part of the 20th century, scientists, among which Becquerel, Marie and Pierre Curie and Rutherford, discovered radioactivity. There are three primary types of decay modes:

1. α-decay: an α particle, or ^{4}He, is ejected from a nucleus.

2. β-decay: an electron (also known as a β) is ejected and a neutron turns into a proton. The analogous decay with a positron is also possible.

3. γ-decay: an excited state nucleon reverts to its ground state, releasing an energetic photon (γ).

There can be other classifications for nuclear decay processes, such as whether the nucleus breaks up, as is the case for α decays, or whether is suffers some internal transmutation, like during β decays.

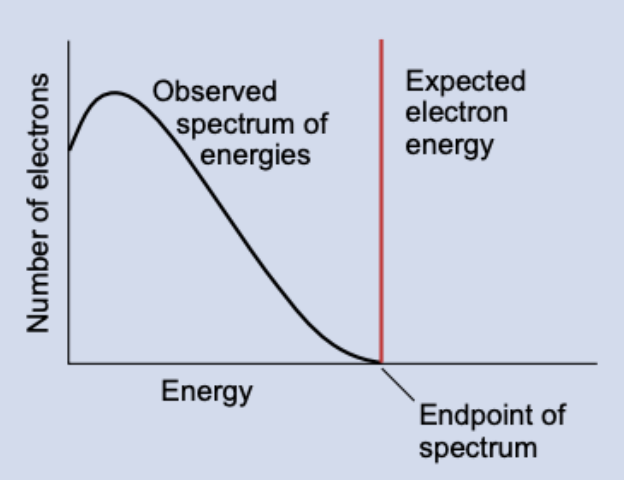

Figure 1.1: Energy spectrum of electrons emitted in beta decay, expected and observed [108]

Perhaps the most puzzling of these was the beta decay. While alpha and gamma decay energies can be easily predicted via energy conservation, the same cannot be said of the beta decay. Its spectrum was not at one given energy; rather, it was diffuse as can be seen in Fig. 1.1.

Beta decay was expected to be mono-energetic because the mass difference between a parent and daughter nucleus is a constant. If all the energy were to go to the electron, then the total energy of the electron must also be constant, and equivalent to the mass difference mentioned above. If this was not the case, and the energy of the electron can be lower than that difference, then there is, at first glance, some 'missing' energy. At that time, this was considered to be due to a violation of energy conservation. Hypotheses of quantum-level energy non-conservation predated the proposal of a neutrino. However, Pauli did not agree these ideas and, in the end, this opposition led to the prediction and discovery of a new particle, the neutrino. Perhaps energy conservation held on a quantum level, after all [34].

1.2 Neutrinos: Background and Proposals

In the 1920s, matter seemed simple. The consensus was that matter was composed of two kinds of elementary particles, the 'proton' and 'electron'. This inference about the nature of things was due to the fact that protons were emitted under alpha-particle bombardment [110], while electrons had been discovered in Thomson's cathode-ray experiments in 1897 — the first subatomic particle to be discovered [117]. Nothing else had been observed. However, in December of 1930, Pauli sent an open letter [104] to, among others, Hans Geiger and Lisa Meitner. The letter detailed a 'desperate remedy' to save the exchange theorem of statistics (related to the conservation of angular momentum) and the energy conservation theorem, due to observed inconsistencies in the statistics of ^{14}N and ^{6}Li as well the fact that the beta decay energy spectrum was continuous. He described a neutral particle inside the nucleus, which he called a 'neutron'. The 'Pauli-neutron' was thought to be a spin-1/2 fermion which had mass, albeit less than a hundredth of the mass of the proton. This small mass was theorized due to the fact that a nuclear mass change had not been seen and this particle needed to be light enough to carry away roughly half the kinetic energy, on average. This particle would solve the issue of the beta-spectrum's continuity, as both a 'Pauli-neutron' and electron would be emitted. What we now know as a neutron was finally discovered in 1932 by Chadwick [39] and had different properties.

In mid-1933, experiments by Ellis and Mott [55] showed there was a sharp upper limit in the beta-spectrum, corresponding to the energy difference between parent and daughter nucleus. This upper limit is called the Q-value; it was consistent with the idea of a neutrino [34].

Also in 1933, Enrico Fermi first published his theory of beta decay [71] [72]: that a neutrino ('small neutral one') accompanies the beta-decay electron, both being created at the moment of emission. This theory remains the basis of the current understanding of beta-decay, although more details have since been put to paper. The main difference between Fermi's theory and Pauli's earlier model is that Pauli had predicted the neutrinos to be

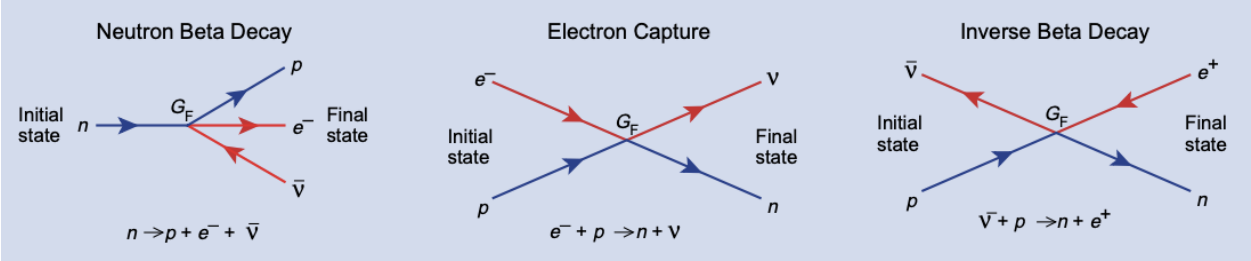

Figure 1.2: Examples of 4-point interactions [108]

present in the nucleus itself, rather than be spontaneously created during the beta decay process itself [34].

Fermi's interaction is an explanation of β-decay; in this theory, four fermions directly interact with one another, at one vertex of the associated Feynman diagram. As can be seen in the leftmost part of Fig. 1.2, the neutron decays into proton, electron and neutrino, with a rate related to the Fermi coupling constant G_F. This is a limiting case of the current theory, specifically a low-energy effective field theory (EFT). It was a crucial point in the development of Quantum Field Theory (QFT), and the precursor of the weak interaction between a proton-neutron and an electron-antineutrino, mediated by a vector boson, W^-.

1.3 Discovery of the Neutrino and Some of Its Properties

It was all well and good in theory, but did it work in practice? Yes, it did. In 1956, Cowan, Reines et al. confirmed the existence of the neutrino [41]. Due to the neutrino's lack of charge and predicted massless nature, it was rather hard to detect. The Cowan-Raines experiment used then-hypothetical electron anti-neutrinos from a nuclear reactor, in their third attempt at a neutrino measurement. The detector consisted of two Cadmium-loaded water tanks in between three large scintillator tanks; interactions with the protons were observed. The reaction itself was as follows: $\overline{\nu} + p \rightarrow n + e^+$. The positron would then be annihilated by a nearby electron, releasing energy in the form of two counter-propagating photons. A photon cascade would be released upon the neutron's capture by a Cd nucleus. These three photon

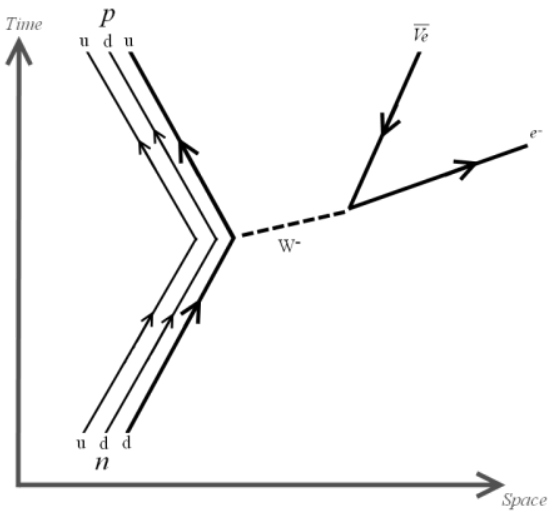

Figure 1.3: Feynman diagram of β decay

signals provided a unique, identifiable signature, using the delayed coincidence technique. Thus, the neutrino's existence was confirmed. Cowan and Reines would receive the Nobel Prize almost 40 years later, in 1995.

While parity (P) symmetry is conserved for the strong and electromagnetic forces, Lee & Yang suggested in 1956 that P could be violated by the weak force [87]. Wu, in 1957, experimentally showed P violation by looking at the correlation between the directions of deexcitation γs and the electrons emitted in the decay of ^{60}Co in a strong magnetic field at low temperature. Wu found an anti-correlation, implying P is violated [57]. In 1958, Goldhaber observed the weak force maximally violates P, and weak interactions involve only left-hand neutrinos [74]. Since neutrinos were presumed massless, this implied there was no experimental evidence about the existence of right-handed neutrinos – therefore it was assumed all neutrinos were left-handed.

Discoveries of the fundamental properties of neutrinos would continue through the 1960s.

In 1962, Lederman, Schwartz and Steinberger would confirm the existence of the muon neutrino, showing that several distinct types of neutrinos exist. The third type, the tau neutrino, would be discovered as late as 2000 by the DONUT Collaboration.

1.4 Neutrino Oscillations

Starting in 1964, the first signs of trouble, or 'new physics', appeared. There was seemingly a Solar neutrino problem, first seen in the Homestake experiment [29]; that is, there was a large discrepancy between the predicted flux of neutrinos from the Sun and the number measured in the Homestake mine. The observed flux was much lower, about two-thirds lower, than calculated in the Standard Solar Model.

Ray Davis et al. constructed a 0.61 kiloton experiment in the Homestake mine in Lead, South Dakota. The reaction the experiment would be looking at is the capture of an electron neutrino on Chlorine-37: $\nu_e + {}^{37}Cl \rightarrow {}^{37}Ar + e^-$. The few Argon-37 atoms would be collected in a gas counter; Argon-37's radioactive decays would be counted and therefore the number of neutrinos could be estimated. After 25 years of data taking, the rate was measured at $2.55 \pm 0.17(stat) \pm 0.18(syst)$ solar neutrino units, or SNU. This unit equals the neutrino flux producing 10^{-36} captures per target atom per second. Meanwhile, the theoretical predictions ranged from 8.0 ± 1.0 SNU to 6.4 ± 1.4 SNU [77], corresponding to a factor 2 to 3 disparity.

Solar neutrinos, which the Standard Solar Model predicted as electron neutrinos only, were expected to travel unimpeded through the Sun's outer layers, and then the Earth's atmosphere, and rock, until reaching Ray Davis' chlorine-based detector. However, a large number of neutrinos, roughly two-thirds of those expected, would appear to have vanished. Later experiments, like SAGE/GALLEX and early runs of Kamiokande, confirmed the discrepancy. SAGE (Soviet-American Gallium Experiment) measured the radio-chemical solar neutrino flux in Gallium-71 using inverse beta decay. The resulting Germanium-71 was extracted from the liquid Ga and measured using proportional chambers [45]. GALLEX (and its follow-up GNO) was a very similar experiment based at the Laboratori Nazionali del-Gran Sasso (LNGS), which used the same reaction and also used proportional counters to

measure the ^{71}Ge content in $^{71}GeH_4$ [44]. The results from these experiments suggested that the neutrino fluxes were incompatible with the Standard Solar Model (SSM), leading to the creation of SNO (Sudbury Neutrino Observatory) and Super-Kamiokande.

In the end, after more than three decades of progress, it was proven that there was actually no problem. The solution depended on one of the neutrino's fundamental properties. If the neutrinos had mass, unlike described in the Standard Model of Particle Physics, then they could oscillate between their three different flavor/interaction eigenstates: electron, muon and tau neutrinos. The oscillation hypothesis was finally confirmed in 2001 and 2002, when SNO [46] and Super-Kamiokande [64] both detected solar neutrino oscillations. Atmospheric neutrino oscillations (muon-tau neutrinos) had been discovered by Super-Kamiokande in 1998 [47].

1.4.1 Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix

The mathematical object that describes the nature of neutrino oscillations is the unitary Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix:

$$U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{1.1}$$

Here, c_{ij} and s_{ij} are the cosines and sines of the mixing angles between the eigenstates, θ_{ij}. δ_{CP} is the CP phase, which is a CP-violating term. If the neutrinos are Majorana particles, the PMNS matrix is allowed to have additional CP phases – specifically, the α_1 and α_2 terms are the Majorana CP phases. The PMNS matrix relates the neutrino's flavor eigenstates $\nu_{e,\mu,\tau}$ to its mass eigenstates $\nu_{1,2,3}$. This is necessary as when neutrinos are produced, or interact in any way, they are in their flavor eigenstates, whereas when they propagate they must be in their mass eigenstates, due to energy conservation [111] [105]. During propagation, then, a neutrino produced as, e.g., an electron neutrino can convert to a muon and/or tau neutrino,

and back again, with some probability that depends on the difference in the squares of the masses 1, 2, 3. In a 2-neutrino approximation, the transition probability is:

$$P_{\alpha\to\beta} = \sin^2(2\theta_{ij}) \sin^2\left(\frac{(\Delta m^2)_{ij} c^3}{4\hbar E} L\right) \tag{1.2}$$

The $(\Delta m^2)_{ij}$ is the difference of the squares of masses of the neutrino mass eigenstates i and j. E is the energy of the neutrino. L is the distance the neutrino travels from its source.

1.5 What Do We Really Know About The Neutrino?

1.5.1 Inverted & Normal Hierarchy

Not only do we not yet know the absolute neutrino masses (because of the unbearable lightness of neutrinos, they are very difficult to weigh, and we are stuck with an upper bound of roughly 0.8 eV as seen by the KATRIN experiment [19]) — we don't even know which is the heaviest mass eigenstate, i.e. the mass ordering, or hierarchy. It could be that $m_1 < m_2 < m_3$ or that $m_3 < m_1 < m_2$, as in Fig. 1.4. The sign of the solar mass-squared difference Δm^2_{21} is known to be positive from the results of the solar-neutrino measurements previously mentioned. The sign of the atmospheric mass-squared difference (Δm^2_{32}) remains yet unknown, with its absolute value being roughly a factor of 30 higher than the solar one. Current experiments indicate that the normal ordering is favored at $\sim 2.7\sigma$.

1.5.2 Neutrino Mass Weirdness [77]

It is a strange feature of particle physics that neutrinos should have the small masses that they do. No symmetry requires the neutrinos to be massless, however, the more than 6 orders of magnitude separating the electron mass from the much lighter neutrinos raise a serious question: why should neutrinos be so light? The current KATRIN upper limit on the neutrino mass is 0.8 eV [19], while the electron's rest mass is 511 keV.

In the Standard Model, neutrinos are regarded as strictly massless, as discussed previously. In this case, there is simply no interaction with the Higgs, which gives masses to the

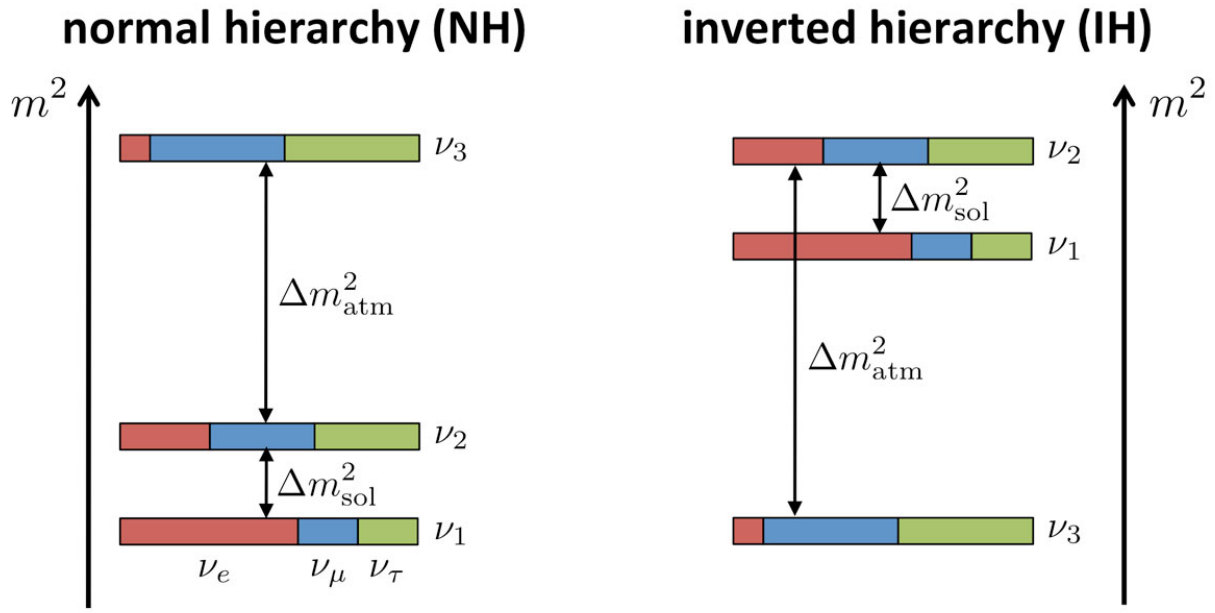

Figure 1.4: Neutrinos could have either normal or inverted hierarchy (or mass ordering). Credit: JUNO Collaboration/JGU-Mainz

other SM particles; Dirac mass terms similar to those of other particles cannot be constructed due to the absence of right-chiral neutrino fields in the SM. An option to consider, that would yield non-zero masses for the neutrino would be the existence of Majorana masses for the neutrino. In extended models, it is natural to have the same mass scale for fermions of the same family (e.g. the electron neutrino, electron, up quark and down quark), due to similar Yukawa couplings with the Higgs being the 'simplest', or most 'natural' explanation. In this scenario, the small neutrino masses can be explained via the Majorana masses, through the See-saw Mechanism.

1.5.3 See-Saw Mechanism

There are two distinct classes of neutrino mass terms allowed in the Lagrangian, called Dirac (m_D) and Majorana (m_R) masses. The Dirac mass relates left-handed (ν_L) and right-handed (ν_R) neutrinos while the Majorana term has interactions between states of the same chirality. In the simplified 1 neutrino model, the mass terms look as follows:

$$\mathcal{L}_D = -m_D \overline{\nu_R} \nu_L + h.c. \tag{1.3}$$

$$\mathcal{L}_R = -\frac{m_R}{2} \overline{(\nu_R)^c} \nu_R + h.c. \tag{1.4}$$

The simplified 1-neutrino mass Lagrangian, after diagonalizing the mass and re-writing, is:

$$\mathcal{L} = -\frac{1}{2}\frac{m_D^2}{m_R}\overline{\nu}\nu - \frac{1}{2} m_R \overline{N} N \tag{1.5}$$

Note that the final expression for the Lagrangian (Eq. 1.5) shows that the light and heavy neutrinos ν and N are Majorana particles, as the terms mix ν and $\overline{\nu}$ and N and $\overline{N}$ respectively. Many classes of GUT contain heavy Majorana singlets, which predict the see-saw mechanism, leading to very light-mass Majorana neutrinos.

The masses obtained are $\frac{m_D^2}{m_R}$, or m, for the light neutrino, and m_R, or M, for the heavy neutrino. The choice for M is usually $\sim 10^{15}$ GeV, around the Grand Unified Theory (GUT) scale, because the right-handed singlet's Majorana mass term does not originate via the Higgs mechanism and therefore wouldn't be expected to have a Higgs-like mass. Then, the natural mass scale would be that of new physics, in most cases considered the GUT scale. With mass m_D at an order of 100 GeV, a 'natural' choice, similar to the Higgs field's vev of 246 GeV, the light neutrino mass m is on the 0.01 meV scale, consistent with the latest experimental results. This is the See-saw Mechanism. The factor $\frac{m_D^2}{m_R}$ is the small parameter characterizing the scale difference between neutrinos and the other particles.

1.5.4 Dirac vs Majorana: is the neutrino its own antiparticle?

Neutrinos do not carry electric charge or color charge. A curious observation is that all neutrinos so far observed have been left-handed, i.e. that their momentum and spin are opposite in direction. While the CPT theorem tells us that there must be a right-chiral anti-neutrino, it is unknown if this particle would be the same as a right-handed neutrino. In theory, if one were able to Lorentz-boost past a neutrino, we would witness a flip in the direction of its momentum, and observe a right-handed neutrino. The question is, whether the right-handed neutrino interacts with charged leptons (in which case it is the right-chiral Majorana neutrino) or is sterile (the right-chiral Dirac neutrino) [83]. In the ultra-relativistic limit, which applies to neutrinos due to their extremely low mass, this boost is next to impossible. Therefore, this is remains a very hard question to approach directly, and different solutions have to be investigated.

A possible way to figure out whether a neutrino is its own anti-particle is to look for lepton number violating (LNV) processes. While lepton number so far appears to be a conserved quantity, it might not actually be one. There is no symmetry generating lepton number conservation, according to Noether's Theorem [100], therefore it is referred to as an accidental symmetry. If a neutrino has any charge, including lepton number, then it cannot be Majorana, i.e. its own anti-particle. As such, the observation of a LNV process, such as neutrinoless double beta decay, would immediately indicate the Majorana nature of the neutrino.

Why should we care if the neutrino is Majorana?

At first glance, it might seem that the nature of the neutrino might be more a matter of bookkeeping than actual interest. However, this is not true, as our very existence might be linked to this property.

1.6 Baryogenesis

From observations of our Universe, such as Baryon Acoustic Oscillations (BAOs) in the Cosmic Microwave Background (CMB) and Big Bang Nucleosynthesis (BBN), we know that there are more baryons than anti-baryons in the Universe [42]. BAOs are the observed density fluctuations of the baryonic matter content of the Universe, created by acoustic waves in the primordial plasma of the early Universe. BBN is the production of light, non-^{1}H, nuclei during the early Universe, within roughly 10 minutes after the Big Bang. It is responsible for the creation of, and relative proportions between, light isotopes such as ^{3}He, ^{4}He and ^{7}Li [43].

We also know that our entire observable Universe is made of matter rather than antimatter. If a matter and an antimatter particle collide, they emit gamma-rays; since we do not observe this gamma-ray emission zone anywhere in the sky, the matter zone (also known as the 'sequestration scale') is at least larger than the cosmological horizon. However, there is no feasible mechanism that would allow such a large sequestration scale [42].

A question that remains is whether the asymmetry existed from the very beginning. One example of fine-tuning is such a possibility: in this case, in the Universe's initial conditions there would be an extra matter particle in roughly 1 billion. There is, however, no theory explaining why the Universe would have this imbalance 'written in'. Another point against the initial asymmetry is the idea that any such pre-existing asymmetry would be diluted by inflation by a factor of 10^{26} or worse. Therefore, assuming the asymmetry is an initial condition actually makes everything harder to explain, as the initial conditions would need even more asymmetry. Thus, the baryon asymmetry is thought to be generated dynamically through a process called baryogenesis. It is also thought that the Universe started with equal numbers of matter and anti-matter particles [42].

The aformentioned BBN and CMB BAOs can help us quantify the asymmetry. Both agree that the Universe is roughly 4.6% baryons, which corresponds to a baryon-to-photon ratio of roughly $6 * 10^{-10}$, as can be seen in Fig. 1.5. In this figure, the baryon-to-photon

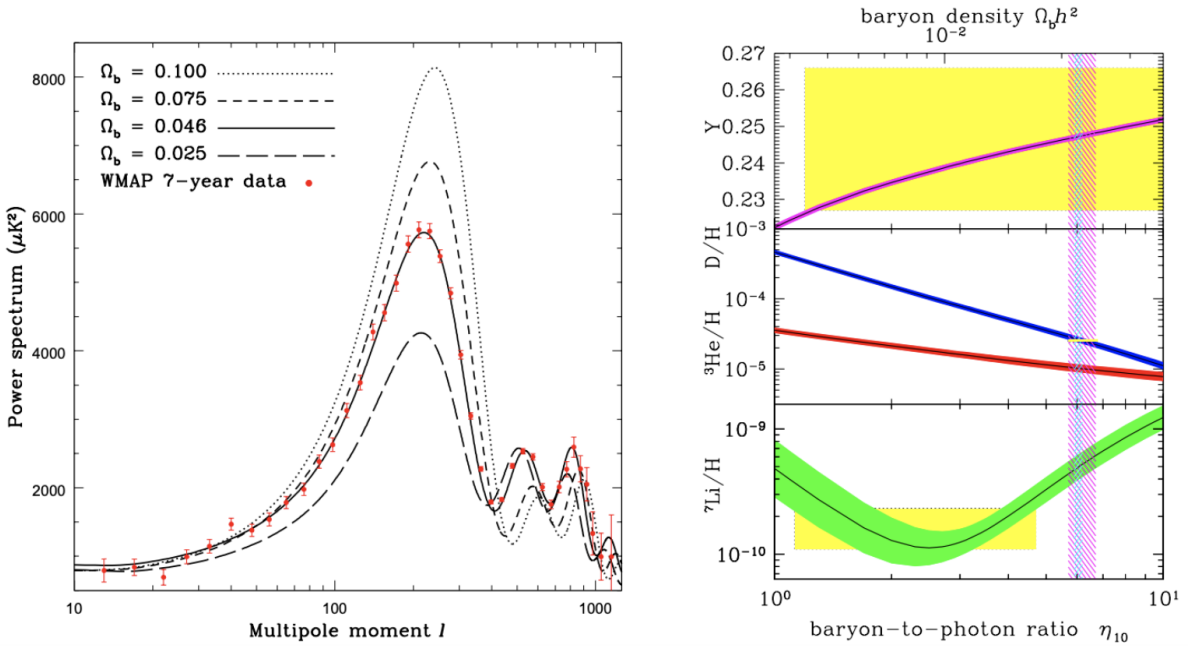

Figure 1.5: Sensitivity of the first acoustic peak of the CMB BAO to the average baryon density of the Universe (left) and concordance between the cosmic abundance of light nuclei and the baryon density, as required by Big Bang Nucleosynthesis (right). Taken from Cline [42]. The blue and pink bands are the 1 and 2-σ bands around the best-fit value to the baryon-to-photon ratio.

ratio η_{10} is in units of 10^{-10}. This is related to the baryon asymmetry parameter that will be discussed in the next section. It is a measure of the asymmetry if we assume that baryons and antibaryons have had enough time to annihilate, and be converted into photons.

What could have created this asymmetry? There are several proposed explanations, including GUT baryogenesis, Electroweak baryogenesis, Affleck-Dine mechanism [12], and Leptogenesis. The latter will be discussed in more detail in the next section.

These explanations must satisfy three conditions known as the Sakharov conditions [112]: 1) baryon number must be violated, 2) there must be a loss of thermal equilibrium, as otherwise the rates for creating and destroying baryons would be the same, and 3) C and CP violation, as one needs some way to prefer matter to anti-matter (C) and tell apart forward

and backward in time; otherwise, processes containing baryons would proceed at the same rate as C-/CP-conjugate processes with anti-baryons.

1.7 Leptogenesis

Leptogenesis is "a mechanism [...] to generate cosmological baryon number excess. The lepton number excess originating from Majorana mass terms may transform into the baryon number excess through [...] EW processes at high T" [73]. This theoretical model thus relates the matter-antimatter asymmetry to properties of the neutrinos.

There is a connection between baryon number B and lepton number L in the high temperature, symmetric phase of the Standard Model. Since weak interactions are chiral in nature, B and L are not conserved. Around the critical temperature of the electroweak phase transition, T_c, B and L violating processes come into thermal equilibrium. In the SM, there are effective interactions of all left-handed fermions which violate baryon and lepton number by 3 units, called sphaleron processes, which generate cosmological baryon asymmetry in the hot early universe. Specifically, B and L violating processes are in thermal equilibrium for temperatures in the range 10^{10} and 10^{12} GeV.

Decays of the N_1 — the lightest see-saw heavy neutrino — can generate a lepton asymmetry. There are no SM gauge interactions for N_1, so N_1 can propagate out of equilibrium with thermal quarks, leptons and Higgs. Its Lagrangian violates the discrete C and CP symmetries and lepton number. Sphaleron processes partially convert the lepton asymmetry into a baryon asymmetry. Thus, baryon number is violated through the sphaleron process in the early Universe [35]. Therefore all of the Sakharov conditions are satisfied and lepton asymmetry can be generated.

Once the temperature drops below the mass M_1 of the particle N_1, the heavy neutrinos cannot follow the change of the equilibrium distribution. There are too many heavy neutrinos, implying a deviation from thermal equilibrium. Eventually, the heavy neutrinos decay, generating a B-L asymmetry via CP-violating processes. One can predict a value for the baryon asymmetry:

$$\eta_B = -\frac{c_s}{f} N_{B-L} = -\frac{3}{4}\frac{c_s}{f}\epsilon_1 \kappa_f \sim 10^{-2}\epsilon_1\kappa_f \text{ [35]} \tag{1.6}$$

Here,

- c_s is the fraction of lepton asymmetry converted into baryon asymmetry via the sphaleron process.

- f is a dilution factor due to the increase in photon number density between leptogenesis and recombination.

- N_{B-L} is the amount of B-L asymmetry in a comoving volume element that contains one photon at the time of leptogenesis.

- ϵ_1 is the CP asymmetry (on the order of 10^{-5} to 10^{-6}).

- The κ_f term represents the effect of washout processes in the plasma, obtained via solving the Boltzmann eq. in plasma (on the order of 10^{-1} to 10^{-2}).

Putting all the terms together yields a value of the baryon asymmetry η_B on the order of 10^{-8} to 10^{-10}, as observed. Therefore, the leading theory for leptogenesis predicts the existence of a particle that would give the neutrinos Majorana masses via the see-saw mechanism. A heavy neutrino mass hierarchy comparable to the one of quarks and charged leptons, which leads to a small CP asymmetry, together with the kinematic factors f and κ_f can explain the observed matter-antimatter asymmetry [35].

1.8 Neutrinoless Double Beta Decay

The question that remains is, how can we determine if the neutrino is indeed its own antiparticle. One promising way to test this is to search for the process known as neutrinoless double beta decay [113].

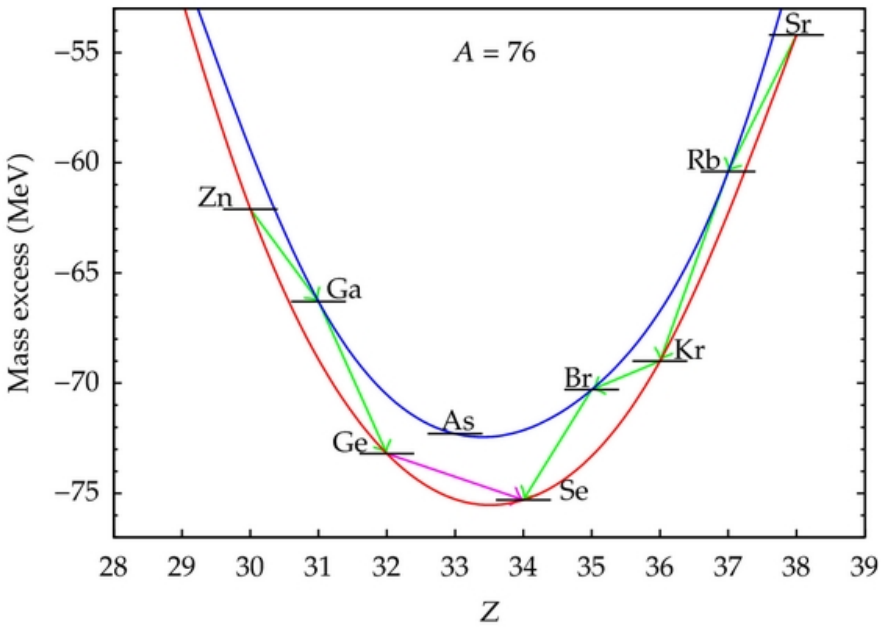

Figure 1.6: Why double beta decay is allowed in certain cases – here for ^{76}Ge (Menendez). Since the mass excess is higher for As than for Ge, beta decay cannot proceed. However, Se has a lower mass excess, so double beta decay can then proceed.

Standard model double beta decay is just what it sounds like: two beta decays happening at once. This process can be more easily observed when single beta decay is energetically forbidden (see Fig. 1.6), but the decay to the second nearest neighbor is allowed, i.e. the latter has lower energy. It is at its core a rare process with lifetimes around 10^{20} years.

Neutrinoless double-beta decay ($0\nu\beta\beta$) is a hypothetical lepton-number violating nuclear process involving the emission of two electrons with no neutrinos. Its discovery would indicate that the neutrino is a Majorana fermion. The experimental signature of $0\nu\beta\beta$ is a peak in total electron kinetic energy at the Q-value ($Q_{\beta\beta}$) of the decay [26].

The Schechter-Valle, or Black Box (Fig. 1.7), theorem indicates that no matter what the underlying, fundamental process or mechanism actually is, neutrinoless double beta decay indicates the Majorana nature of the neutrino. Specifically, any neutrinoless double beta

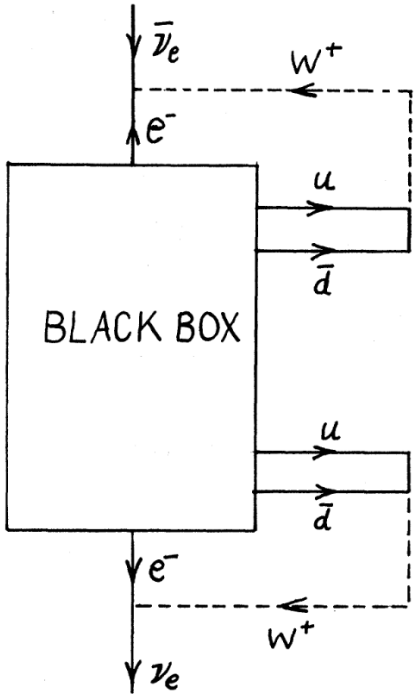

FIG. 2. Diagram showing how any neutrinoless double-β decay process induces a $\bar{\nu}_e$-to-ν_e transition, that is, an effective Majorana mass term.

Figure 1.7: The Black Box of Schechter-Valle [113]

decay process induces an effective Majorana mass term [113].

We will limit discussion to the light neutrino exchange model, as an example. The two neutrinos emitted in standard double beta decay are no longer in the final state of the process; in effect, a virtual neutrino is exchanged between the two weak vertices in the tree-level Feynman diagram, as in Fig. 1.8.

As briefly mentioned, this is similar to double beta decay except that the final state is missing any neutrinos. This takes us all the way back to the initial mention of the diffuse beta decay spectrum; we must now look for something that scientists in the 1920s might have expected to see: a mono-energetic peak at $Q_{\beta\beta}$ (the Q-value), specifically 2039 keV for ^{76}Ge. However, it would not displace the diffuse double beta decay spectrum, but appear

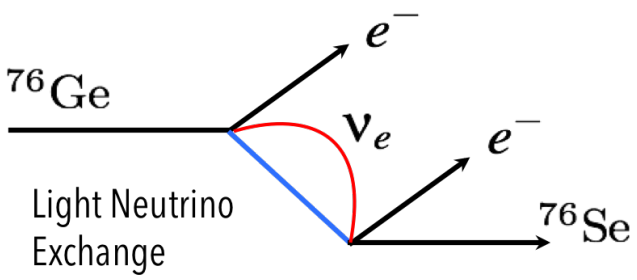

Figure 1.8: Feynman diagram of Germanium-76 neutrinoless double beta decay

alongside it. Theory indicates that, if it exists at all, neutrinoless double beta decay is a very rare process. However, a robust program involving many low-background experiments is engaging in a search for $0\nu\beta\beta$ in multiple isotopes, with half-life limits in some cases surpassing 10^{26} y [21, 14, 3, 24, 11, 28, 20, 25]. Well-motivated theoretical predictions yield a probable half-life for $0\nu\beta\beta$ of 10^{26} to 10^{30} years, as described below, possibly within the reach of next generation experiments [15].

A half-life measurement is sensitive to a coherent sum of the neutrino masses, $m_{\beta\beta}$, and nuclear matrix elements (NME) that must be calculated through many body nuclear theory [56]. To calculate the expected half-life of neutrinoless double beta decay, one needs to include phase space factor $G_{0\nu}$, an effective Majorana mass $m_{\beta\beta}$ and an NME $M_{0\nu}$. The equation relating these factors to the half-life is:

$$[T_{1/2}^{0\nu\beta\beta}]^{-1} = G_{0\nu}\ |M_{0\nu}|^2\ |m_{\beta\beta}|^2. \tag{1.7}$$

$|m_{\beta\beta}|$ is usually taken to be on the order of 0.001 to 0.1 eV as in Fig. 1.9. The inverse of $G_{0\nu}$ depends on the isotope undergoing neutrinoless beta decay but is roughly $10^{25}\ yr\ eV^2$. Determining the value of the NME $M_{0\nu}$ the most difficult part of the computation, due to the complicated nuclear structure. Gamow-Teller and Fermi terms include initial and final nuclear wave functions, which have no exact solution. Different models can lead to drastic changes between computed $|M_{0\nu}|$, for example between roughly 2-6 for ^{76}Ge, as can be seen

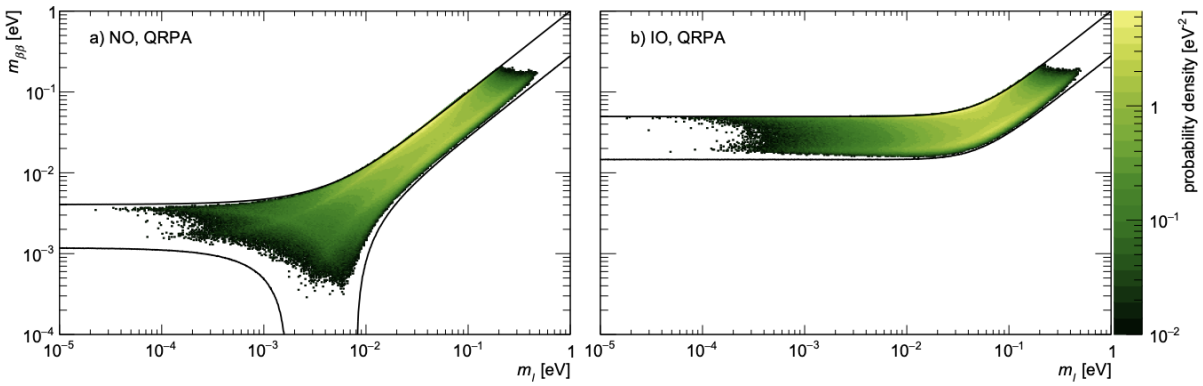

FIG. 1. Marginalized posterior distributions for $m_{\beta\beta}$ and m_l for NO (a) and IO (b). The solid lines show the allowed parameter space assuming 3σ intervals of the neutrino oscillation observables from nu-fit [12]. The plot is produced assuming QRPA NMEs and the absence of mechanisms that drive m_l or $m_{\beta\beta}$ to zero. The probability density is normalized by the logarithm of $m_{\beta\beta}$ and of m_l.

Figure 1.9: $m_{\beta\beta}$ phase space for inverted and normal ordering [15]

in Fig. 1.10. Another issue is g_A quenching, which is used to bring theoretical predictions in line with experimental observations. For $0\nu\beta\beta$, the factor is probably no more than 20-30% in $M_{0\nu}$, corresponding to a factor of $\sim$2 in half-life computations. Putting everything together yields the aforementioned value of about 10^{26} to 10^{30} years [15, 56].

1.8.1 How to actually observe $0\nu\beta\beta$?

The previously mentioned mono-energetic peak consists of the summed kinetic energies of the two emitted electrons that are captured within the active volume of a detector. Therefore, measuring the summed energy at the Q-value, and making sure it can not be explained by some background process, is necessary and sufficient for discovering a $0\nu\beta\beta$ process. As a rare event with low backgrounds, this is described by Poisson statistics.

Several isotopes can be used to search for $0\nu\beta\beta$, including ^{76}Ge, ^{136}Xe, ^{130}Te. Each isotope has advantages and disadvantages and none is intrinsically better suited for the search.

- Q-value: A higher Q-value is advantageous as it allows for less background in the

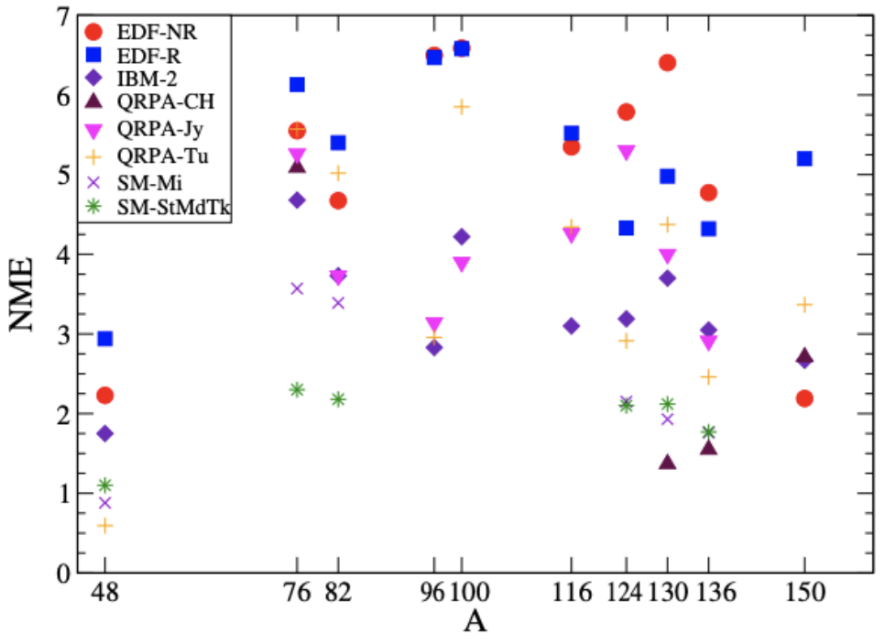

Figure 1.10: NMEs from Dolinski & Poon [53]

region of interest around this value. This is due to the fact that many processes can down-scatter, i.e., lower the energy of a background process, but there are none that allow the energy to go up; while the Q-value depends strictly on the isotope used, most backgrounds remain the same, hence a high Q-value is desired. The highest prominent gamma ray in natural radiation is ^{208}Tl's 2615 keV gamma. Ideally, the Q value would be above this line.

- Natural abundance of active isotope: most elements have a very small percentage of the desired isotope occur naturally. Therefore, it's advantageous to search using either those elements with a high fraction of the active isotope occurring naturally, like ^{130}Te (34%) or ones that can be enriched in a relatively straight-forward manner, like ^{76}Ge and ^{136}Xe. Having a high natural abundance, however, saves the cost of enrichment.

While ^{76}Ge has a low natural abundance and a relatively low Q-value, Germanium detec-

tors are an established technology and offer the best energy resolution, which is invaluable in rejecting background. A description of Germanium detectors and how they work, as well as details on the MAJORANA DEMONSTRATOR will be given in the following chapters.

1.9 Conclusion

Due to its electric charge neutrality, the neutrino is the one fundamental particle that could be Majorana, i.e. its own anti-particle. Overall, there is a good chance to discover whether this is so, using the next generation of experiment, such as LEGEND-200 and LEGEND-1000. Calculated distributions for $m_{\beta\beta}$ and m_l (lightest neutrino mass) indicate that the inverted hierarchy will be fully covered, and the normal hierarchy partially covered if LEGEND-1000 will reach its design sensitivity. While the normal ordering is only partially covered, recent results indicate that large swaths (more than 50%) of phase space will be available for testing, especially in the absence of mechanisms that drive one of the masses to 0. This can be seen in Fig. 1.11, which presents expected discovery probabilities as of 2017 for the next generation of germanium, tellurium and xenon based experiments. In all cases the expected discovery probabilities surpass ∼80% for the inverted ordering and ∼45% for the normal ordering [15].

Fig. 1.12 shows the most current expectation of backgrounds and exposure for next generation experiments. A fair number of projects are expected to be surpass the half-life sensitivities required to test the bottom of the inverted hierarchy, specifically CUPID and Amore-II in ^{100}Mo, nEXO and NEXT-HD in ^{136}Xe and LEGEND-1000 in ^{76}Ge [16].

If neutrinoless double beta decay is seen, regardless of the process underlying it, neutrinos must be Majorana particles and matter creation will have been observed for the first time. Theoretically, there are also good reasons to suspect the neutrino is a Majorana particle. Our best models for new particles around the grand unification scale, the matter-antimatter asymmetry, via leptogenesis, as well as the observed lightness of the neutrino masses predict neutrinoless double beta decay and Majorana masses to exist, and anticipate the Majorana nature of neutrinos. All that remains is to see whether the supposition is true.

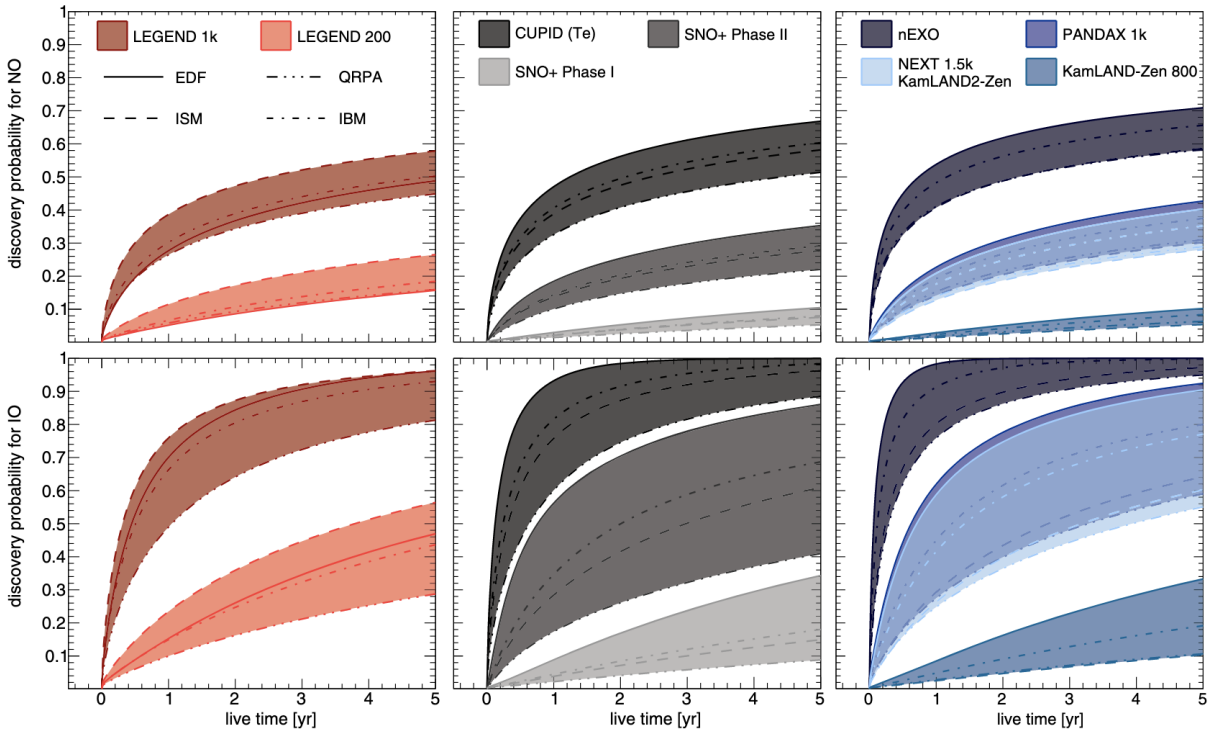

FIG. 4. Discovery probability as a function of live time for a selection of next-generation experiments grouped according to the target isotope (from left to right, ^{76}Ge, ^{130}Te, ^{136}Xe), assuming the absence of mechanisms that drive m_l or $m_{\beta\beta}$ to 0. The top panels show the discovery probability for NO, the bottom panels for IO. The variation of the NME among models is represented by the shaded regions.

Figure 1.11: Discovery probability for different $0\nu\beta\beta$ experiments as of 2017 [15]

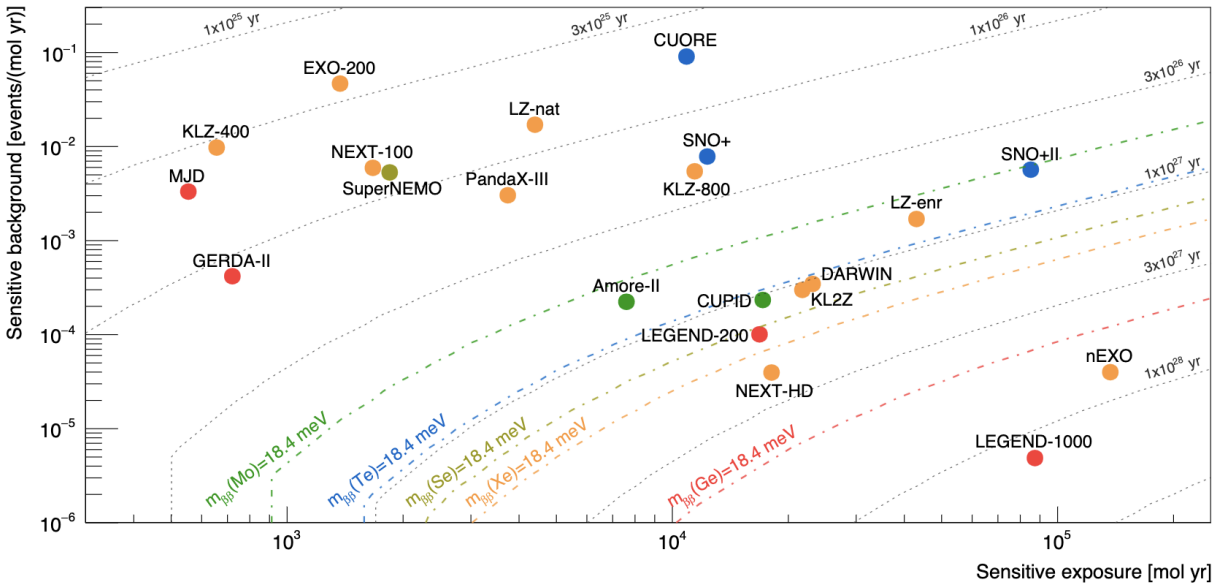

Figure 1.12: Sensitive background and exposure expectation for different $0\nu\beta\beta$ experiments as of 2022 [16]. As can be seen, several next generation experiments are expected to be able to cover the inverted ordering, i.e. to be able to claim discovery of $0\nu\beta\beta$ in the case of inverted mass ordering.

Chapter 2

EXPERIMENTAL DETAILS

2.1 Germanium Detector Technology

This chapter will include a description of germanium detectors, such as their history, production, properties and performance.

The longest-running and most well tested technology for $0\nu\beta\beta$ decay searches is the High Purity Germanium (HPGe) detector. The first $0\nu\beta\beta$ decay experiment based on Ge detectors was performed in 1967, by Fiorini et al. [16] [58].

HPGe detectors are single-crystal semiconductors. A detector is grown using the Czochralski method [32] from Ge material, which can be enriched up to 92% in ^{76}Ge. The detectors used in current and upcoming experiments are p-type crystals. The semiconductor junction is formed between the n+ electrode and the p-type crystal. The detector is then fully depleted by applying a reverse bias of a few thousand volts [16]. A typical p-type point contact (PPC) enriched Ge (enrGe) detector in the MAJORANA DEMONSTRATOR can be seen in Fig. 2.1.

In general, $0\nu\beta\beta$ decay experiments use one of three types of detectors [16]:

- Solid-state detectors with an embedded source.
- Monolithic liquid or gas detectors with an internal (dissolved) source.
- Composite detectors with external sources.

As germanium detectors are solid, we will limit the discussion to solid detectors. Solid-state $0\nu\beta\beta$ detectors consist of crystals grown from material which contains the $\beta\beta$ isotope. Masses are typically on the order of a kilogram, depending on the element used. Volumes

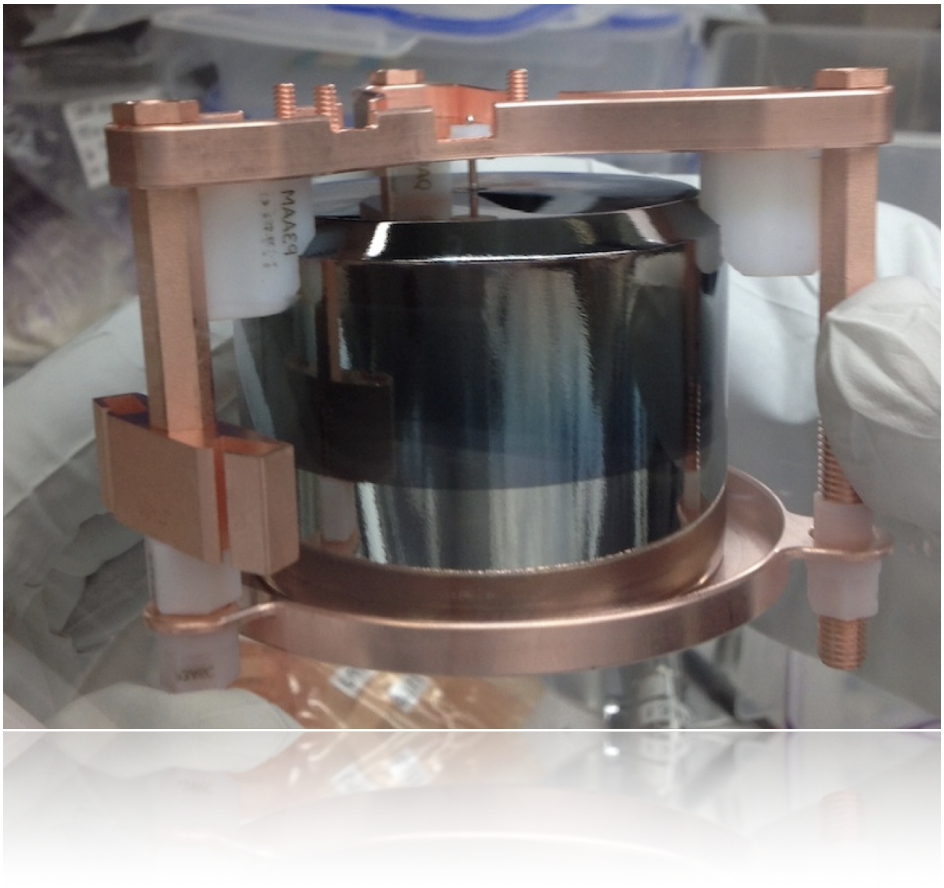

Figure 2.1: An enriched Ge PPC detector.

are usually on the order of a few hundred cubic centimeters. Crystal-based experiments use calorimetric measurements, as the important thing is to measure, ideally, the entirety of the electrons' energy. Primary readout channels tend to be ionization and phonons, with energy resolutions at the .1% level. A main feature of solid detectors is granularity. This permits the total detector mass to be split into several detectors, which allows for the total mass to be increased in steps as more material becomes available. A downside is that the production and operation of several (up to hundreds) of detectors is cumbersome and challenging [16].

Generally speaking, event reconstruction in $0\nu\beta\beta$ decay experiments can use four primary channels: ionization, phonons, scintillation and Cherenkov light. Germanium detectors use

ionization. Energetic charged particles traversing Ge lose energy due to ionization processes in which charge carriers, electrons and holes, are produced. These charge carriers flow along the electric field and are read out as an electrical signal. In semiconductors, the number of produced charge carriers depends on how much energy is lost by the incident particles compared to the mean energy necessary to create an electron-hole pair. The more energy is deposited in the semiconductor, the more charge carriers are produced. The intrinsic energy resolution is determined by the variance in the number of charge carriers, which exhibit sub-Poisson fluctuations characterized by the Fano factor, F. The best resolution to measure the deposited energy E_{dep} is:

$$\mathrm{FWHM} = 2.355\sqrt{\mathrm{FwE_{dep}}} \tag{2.1}$$

Here, w is the mean energy required to produce a charge carrier. The energy resolution also depends on the efficiency of the charge collection, which depends on the detector technology. In the ionization channel, charge collection tends to be slower than the electronic readout response. This allows the charge arrival time information (for instance, via the pulse shape of a waveform) to be used to discriminate different types of event topology [16].

2.1.1 Production & The Czochralski Method

Germanium has a long history of use in industrial applications, starting in the 1940s in developing better crystal rectifiers for use in radars. The first transistor was built at Bell Labs in 1947, using one of the first germanium crystals with good semiconductor properties. To develop more reliable transistors, high purity germanium without grain boundaries was required. Teal and Little of Bell Labs developed the growth of single crystals using the crystal pulling technique named after Czochralski, who had used this method of pulling wires from a melt (see Fig. 2.2) in the early 20th century to determine the crystallization rate of alloys. The first such single Ge crystal was grown in 1948 [32].

^{76}Ge has a natural abundance that is rather small compared to other $0\nu\beta\beta$ decay nuclei, of circa 7.8%. As such, to use ^{76}Ge in a HPGe detector and set competitive $0\nu\beta\beta$ decay limits,

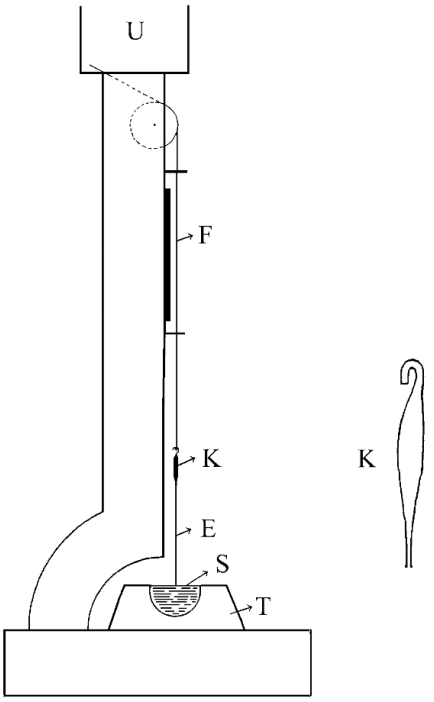

Figure 4. Czochralski's crystal drawing apparatus of 1916.[5] The clockwork motor (U) draws the single crystal (E) out of the melt (S) which sits in the crucible (T). The seed crystal resides in a capillary (K), which is attached to a thread (F). The capillary (K) is shown on the right-hand side of the apparatus, magnified six times. From a diagram by J. Czochralski.[5]

Figure 2.2: Schematic of the Czochralski method [61].

one must start by enriching germanium in ^{76}Ge. The germanium used for the MAJORANA DEMONSTRATOR is enriched in Zelenogorsk, Russia at the Electrochemical plant, using the large centrifuge. Specifically, GeO_2 is processed until reaching $> 87\%$ in ^{76}Ge. The enriched GeO_2 is held in shielded storage and sent to Oak Ridge by land and sea, in a special heavy iron shield, in a steel shipping container, to minimize cosmogenic activation [63]. Germanium oxide is then reduced to pure germanium semi-metal using a reduction furnace, and is then zone refined, to reach a specific resistivity and level of impurity.

The Czochralski method requires an extremely pure sample of germanium as the input. The concentration of impurities in the starting material is typically below 10^{13} cm^{-3}. The pu-

rification of germanium is mainly based on segregating impurities at a solid-liquid interface. The semi-metallic material is melted, and most impurities in germanium can be dealt with by using controlled directional solidification. In this process, impurities in the germanium sample remain in the liquid phase and are swept towards the end of the bar during solidification. Therefore the blocks of solid germanium that result are purified. This is also known as zone refining, and ingots that have been zone refined are the starting point for growing electronics-grade single Ge crystals. A seed crystal is introduced and is pulled slowly to form a crystal boule, which is a single crystal. Current technology allows the production of single crystals several kilograms in mass. The growth process depends strongly on the properties of the starting material as well as the type of applications for which the crystal is grown. For radiation detection, such as gamma-ray detection, germanium crystals should have an electrically active impurity concentration on the order of 10^{10} cm^{-3} or less. This requires the use of specialized techniques, such as working under a hydrogen atmosphere to prevent incorporation of SiO_2 precipitates, and tools, such as a fused silica crucible. Dislocation-free HPGe grown in hydrogen is unsuitable due to a divacancy-hydrogen complex with a concentration of 10^{11} cm^{-3}, which itself acts as a trap. Dislocation density is crucial for detector fabrication. It is measured by counting the number of dislocation lines that thread a unit area of surface. Alternatively, it can also be defined in terms of the total dislocation length per unit volume – in either case, its unit is inverse area [101]. If the dislocation density is above 100 cm^{-2}, the divacancy-hydrogen trap concentration becomes good enough for detector performance. If the dislocation density rises above 10000 cm^{-2}, the dislocations act as traps themselves. Matching the dislocation density over a large volume on the order of 500 cm^3 is one of the biggest challenges in crystal growth [32].

The germanium is made into a detector blank by cutting, machining and grinding, in a lubricant bath. To make germanium detectors, etching the surfaces of the detector crystal blanks is required, resulting in a loss of on the order of a few percent of the germanium per etch. The germanium crystal is etched with nitric and hydrofluoric (HF) acid solutions. All but one surface of the crystal blank is diffused with lithium, in order to form the p–n junction

of the p-type germanium semi-conductor diode. These steps produce a significant quantity of scrap germanium that must be recovered, especially in the case of enriched material [63].

As will be discussed in more detail later, $0\nu\beta\beta$ decay experiments are conducted in deep underground laboratories, in order to be shielded from cosmic rays and muons. These cosmic rays produced in the atmosphere can induce spallation as they interact with the germanium material or detector. The emitted nucleons can be highly energetic and can cause secondary processes such as fission or secondary spallation. They can activate cosmogenic nuclei in the material prior to deployment underground. Cosmogenic nuclei are troublesome for experiments as they can mimic $0\nu\beta\beta$ decay processes, so it is common practice to minimize and track the above-ground exposure of all materials near or in the detector. Germanium detectors are produced in commercial facilities, rather than fabricated underground on-site, as is the case with some other components. As such, the exposure of the germanium material to cosmic rays and muons must be carefully managed. For instance, despite the longer travel time, germanium and the detectors themselves, should be shipped by land and sea rather than by air, due to the much lower flux at sea level as compared to 10 km in the air. To reduce short lived isotopes, detectors are usually also stored underground before being installed in an experiment [16].

2.1.2 Operational Characteristics

The germanium band-gap is about 0.7 eV, meaning that at room temperature the operation of Germanium detectors is impossible, as the thermally induced leakage current would be too big. Germanium detectors also require high voltage (HV) and a readout amplifier. Therefore, germanium detectors are operated at cryogenic temperatures, usually around the LN temperature of 77 K. This is because the noise associated with the leakage current should not spoil the germanium detectors' intrinsic superb energy resolution [84].

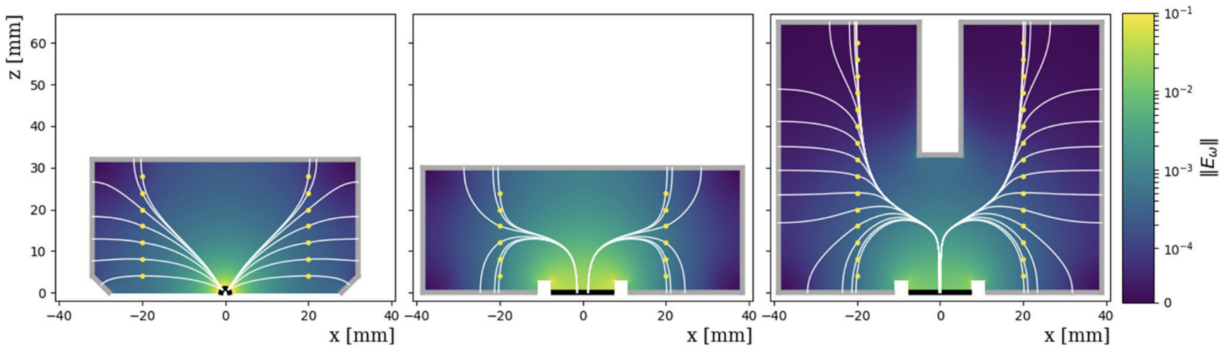

Figure 2.3: Schematic of different detector geometries used in the MAJORANA DEMONSTRATOR. From left to right they are the PPC, BEGe and ICPC geometries. The weighting potential is shown for each type of detector, as well as charge collection trajectories for various energy deposition locations. The n+ contacts are shown in gray and the p+ contacts are shown in black. Figure from Ref. [65].

2.1.3 Properties of the P-type Point-Contact (PPC) Detectors

There are several geometries of HPGe detectors used for the MAJORANA DEMONSTRATOR. These are the Broad Energy Germanium (BEGe), P-type Point Contact (PPC) and Inverted Coaxial Point Contact (ICPC) detectors. They are operated in a very similar manner but have different geometries of their passivated surfaces and contacts. BEGes are natural germanium detectors, and ICPCs have only functioned as part of MAJORANA DEMONSTRATOR in the last few months of running. Therefore, the PPC detectors are the most important to describe. The current mass of germanium detectors ranges from 1-3 kg, with <1 kg being typical for a MAJORANA DEMONSTRATOR PPC detector and 2-3 kg being typical for the newer inverted coaxial point contact (ICPC) detectors.

Charge Collection and Pulse Shape

Electrons and holes produced within the crystal by ionization drift along the electric field, inducing a current. The current integral is proportional to the energy deposited within the detector, and its time-structure carries information about the event topology [16].

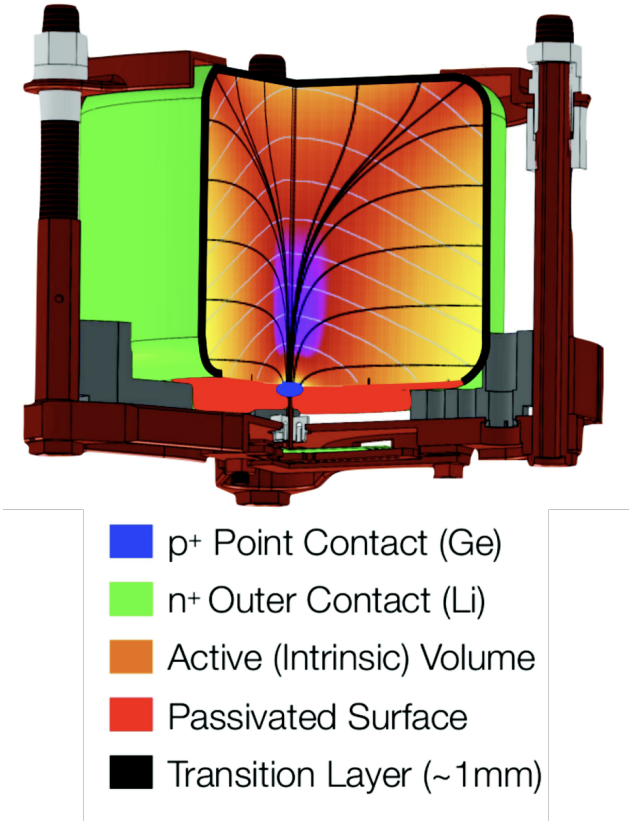

Figure 2.4: Schematic of a MAJORANA DEMONSTRATOR PPC detector. Also depicted are the drift paths of the electrons and holes.

Ultimate time resolution depends critically on the overall average rise time and significant variation of the pulse shape between events [84]. Charge collection should occur in the minimum possible rise time, in order to maximize the performance of pulse shape discrimination techniques. To obtain this, we should have sufficiently high E fields so that the saturation drift velocity is achieved for both holes and electrons, i.e. a field of $\sim 3 \times 10^5$ V/m.

The shape of the pulse rise in HPGe detectors can vary substantially event to event. These variations are due to the positions at which the charge carriers are formed in the detectors. There will be separate collection times for holes and electrons, due to the different lengths the charges must travel in the detector to the respective contact.

Implications of the Shockley-Ramo Theorem

It is important to note that the signals from HPGe detectors are due to the movement of the charge carriers created by the incident particle or radiation. The detector signal is not formed only after the charge carriers deposit their energy on the electrodes. Rather, the pulse begins to form immediately once the charge carriers start their motion towards the electrodes. There is no delay before the onset of the pulse due to the time required for the carriers to arrive at the collecting electrode [107] [84].

The Shockley-Ramo Theorem states that the current induced on the electrode is due to the "instantaneous change of electrostatic flux lines which end on the electrode". Specifically, the equation:

$$i = q\vec{v} \cdot \vec{E_v} \tag{2.2}$$

calculates the instantaneous current received by a given electrode due to a single electron's motion [107]. Here, $\vec{E_v}$ is the component of the weighting, or Ramo, field in the direction of motion of the particle at its instantaneous position. This is emphatically not the electric field – in 2-electrode configurations the weighting field and electric field are the same shape, but they are not identical! More specifically, it is the component of the field with the charge removed, the electrode in consideration at unit potential and, other electrodes (if any) at zero potential. It depends only on the geometry of the electrodes and determines how the charge's motion couples to a specific electrode to induce a signal. $\vec{v}$ is the instantaneous velocity of the particle. q (sometimes, e) represents the charge of the particle. An implication is that the integral of i along the entire drift is equal to the charge q, i.e. the entire drifting charge is collected. It is important that i includes the motion of both holes toward the low-V contact and electrons toward the high-V contact. As a one sentence summary, the summed, integrated instantaneous current i for all electrodes in a given detector creates the pulse-shape we use to calculate parameters such as the energy of the event. Applying the Shockley-Ramo theorem, detectors with the PPC electrode configuration have a sharp response near the point contact, which is beneficial for discriminating $0\nu\beta\beta$ events.

2.1.4 Calibration

To do spectroscopy with germanium detectors, the pulse height scale must be calibrated in terms of energy such that the peaks appearing in the spectrum can be identified. Generally, a calibration gamma-ray source (e.g., from ^{208}Tl in the ^{228}Th chain) is used to create peaks at known energies. In principle, one should use sources which have peaks at energies similar to the region of interest, i.e. for $0\nu\beta\beta$ around 2039 keV. Due to typical spectrometer non-linearity over the full energy range, it is also ideal if the calibration source has multiple peaks throughout the measured energy range [84].

Performance

HPGe detectors' design enables high efficiency in detecting $0\nu\beta\beta$. Another advantage of germanium detectors is having the best energy resolution (FWHM) in the region of interest around $Q_{\beta\beta}$. Due to the presence of the electrodes on the detector surface, the active volume of such a detector is reduced to ∼90%. It can also lead to energy loss for some fraction of $0\nu\beta\beta$-decay events, although this efficiency is also typically ∼90%. Another loss in efficiency is due to $0\nu\beta\beta$-decay tagging in the analysis, limited by the pulse-shape, or current time-structure, methods which discriminate signal and background, ϵ∼80-90% [16]. Performance, specifically in the MAJORANA DEMONSTRATOR, will be covered in more detail in later chapters.

2.2 Overview of the MAJORANA DEMONSTRATOR

The MAJORANA DEMONSTRATOR is operating at the 4850-foot level (4300 m.w.e.) of the Sanford Underground Research Facility (SURF) in Lead, South Dakota. Among the specific goals of this experiment are to:

1. Demonstrate a path forward to achieve a background rate at or below 3 counts/(ROI-t-y) in the 4 keV region of interest (ROI) around the 2039 keV Q-value for ^{76}Ge $0\nu\beta\beta$ decay.

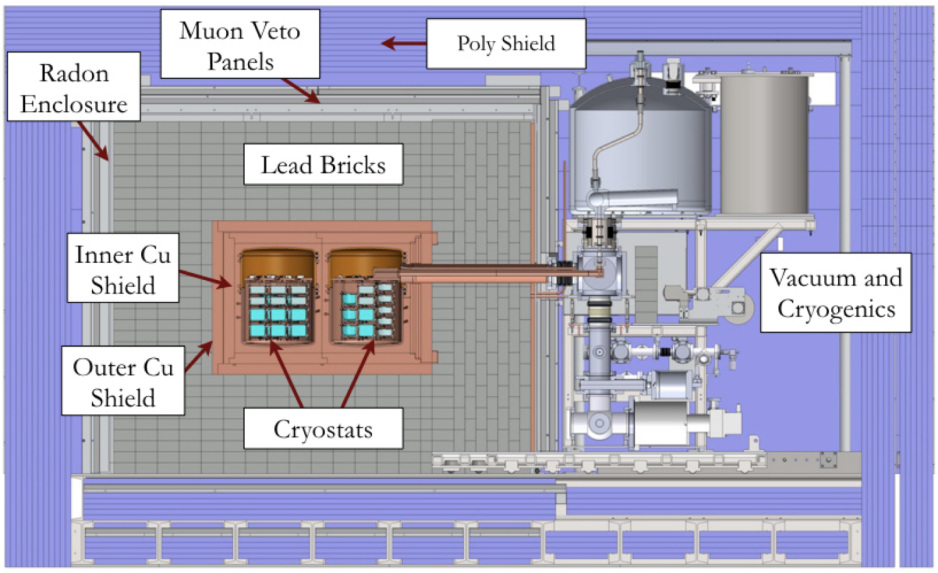

Figure 2.5: Schematic of the MAJORANA DEMONSTRATOR Experiment

2. Show technical and engineering scalability toward a tonne-scale instrument.

3. Set limits on the half-life of $0\nu\beta\beta$ and on $m_{\beta\beta}$ which are competitive with leading searches such as GERDA [14].

4. Search for additional Beyond the Standard Model (BSM) physics.

The MAJORANA DEMONSTRATOR is searching for $0\nu\beta\beta$ in ^{76}Ge, taking advantage of the useful properties of the isotope [4].

This chapter will describe the design of the MAJORANA DEMONSTRATOR in detail; a schematic of the experimental set-up to guide the description can be seen in Fig. 2.5.

2.2.1 Hardware

The MAJORANA DEMONSTRATOR uses two shielded modules of high purity germanium (HPGe) detectors, with each module operated in a separate vacuum cryostat (Fig. 2.6). ~30 kg of the detectors are enriched to 88% in ^{76}Ge.

Figure 2.6: A MAJORANA DEMONSTRATOR cryostat. The white tubing is the track used for deploying the calibration source during calibration runs. The cross arm is visible on the upper left of the cryostat, penetrating the shielding structure from the left.

The cryostats and the structural support of the arrays are constructed from ultra-low background underground electroformed copper (UGEFCu) produced on-site at SURF and specially-selected plastics. Examples of UGEFCu parts are given in Figs. 2.7, 2.8. The success of, among others, double-beta decay experiments depend on the purity of materials selected for construction. While some materials of the desired purity can be found, most must undergo significant purification to create ultra-pure material before being installed near HPGe detectors. Sources of background can include cosmic rays but also radioactivity in materials used in assembling the experimental setup, e.g. in the detectors or shielding. Several previous measurements indicated that a full-spectrum count rate of detectors below 500 m.w.e (meters of water equivalent), with increasing depth, did not result in a further decreasing count rate [59]. Therefore, materials in the detector and shield were concluded to

have been the dominant background sources. Reduction of impurities in copper is crucial as this material can be the most massive component in many ultra-low background cryostats. Most materials which are desirable for constructing detector parts (e.g. cryostats) suffer from both activation due to reactions with secondary neutrons produced by cosmic rays as well as contamination with U and Th during production, handling and transit. Although some starting Cu anode material, such as from Outo Kumpu, has relatively low radioactivity, 0.9 ± 0.3mBq/kg in ^{232}Th, compared to 0.7 ± 0.5mBq/kg after the Hoppe electroforming method referenced [59], the cosmogenic activation is still problematic. Copper suffers activation due to reactions with cosmic-ray-produced secondary neutrons, like ^{60}Co, as well as contamination with ^{238}U and ^{232}Th during the production and handling processes. The most damaging is the creation of ^{63}Cu(n,α)^{60}Co in copper, due its roughly 5 year decay half-life and its 2.5 MeV Q_β. As such, it is crucial to produce radio-pure copper underground, ideally on site, as cosmic-ray exposure effects are long lived, and preventing activation is complex. A schematic (Fig. 2.10) and a picture of the actual electroforming apparatus (Fig. 2.9) are shown.

Low background front end electronics, as well as signal and high-voltage (HV) cables and connectors were developed for the experiment [31, 6]. The background goals for the MAJORANA DEMONSTRATOR require HV cables to be extremely low-mass miniature coaxial cables, produced by Axon Cable in collaboration with MAJORANA DEMONSTRATOR. They have an outer diameter of 1.2 mm, with a conductor diameter of 0.152 mm and an inner dielectric diameter of 0.77 mm. The central conductor is made of bare Cu, with the helical ground shield being 50AWG Cu, whereas the inner dielectric and outer jacket are FEP (Fluorinated Ethylene Propylene). The HV feedthroughs must be small as to integrate 80 connections in 4 flanges. However, these cables and connectors must be able of supplying voltages up to 5 kV to the HPGe detectors, without discharging, which can damage the front-end electronics or mimic detector signals. The occurrence of surface micro-discharges at imperfections and interfaces, due to high voltage required a unique HV component testing program which confirmed the good performance of the Axon connectors for the purposes of MAJORANA

Figure 2.7: Underground Electroformed Copper Detector Holder, with a natural Broad Energy Germanium (BEGe) detector. The low-mass front end (LMFE) is visible near the center of the three-pointed copper structure, held in place by a spring clip. The connector cables are visible towards the top.

Demonstrator. During testing, a configuration with a low expected micro-discharge rate and no breakdowns has been found and was subsequently used in the Majorana Demonstrator. In addition, the micro-discharge effect in the Majorana Demonstrator will not affect physics measurements as the polarity of waveforms is flipped for discharge events compared to physics events [6]. A low-noise resistive-feedback front-end electronics assembly, known as the low-mass front-end (LMFE), was produced to be used with p-type point contact (PPC) detectors in low background experiments. The assembly was designed to

Figure 2.8: Underground Electroformed Copper Parts, held in a nitrogen atmosphere, awaiting the assembly of detector units.

have low mass, and low radioactivity, allowing it to be mounted close to the detector, to match the low PPC capacitance and have minimal stray input capacitance. It is fabricated on fused silica and its parts are a low-noise JFET, an amorphous Ge thin film and a feedback capacitor based on the stray capacitance between traces. The signals from the source, drain, and feedback of the LMFE are connected to a low noise discrete preamplifier. The LMFE is mounted in a tension-fit copper clip and the electrical connection to the detector is established via a 10-20 mm long copper pin. The purpose of the LMFE board is to provide the first stage of amplification to signals coming from the HPGe detectors [31].

The modules are placed inside a multi-layered low-background shield, consisting of, from inside-out, 5 cm of UGEFCu, followed by 5 cm of commercially sourced C10100 copper, followed by 45 cm of high-purity lead. This shield construct is enclosed in a steel radon exclusion box that is constantly purged with liquid nitrogen (LN) boil-off gas. This volume is then surrounded by an active muon veto with nearly 4π coverage [36], which is enclosed in borated and pure polyethylene neutron shielding. The vacuum and cryogenic hardware and control

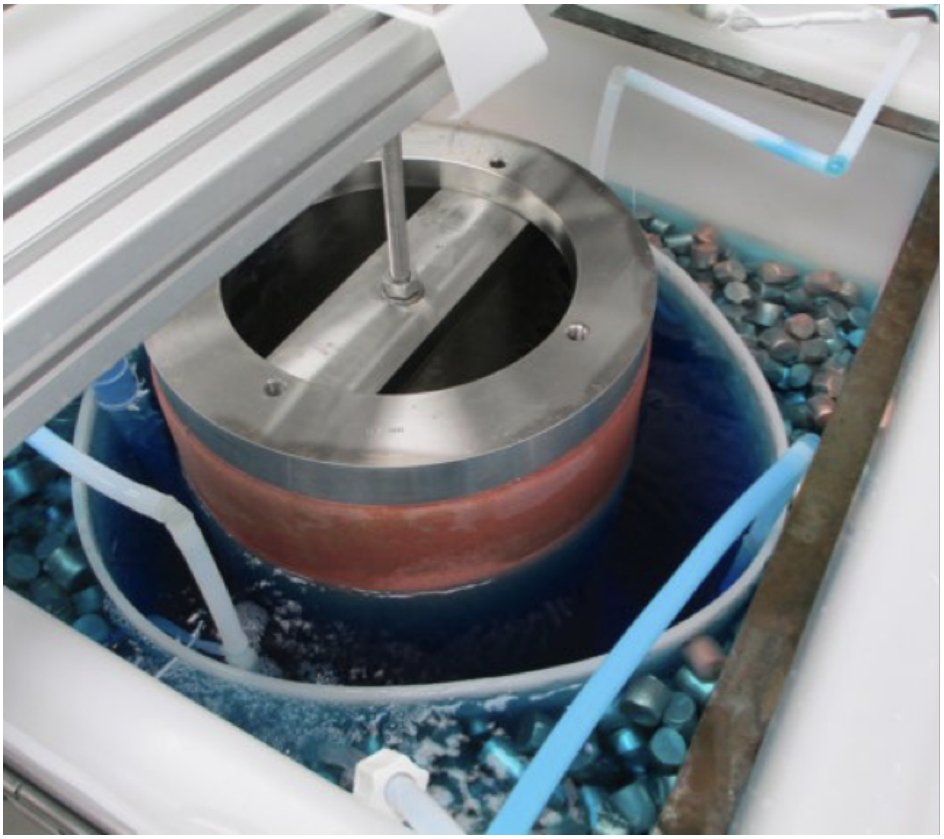

Figure 2.9: Underground electroformed copper (UGEFCu) bath. Commercially available copper nuggets are placed in a bath of copper sulfate and sulfuric acid in constant circulation. Dissolved copper is driven by an electric field onto a mandrel, whereas contaminants in the commercial copper, like U and Th, remain in the bath. The setup is located underground at the 4850' level of SURF to minimize cosmogenic activation.

electronics are placed between the polyethylene shielding and radon enclosure. A shielded crossarm penetrates the shield and connects these to the cryostats and detector electronics. In addition, a ^{228}Th line source is used for each module, with a penetration and helical track enabling its deployment for detector calibration [8]. To ensure low backgrounds, components of the modules and shield were subject to an extensive radioassay campaign [7]. To avoid cosmogenic activation, time spent on the surface by detectors and detector components was minimized, and their locations were logged in a part tracking database [5].

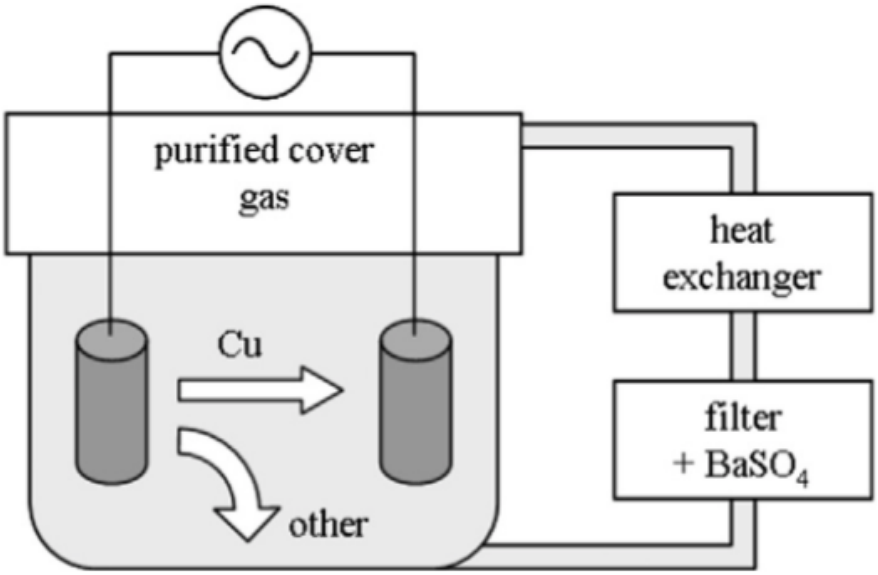

Figure 2.10: A schematic about what is really going on in the copper bath. The electroplating process shows the movement of copper ions across a bath of copper sulfate and sulfuric acid, under the influence of a low voltage. The $BaSO_4$-seeded filter is used to scavenge Ra from the solution [59].

Detector signals are digitized using ADCs developed for the GRETINA Experiment [62], with 14 bit resolution and a sampling rate of 100 MS/s [23]. Each detector is assigned two ADC channels, one for a high gain signal with dynamic range up to $\sim$3 MeV, and one for a low gain signal with $\sim$1/3 as much gain. The high gain channels are used for most analyses; use of low gain channels is reserved for high energy events such as muons and alpha decays. An internal trapezoidal filter is used to trigger each ADC channel independently, upon which 2016-sample waveforms are recorded to disk, giving a $\sim$20 μs acquisition window. The digitizers can be run in multi-sampling mode, which presums every 4 samples after the rising edge of a waveform, extending the sampling window to 38.2 μs. A pulser signal is fed to the front-end electronics every 10 s in order to monitor the electronics stability and detector livetime. Many detector channels spent significant periods of time non-operational

due to connection failures. Data was collected nearly continuously in 1 hr long runs while the detectors were operating.

2.2.2 History

The enriched detectors of the MAJORANA DEMONSTRATOR took data between 2015 and 2021, when they were removed for deployment in the Large Enriched Germanium Experiment for Neutrinoless Double-Beta Decay (LEGEND). An upgrade of a module in 2020 with improved connectors and cabling successfully made all of its detectors operational, and it allowed the deployment of four inverted coaxial point contact (ICPC) enriched ^{76}Ge detectors, with a total mass of 6.7 kg, to study their performance prior to use in LEGEND. The larger detectors in Fig. 2.11 are the ICPCs. The enriched PPCs used for most of the duration of the MAJORANA DEMONSTRATOR experiment have an isotopic abundance of $87.4 \pm 0.5\%$ in ^{76}Ge while the ICPCs' isotopic abundance is $88.0 \pm 1.0\%$ [52].

2.2.3 Signal Amplification and Digitization

Whenever an ionization event happens within the HPGe detectors, a cloud of electron-hole pairs is produced. Under the kilovolt-level high voltage (HV), electrons drift towards the n^+ contact and holes towards the p^+ contact. This movement induces an electric current at each contact, the integral of which is the induced charge signal. We can measure this charge signal, which is directly proportional to the energy of the event. The LMFE board, a charge-sensitive amplifier resistively coupled to ground, measures the signal. To minimize electronic noise and crosstalk, the LMFE must be mounted as close to the detectors as possible – specifically, on each detector mount. Cables connect the LMFE with a second-stage amplifier outside the lead shield. This amplifier produces two different gains of signal, low-gain and high-gain, i.e. there are 2 channels per detector. The high gain channel is mostly used for $0\nu\beta\beta$ studies as it has better noise characteristics at lower energies; however, it saturates above 3-4 MeV. The low gain channel is used above this energy, for instance in alpha studies. The output from the circuit is a waveform of characteristic peak-shape: an

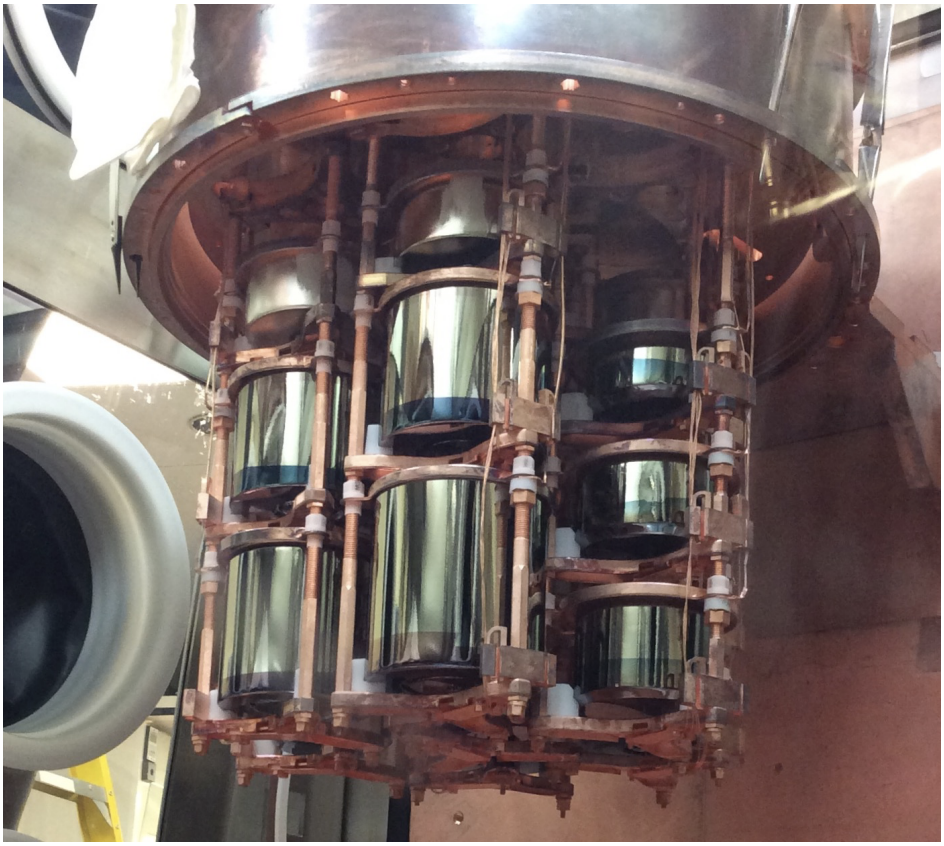

Figure 2.11: ICPC detectors deployed in Module 2

exponential impulse with a decay time on the order of 72 microseconds. The LMFE board also features a capacitively-coupled pulser. This can inject artificial pulses of known energy, which can be used to measure the stability of the electronics [40]. Example waveforms can be found in Fig. 2.12.

Signals are then digitized using GRETINA cards developed for the eponymous experiment. The MAJORANA DEMONSTRATOR signals are digitized at a frequency of 100 MHz, with 14 bits of precision. Each channel is triggered using an energy threshold, recording the waveforms. The MAJORANA DEMONSTRATOR digitizers can multi-sample. This means they can have a high frequency of sampling around the rising edge of the waveform to help pulse shape analysis techniques which use the waveform's rise-time. They can utilize a low sample

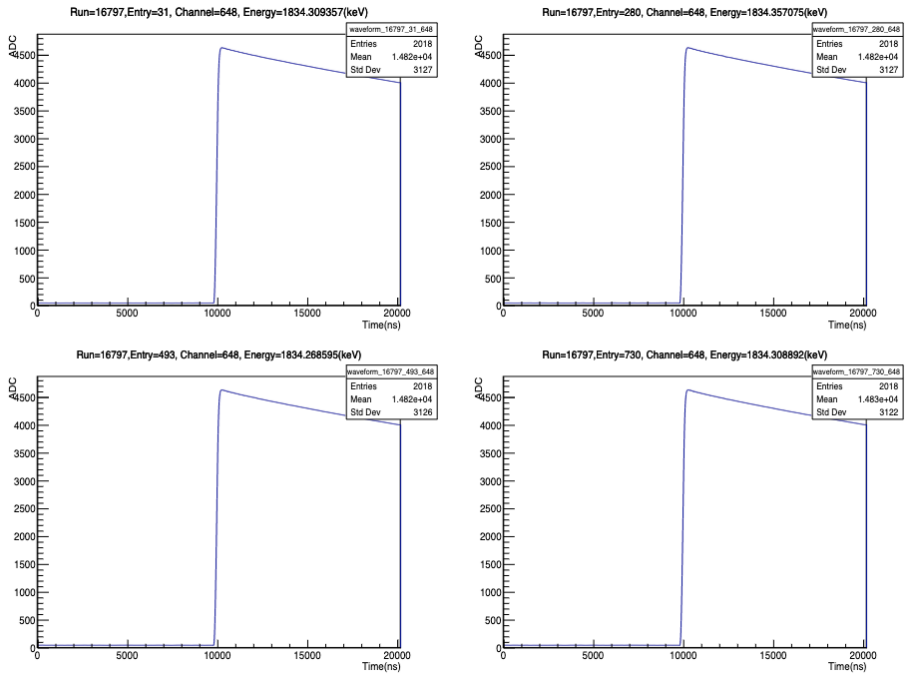

Figure 2.12: Pulser waveforms in run 16797 of DS3.

frequency of a waveform's exponential tail, allowing for longer traces, which can improve techniques based on the exponential tail information. In the MAJORANA DEMONSTRATOR, multi-sampling is either disabled, resulting in roughly 20 μs waveforms with a single sampling frequency of 100 MHz, or enabled with the falling tail sampled at 25 MHz, resulting in circa 40 μs waveforms, depending on the dataset. Each digitized pulse is recorded and then reanalyzed offline.

2.2.4 Performance

Excellent energy performance has been achieved with the MAJORANA DEMONSTRATOR HPGe detectors, including low energy threshold, great linearity, and a FWHM energy resolution (2.5 keV) that is approaching 0.1% at the double beta decay Q-value (of 2039 keV), the best in all $0\nu\beta\beta$ experiments. The MAJORANA DEMONSTRATOR has successfully demonstrated

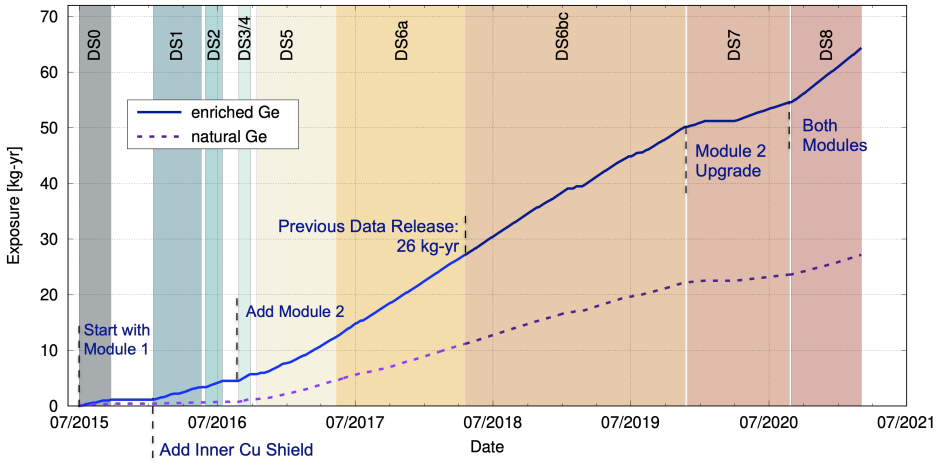

Figure 2.13: MAJORANA DEMONSTRATOR total exposure versus time

the feasibility and advantages of the ton-scale LEGEND project, and it has also been highly productive and competitive in a broad range of physics topics. The experiment reached an exposure of ∼65 kg-yr (see Fig. 2.13) before removal of the enriched detectors for the LEGEND-200 experiment at LNGS.

Chapter 3

MAJORANA DEMONSTRATOR **ANALYSIS**

3.1 Introduction

3.1.1 Basic Features of $0\nu\beta\beta$ decay

The main experimental signature of the $0\nu\beta\beta$ decay signal is a peak at $Q_{\beta\beta}$ in the energy spectrum of the summed electron energies. $0\nu\beta\beta$ decay is a nuclear decay, and therefore a random process obeying Poisson statistics. As the lifetime of $0\nu\beta\beta$ decay is larger than the age of the Universe, $0\nu\beta\beta$ events occur at an effectively constant rate in time across the life of a given experiment. The rate is also proportional to the fraction of active material. Such decays are also expected to be distributed uniformly throughout the active material.

The final state of $0\nu\beta\beta$ decay is a nuclear recoil and two emitted electrons. The electrons' mass is much smaller than that of the nucleus, implying that the nuclear recoil energy is negligible. It follows that the sum of the electron energies is equivalent to $Q_{\beta\beta}$, the available total energy. The electrons deposit their energies in a localized manner, depending on their attenuation length in a given material. For solids, this is on the order of 1-10 mm, and circa 2 mm in germanium. In principle, the kinetic energies and momenta of the electrons, as well as the time and position of the decay, are measurable quantities [16].

Naturally, $0\nu\beta\beta$ decay competes with $2\nu\beta\beta$ decay, for any and all isotopes. In the latter, two electrons and two electron anti-neutrinos are emitted. Since the anti-neutrinos escape without interacting with the detector, the sum energy measured is that of the two electrons, $<Q_{\beta\beta}$.

Measuring the sum electron energy is a necessary condition for discovery: $0\nu\beta\beta$ decay will feature a peak at $Q_{\beta\beta}$, while the $2\nu\beta\beta$-decay mode exhibits a continuum from zero to $Q_{\beta\beta}$. For detectors with good enough energy resolution, with no sources of background,

energy measurement would also be a sufficient condition for discovery. Energy resolution is therefore critical in minimizing the background level around the $Q_{\beta\beta}$.

There also exist emissions originating from the ^{238}U and ^{232}Th decay chains (α, β, γ particles) with energies greater than $Q_{\beta\beta}$. These backgrounds can occur from radioactive particles impinging on the surfaces of our detectors. This is minimized but not eliminated completely. Since the backgrounds occur on the outside of our HPGe detectors, we can use background rejection techniques to reduce their detrimental effect on our analysis.

The event topology of a $0\nu\beta\beta$ decay is clearly defined for each detector technology: an energy deposition contained a volume $O(10)$ mm^3. Depending on the detector's spatial resolution, several particles might be distinguishable. For instance, γ particles have a longer range and often Compton scatter, generating energy depositions at several different locations within a detector or detectors. Muons are found by their interaction with several muon veto plates. $0\nu\beta\beta$ decays are also homogeneous in time and not correlated with other events, unlike, e.g., delayed coincidences between decays of isotopes in the ^{238}U and ^{232}Th decay chains, or the decay in cosmogenically activated isotope occurring soon after a muon passed through the veto.

3.1.2 Basic Details of the Majorana Demonstrator *Analysis*

An offline analysis is performed in the Majorana Demonstrator to calculate waveform energies and pulse shape discrimination (PSD) parameters that will be used to reject likely backgrounds. $0\nu\beta\beta$ events deposit their energy in a small, circa 10 mm^3, volume in the bulk of a detector, so background rejection techniques seek to remove events with different topologies [26]. This chapter contains details about the the exposure, detection efficiency due to data cleaning, containment and PSD cuts, as well as other relevant information about the Majorana Demonstrator analysis.

In the case of the Majorana Demonstrator analysis, which is a rare-event peak search of $0\nu\beta\beta$ decays in a region of interest around $Q_{\beta\beta}$, the lower limit on the half-life is estimated as:

$$T_{0\nu}^{1/2} > \frac{\ln(2) N T \epsilon}{S(C,B)} \tag{3.1}$$

Here, N is the number of ^{76}Ge atoms, T is the measurement time (their product NT being the exposure), and ϵ is the detection efficiency of the experiment. $S(C,B)$ is the upper limit on the number of signal events for C observed events in the region of interest, with B expected background events. This chapter describes how the inputs to the MAJORANA DEMONSTRATOR limit are calculated, i.e. how we arrive at our exposure and efficiency numbers, including their uncertainties, for all datasets from DS0 to DS8. The half-life limit calculations themselves, and detailed descriptions of the methods which can be used, are found in Chapter 5.

3.2 MAJORANA DEMONSTRATOR *Data Sets*

MAJORANA DEMONSTRATOR data are divided into 9 datasets (DSs) based on major changes to the experimental configuration.

- DS0 consists of data collected with only Module 1 prior to the installation of the inner UGEFCu shield, and has a higher background rate than other DSs.

- DS1 began with the installation of the inner shield.

- DS2 started when the DAQ configuration was changed to enable pre-summing and therefore multi-sampling.

- DS3 and DS4 occur simultaneously; after Module 2 was first installed in the shield, data were collected from Modules 1 and 2 using separate DAQ systems, in DS3 and DS4, respectively.

- DS5 began when the modules were integrated into a single DAQ system, and is further subdivided into DS5a, which had elevated noise, DS5b, after optimizing the electrical grounding, and DS5c, in which the blindness scheme was re-started.

Data Set	Part Number	Start Date	End Date	Active Mass (kg)	Exposure (kg yr)
DS0	P3JDY	06/26/15	10/07/15	13.06 ± 0.20	1.13 ± 0.02
DS1	P3KJR	12/31/15	05/24/16	12.63 ± 0.18	2.24 ± 0.03
DS2	P3KJR	05/24/16	07/14/16	11.90 ± 0.17	1.13 ± 0.02
DS3	P3KJR	08/25/16	09/27/16	12.63 ± 0.18	0.96 ± 0.01
DS4	P3LQG	08/25/16	09/27/16	8.72 ± 0.12	0.26 ± 0.00
DS5a	P3LQK	10/13/16	01/28/17	20.38 ± 0.30	2.74 ± 0.04
DS5b	P3LQK	01/28/17	03/17/17	20.38 ± 0.30	1.75 ± 0.03
DS5c	P3LQK	03/17/17	05/11/17	20.38 ± 0.30	2.34 ± 0.03
DS6a	P3LTP	05/11/17	04/18/18	20.38 ± 0.30	14.74 ± 0.21
DS6b	P3LTP	04/18/18	11/27/18	18.76 ± 0.27	9.78 ± 0.14
DS6c	P3LTP	11/27/18	11/27/19	18.76 ± 0.27	13.25 ± 0.19
DS7	P3N99	11/27/19	08/28/20	12.04 ± 0.18	4.44 ± 0.07
DS8 (PPC)	P3NF3	08/29/20	03/03/21	18.88 ± 0.27	6.41 ± 0.09
DS8 (ICPC)	P3NF3	08/29/20	03/03/21	6.08 ± 0.09	2.74 ± 0.04

Table 3.1: A summary of the MAJORANA DEMONSTRATOR's data sets. The enriched exposure includes all run/channel selection, environmental cuts, and hardware dead time as described in Section 3.12. The active masses include all detectors used at any point in the dataset, so they are in fact upper limits of the active masses of each dataset.

- DS6 started when multi-sampling was re-enabled. The border between DS6a and DS6b is somewhat artificial – DS6a is the subset of DS6 used in our 26 kg-yr analysis published in 2019. DS6b is newly analyzed in 2022. Between DS6b and DS6c there was a GRETINA card swap.

- DS7 consists of data collected in Module 1 while Module 2 was upgraded.

- DS8 contains data with both Modules after the upgrade completed, and due to differences in the analyses, is grouped into DS8P for PPCs and DS8I for ICPCs.

3.2.1 Blindness

The MAJORANA DEMONSTRATOR's blindness scheme is described in detail in [54]. During operation, where the proper data acquisition (DAQ) bit is set, roughly 75% of the exposure at all energies was blinded, by restricting access to data files, with cycles of 31 hrs of open data followed by 93 hours of blind data. Calibration data and data in non-standard configuration remained open data.

In the standard MAJORANA DEMONSTRATOR analysis, unblinding features three stages.

1. Open outside the background estimation window (1950-2350 keV) for data quality checks.

2. Open the background estimation window except for a 10 keV region centered at 2039 keV.

3. Open the 10 keV window, completing the unblinding procedure.

At each stage, events with multiplicity, defined and discussed in Section 3.5, greater than one and events with energy less than 100 keV are left blind to allow other physics analyses to utilize blindness independently (as an example, low energy beyond the standard model physics searches).

DS0, 3, 4 & 5ab featured only open data. DS1, 2, 5c, 6abc, 7 & 8 feature the traditional blinding scheme described in this section.

3.2.2 Location of the data

The DEMONSTRATOR skim files are located on CORI at `/global/cfs/projectdirs/majorana/data/mjd/surfmjd/analysis/skim`. This directory contains subdirectories for each data set's background and calibration data, separately. The skim files for open data using the improved analysis routines described in the following sections are located in a directory named `GAT-v02-16` within each background or calibration data set's skim file directory. The skim files produced in each of the stages of unblinding are stored in a separate directory at the path above with `_blind`, `_blind_b1`, and `_blind_b14` appended to the name of the data set for the first, second, and final stages of unblinding respectively.

3.3 Containment

Some fraction of $0\nu\beta\beta$ events will have their energies degraded due to Bremsstrahlung, fluorescence, x-rays, or electrons which escape the active volume of the detector. We define the efficiency for $0\nu\beta\beta$ events depositing full energy within the active volume of a single crystal as ϵ_{cont}. The containment efficiency for each data set is listed in Table 3.2. The treatment of the containment efficiency largely followed the same procedure as previous MAJORANA DEMONSTRATOR analyses. To calculate ϵ_{cont}, it is assumed that a $0\nu\beta\beta$ decay originating in the active region of a detector will deposit at least some energy in that detector. This is supported by 200000 simulations. ϵ_{cont} varies by detector due to the size and geometry of the detector. For each detector, ϵ_{cont} was calculated by a simple division of the number of events depositing an energy between 2030 keV and 2050 keV in that detector by the total number of events depositing energy within that detector whose primary decay occurred within the active volume of that detector. The energy width is large enough to ensure 100% of $0\nu\beta\beta$-decay events fall within it. Systematic uncertainties exist, for each detector, due to differences between detector masses and models used to calculate ϵ_{cont}, such as the rounded

corners of the detectors or variation in the radius of a detector along its length. The difference in mass between the physical detector and the model is also considered. There is a linear relationship between the mass and ϵ_{cont}, which is used to correct for the aforementioned differences in active mass. The uncertainty in the active mass of each detector adds to the systematic uncertainty of the efficiency. It is multiplied by the slope of the linear fit. The mass-weighted average of the efficiency across active individual detectors was calculated for each dataset. The statistical and systematic uncertainties associated with the active mass are taken to be independent across detector and thus calculated separately and combined for each dataset. The uncertainty in ϵ_{cont} due to uncertainty in the active mass is likely to be highly correlated across detectors, and considered after calculating the total active mass in each dataset. [66, 38]

3.4 Data Cleaning

Instrumental background events include either events that are solely due to some problem with the detector or readout electronics, or otherwise good physics events that are marred either by such a problem, or by the configuration of the readout. Since instrumental backgrounds may cause problems with the event reconstruction, they must be dealt with at an appropriate stage in the data processing. Removal of instrumental backgrounds is referred to as "Data Cleaning" [82]. There are several types of pathological populations, i.e. specific non-physical waveform shapes, or distortions of good physical waveforms. Each has an associated bit set in the data. Bit 0 indicates a single sample positive spike on baseline. Bit 1 indicates a single sample negative spike on baseline. Bit 2 indicates a single sample positive spike on physics waveforms. Bit 3 indicates a single sample negative spike on physics waveforms. Bit 4 indicates an early trigger. Bit 5 indicates a late trigger. Bit 6 indicates positive saturated waveforms. Bit 7 indicates negative saturated waveforms. Bit 8 indicates pileup [82].

The Data Cleaning (DC) cut is `wfDCBits==0`. It has been shown there is negligible loss of physics events within the ROI due to the data cleaning cut. The sacrifice was estimated

to be less than 0.1% for all data sets [99, 82]. We take the data cleaning efficiency to be $\epsilon_{clean} = 0.999 \pm 0.001$ for all data sets. The efficiency is computed by calculating the fraction of physics waveforms marred by instrumental backgrounds or mis-tagged by the data cleaning cuts (i.e. the sacrifice), and subtracting this value from 1. In principle, this can be done in the split background window of a 360 keV width. However, this can result in a large statistical uncertainty, and so the sideband and superband methods are used. The 1100 keV to 1400 keV sideband window is examined using the same method as for the background evaluation window, but statistical uncertainties on the order of 1-2% are still observed. Thus, the 100-2500 keV superband is used. The additional systematic uncertainty from mis-reconstructed energy can be taken as 0, due to the use of this full superband. The sacrifice and its uncertainties are calculated using the Clopper-Pearson method in ROOT [99, 82].

3.5 Granularity

The granularity, or multiplicity, cut is:

```
mHL==1
```

where mHL is the number of hits (pulses), i.e. the multiplicity, during an event in one detector, seen either in the high or low-gain channel, often referred to in MAJORANA DEMONSTRATOR as the H-or-L mode. The H-or-L mode as implemented in the MAJORANA DEMONSTRATOR analysis allows us to pick the better performing channel for a given detector, without duplication. Then, mHL is added over all detectors to figure out the multiplicity of the event. $0\nu\beta\beta$ events are single-site events, expected to occur at one place in the bulk of one detector, i.e. in the active volume of a single crystal. Full energy, fully contained $0\nu\beta\beta$ events in one detector cannot produce an additional signal in a second detector. This means that no $0\nu\beta\beta$ events should be lost due to the granularity cut. Note that the granularity efficiency is only computed assuming that containment efficiency has already accounted for cases where E has left the detector and hit another detector. Accidental coincidences are the only possible occurrence which can result in events being rejected. In a dataset, the

random coincidence rate between two independent channels is twice of the product of the event rates of the two channels and the coincidence window, $2r_1r_2\tau$. Random coincidence rate between detector i and all other detectors is $2r_i(R-r_i)\tau$. And, between all detectors, the rate is bounded from above by $2R^2\tau$. Here, r_i is the rate for detector i, R is the total event rate for all channels and $\tau = 10\ \mu s$. The average total event rate is therefore 0.042 Hz for high gain channels, and the corresponding rate for high gain channels is 0.0037 Hz. The total random coincidence rate in DS1 is 18 nHz for high gain channels, and 0.14 nHz for low gain channels [122]. Thus, ϵ_{gran}=1 for all data sets.

3.6 Energy

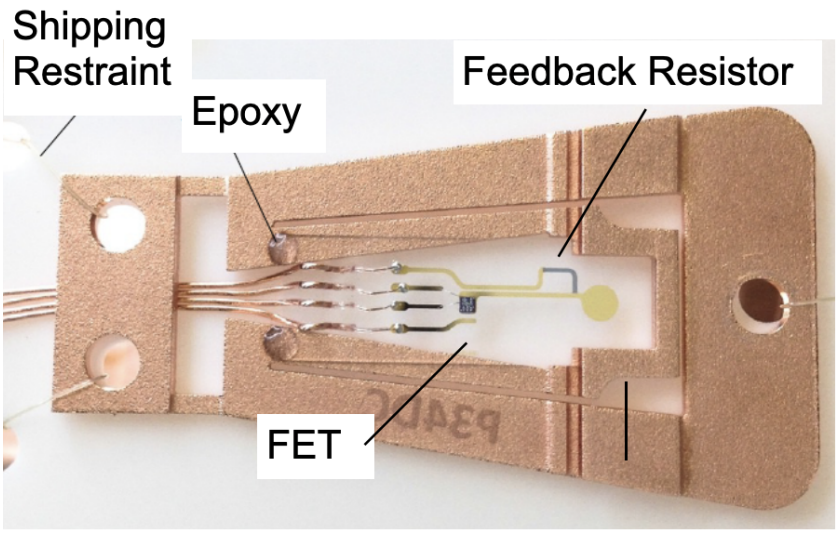

Figure 3.1: Custom low mass front-end boards used in MAJORANA DEMONSTRATOR. It features clean Au and Ti traces on fused silica, an Amorphous Ge feedback resistor, a FET mounted with silver epoxy, underground electroformed Cu and low-background Sn contact pin.

The total charge which is collected on the charge-sensitive pre-amp feedback capacitor seen in Fig. 3.1 is related to the maximum voltage across it, Vmax = Q/C, where the

capacitance is constant. For a detector with linear response, the energy can be estimated from the maximum height of a waveform. However, certain pulse-shaping techniques lead to significantly better energy resolutions. Specifically, event energies are reconstructed from waveform amplitudes which are measured using a pole-zero (PZ) correction and a trapezoidal filter. Each pole-zero corrected waveform is convolved with an impulse function to give a trapezoidal waveform with shaping parameters that optimize the energy SNR. The energy is then measured from the amplitude of the filtered trapezoidal waveform at a fixed time after t_0, called the pickoff time. This start time t_0 of each waveform is calculated by a leading-edge algorithm that uses an asymmetric trapezoidal filter, as depicted in Fig. 3.2, to determine when the waveform is no longer at baseline. By optimizing the PZ time constant for energy resolution rather than a flat-topped waveform, we minimize energy degradation due to charge trapping inside of the detectors [9]. The charge trapping time constant τ_{CT} is varied to maximize the resolution. Energy is calibrated for each detector on a weekly basis by simultaneously fitting the eight most prominent gamma peaks in the ^{228}Tl calibration source energy spectrum. The peaks are fit with a peak shape function and the raw energies are converted into calibrated energies through a linear transformation,

$$E[keV] = \text{slope} \times E[ADC] + \text{offset}. \tag{3.2}$$

Each peak is modelled as the sum of a gaussian component and an exponentially-modified gaussian low-energy tail, plus a background term.

The peak shape is:

$$PS(E) = \frac{1-f}{\sqrt{2\pi\sigma^2}}\exp\frac{(E-\mu)^2}{2\sigma^2} + \frac{f}{2\tau}\exp\left(\frac{\sigma^2}{2\tau^2}+\frac{E-\mu}{\tau}\right)\text{erfc}\left(\frac{\sigma}{\sqrt{2}\tau}+\frac{E-\mu}{\sqrt{2}\sigma}\right) \text{ [91]} \tag{3.3}$$

where μ is the mean of the Gaussian, σ is the standard deviation of the Gaussian, f is the fraction of the amplitude that's taken up by the low charge tail (as in Fig. 3.3) and τ is the decay constant of the tail exponential.

The background terms are:

$$B(E) = b + mE + \frac{H_s}{2}\text{erfc}\left(\frac{E-\mu}{2\sigma^{1/2}}\right) + \frac{3q}{2}\left(E^2 - \frac{1}{3}\right). \tag{3.4}$$

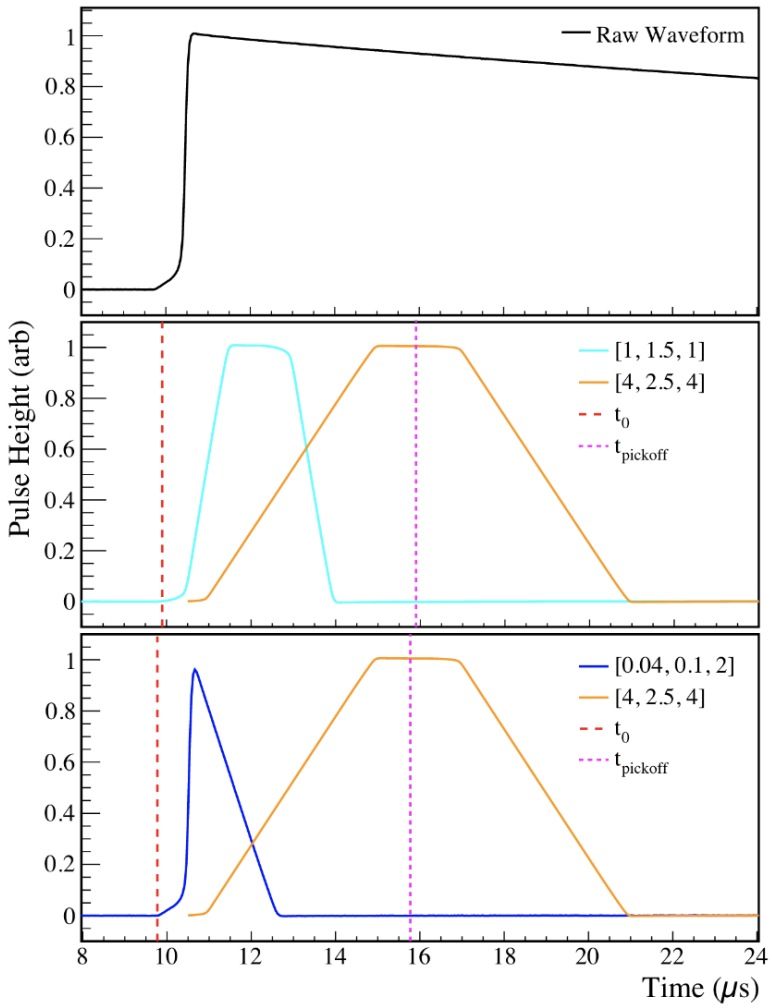

Figure 3.2: Top: a normalized raw waveform. Center: Symmetric leading-edge filtered waveform and energy trapezoidal-filtered waveform. Bottom: asymmetric leading-edge (blue) waveform and energy (orange) trapezoidal-filtered waveform. The red line shows the start time t_0. Fig. taken from Ref. [9].

Here, b is equivalent to the number of counts the flat portion of the background, m is the linear proportionality constant of the background, and q is the quadratic constant of the background. H_s is the height of the step background as a fraction of the peak amplitude,

which appears because we don't expect any physical background events above the peak energy, as there is no physical mechanism which allows for an increase in energy in background events.

The uncertainty on the peak-shape efficiency, which measures how likely an event within the peak is to have its energy reconstructed within the peak energy range, is computed by taking into account the uncertainty in the peak width and the peak position separately. We adjust the peak width by a factor of $1+\alpha$, where α is the fractional uncertainty on the FWHM at 2039 keV. The peak width (FWHM) uncertainty includes both statistical error from the fit and systematic errors from energy non-linearities and drift over time. The peak position is taken into account by shifting the position by uncertainty in the peak position, from energy nonlinearity and gain drift. The energy resolution for each data set and the peak-shape parameters detailed in [91] will be used in setting the $0\nu\beta\beta$-decay limit in Chapter 5.

Another energy improvement was the development of an improved energy estimator. It uses an improved waveform start time (t_0) estimation by applying a quadratic energy scale. This new method relies on using multiplicity ≥ 2 events in calibration to compute the error in start time due to the asymmetric trapezoidal filter, versus energy. The parameters of the correction, which minimize the error at all energies, are then applied to the output of the filter to obtain the new t_0. The correction uses a quadratic calibration with zero y-intercept, meaning no more offset for low energy events and results in a calibration that is linear to within the effect of the digitizer nonlinearity from 0-2614 keV. The quadratic correction reduces the peak shape tail, resulting in an improved ROI efficiency of 1-2% [90]. The energy resolution, for a particular energy, is given by the full width at half maximum (FWHM) of the energy peak. The FWHM is comprised of three components: the electronics noise (Γ_n), the Fano noise (Γ_F), and extra broadening from incomplete charge collection (Γ_q). The complete expression for the FWHM is $FWHM(E) = \sqrt{(}\Gamma_n^2 + E \times \Gamma_F + E^2 \times \Gamma_q)$ The FWHM at 2039 keV for all PPC detectors, including broadening due to gain drift and energy nonlinearities, is 2.52 ± 0.08 keV (excluding DS5a due to high noise), as can be seen in Fig. 3.4. The quadratic and t_0 systematic corrections produce a more linear energy response,

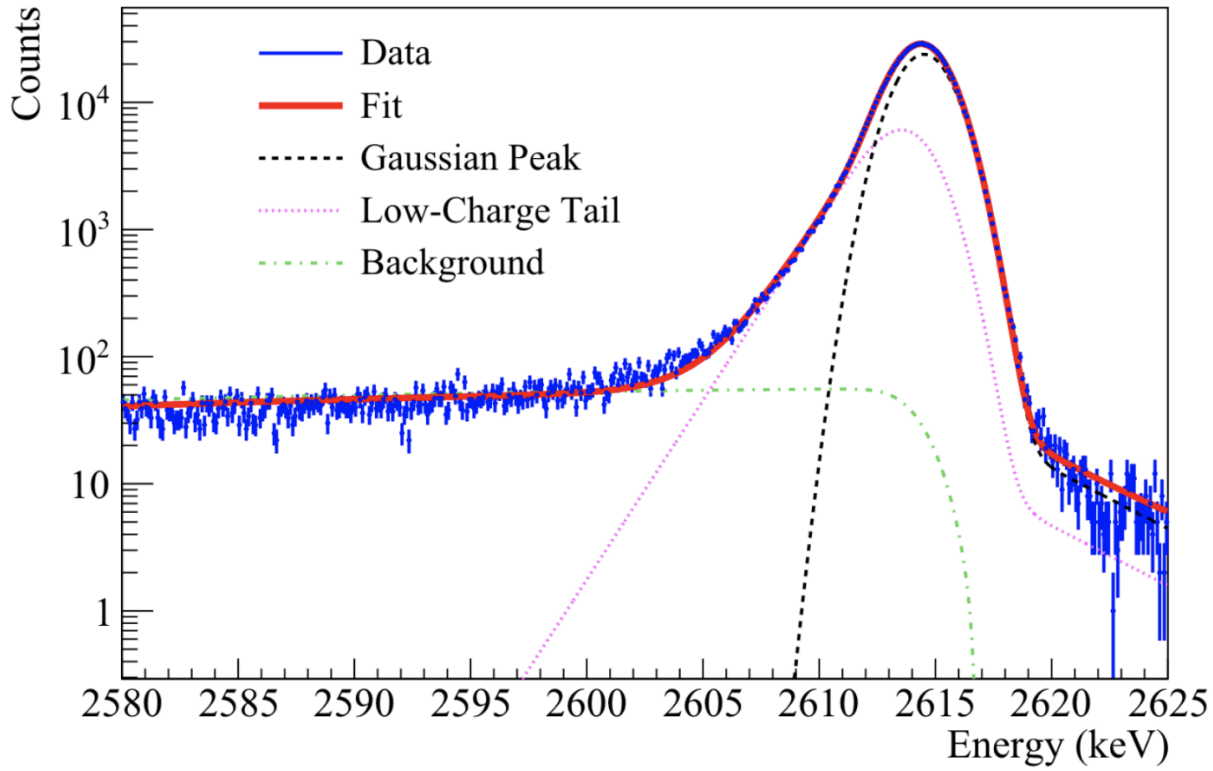

Figure 3.3: A fit of the energy peak shape and the background function to the 2614.5 keV γ peak from 208Tl. The background function is shown in green and the terms of the peak shape are shown in black and magenta. The red line is the sum of the aforementioned terms. Fig. taken from Ref. [21].

especially at energies $<$ 100 keV, and reduce the size of the low energy tail in calibration peaks [2, 21, 26].

3.7 Background Rejection

A crucial aspect of the MAJORANA DEMONSTRATOR's ability to search for $0\nu\beta\beta$ is its excellent ability to do background rejection. $\beta\beta$ events are intrinsically single-site and occur in the bulk of the detector. On the other hand, γ backgrounds often Compton scatter and

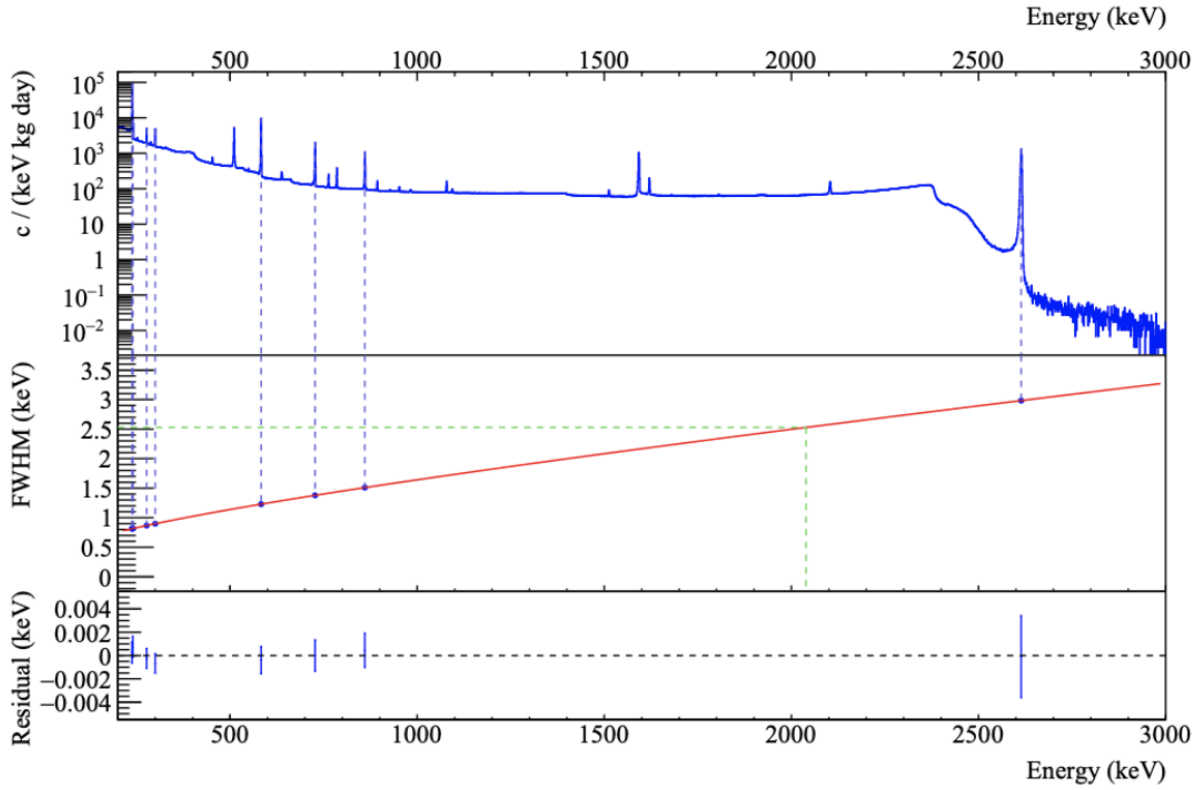

Figure 3.4: Top: combined energy spectrum for all calibrations in DS0 to DS6a for the MAJORANA DEMONSTRATOR analysis result from 2019. Center: exposure-weighted energy resolution for each energy peak used in the calibration fit. Bottom: the fit residuals for the FWHM. Fig. taken from Ref. [21].

present themselves as multi-site events. To select out these events we compare the maximum current amplitude of a waveform, A, to its full charge collection, E, using the parameter we call AvsE.

There are also α backgrounds – these particles are incident on the passivated surface of our detectors and appear with degraded energy. We suspect the α contamination occurs via ^{210}Po decays, ultimately originating from ^{222}Rn exposure on the contact pin and the plastic parts seen in Fig.. 3.5. What happens physically is the charge is trapped at the

passivated surface, and then slowly re-released into the bulk on time scales comparable to the length of our waveforms, resulting in excess slope in the tails of the pulses from those events. This can be seen in Fig. 3.6 and Fig. 3.7. That slope then provides an estimate of this "delayed charge recovery" (DCR), and provide a handle for rejecting events with this feature. Specifically, the bulk events' electronics pole-zero corrected waveforms have a slope of the tail after the rising edge of roughly 0, whereas the α events have a significant positive slope. We use weekly ^{228}Th calibrations to define the efficiency of the DCR parameter – specifically, it retains $98.5\% \pm 0.7\%$ of bulk events. More details about the computation of the DCR parameter can be found in section 3.8 and in Chapter 4.

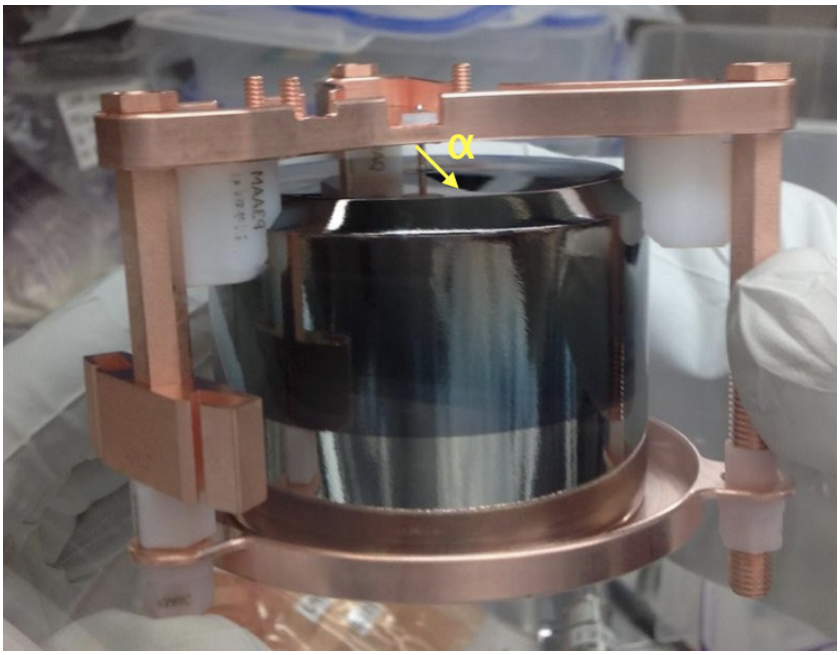

Figure 3.5: Alphas on Passivated Surface are presumed to originate on the contact pin or on the white PTFE plastic parts.

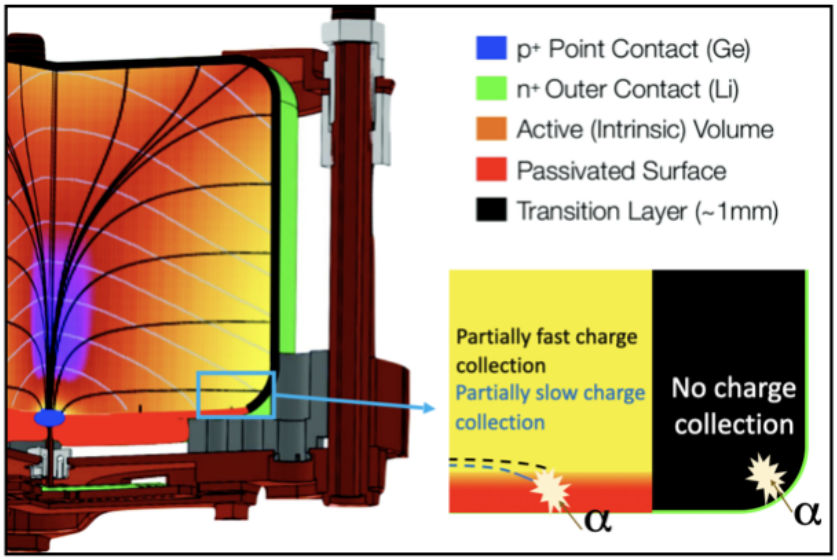

Figure 3.6: Sketch of Alphas on Passivated Surface. Charge is trapped at the passivated surface, and then slowly re-released into the bulk on time scales comparable to the length of our waveforms.

3.8 Surface Alphas & DCR

The work on DCR forms a major part of this thesis and is described in detail in Chapter 4. This is just a brief overview for completeness. The passivated layer of germanium detectors is a fraction of a micron thick, but alphas travel tens of microns. There could also be surface charge that pulls drifting charges to the surface where mobility is lower. There could be extra trapping centers near the surface, as well. There is no full model which explains this effect and each detector appears to exhibit a different effect. There are several test stands (at UW, Munich, and elsewhere) and several articles published [13] [10] [27] to try to measure and understand this effect. More details can be found at the beginning of Chapter 4. The surfaces of PPC detectors surrounding the point-contact have a thin passivation layer. Regardless of the physical reason behind the delayed charge effect, alpha particles impinging on this surface experience increased charge trapping that degrades their reconstructed energy, and can fall

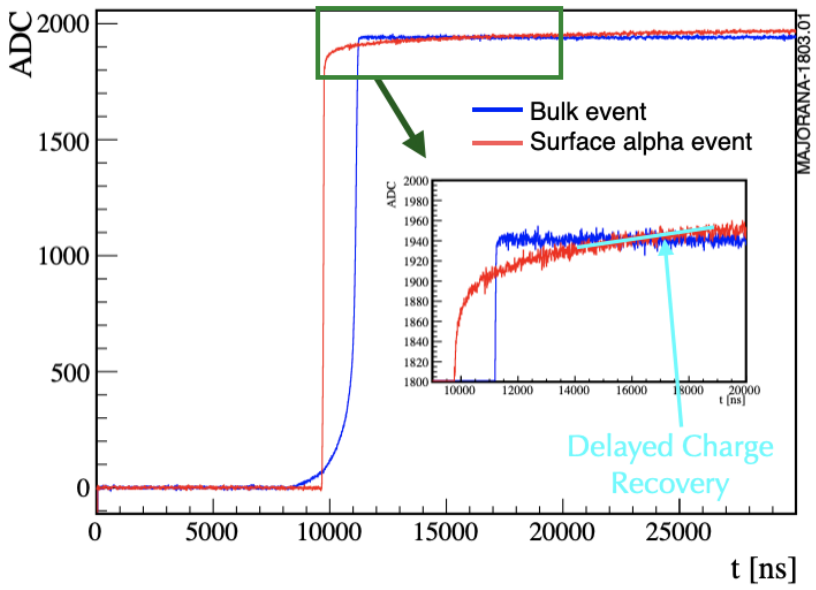

Figure 3.7: Waveforms for bulk and alpha events. Delayed charge is slowly re-released into the bulk on time scales comparable to the length of our waveforms, , resulting in excess slope in the tails of the pulses from those events. From Ref. [21].

into the $0\nu\beta\beta$ ROI, becoming an important source of background. Much of the charge is rereleased slowly, increasing the slope of the falling tail of the waveforms relative to bulk events. Delayed charge recovery (DCR) is a measure of this slope increase, and can be used for surface alpha rejection [26]. For surface alpha events very close to the point contact, the rerelease of charge is smaller; these events may pass the DCR cut. These near-point-contact events, however, have a very fast rise that produce high values of AvsE relative to single-site bulk events. For this reason, a cut on high values of AvsE has been added, as mentioned in Section 3.9 and in Chapter 4.

The DCR cut is tuned for each detector to an efficiency of ∼99% using one calibration. Actually, the overall efficiency is determined, in a separate study, to be $98.5 \pm 0.7\%$ using calibration events in the Compton continuum near the $0\nu\beta\beta$ ROI during all MAJORANA DEMONSTRATOR calibration runs. The calibration runs are also used to estimate the stability

uncertainty on the DCR cut efficiency, as described in Section 4.7 and in Ref. [79], which varies by up to about 1%, as summarized in Table 4.3. Other sources of uncertainty, such as the time-stability systematic and pulse-shape bias, are described in Section 4.7. The process underwent to compute these uncertainties is described in great detail in Section 4.7 and so will not be reproduced here.

Since the DCR parameter is based on the slope between two regions of the waveform (for more details see Sec. 4.3), pileup late on the tail of the waveform can result in a large value of DCR, similar to that of α events which the cut is designed to tag. In the analysis described in [38], the pileup cut as implemented had an unacceptably high rejection of physics events, so it was not used in subsequent analyses. Although inclusion of pileup has a negligible effect on the background data due to its low rate, untagged pileup during calibrations can impact the tuning of the DCR cut, i.e. by essentially adding another waveform on top of the initial waveform, leading to the calculated slope being much higher, resulting in a cut value which accepts more α-like events than expected. As described in [17], the pileup cut has been significantly improved to efficiently cut pileup events with negligible sacrifice of signal efficiency. With the improved pileup cut, the DCR cut is now tuned on events in calibration data after applying the pileup cut, resulting in a cleaner event sample.

The DCR parameter developed in this thesis is described in detail in Sec. 4 and in the DCR unidoc [79]. This parameter is used in the 2023 Majorana Demonstrator PRL final result [26], but different than that used in the 2018 and 2019 Majorana Demonstrator papers [2] [21]. To remove variation over time in the slope of the waveform tail, waveforms are first pole zero (PZ) corrected with a time-dependent decay constant calculated by the Oak Ridge National Lab (ORNL) group in the Majorana Collaboration [76]. In addition to reducing time variation, the resulting distribution of waveform tail slopes has a mean closer to zero and a more symmetric shape. The mean of the distribution for each detector is then shifted to zero while applying a linear correction to account for energy dependence of the DCR slope. The width of the resulting distribution is then scaled such that the DCR parameter is in units of σ such that for a one-sided cut with 99% acceptance the

cut is placed at 2.326.

Figure 3.8 shows the DCR distribution for each detector in the open background data of DS6a. For the previous analysis, the cut value is at 0 which is indicated by the vertical dashed line. With the pileup cut, the time-dependent PZ correction, shifting the DCR mean to 0, and scaling the DCR width to units of σ, we achieve more uniform performance of DCR across detectors. This and other improvements will be discussed in detail in Chapter 4. Detectors which are not present in the left panel of Figure 3.8 but are present in the right panel were previously not used in physics analyses (also referred to as "veto-only") for the entirety of DS6a.

A drift time dependence to the DCR parameter has been found in previous studies. Calibration data for each detector in each data set is used to fit a linear function to the two-dimensional distribution of DCR vs drift time. The angle at which the distribution lies in this plane is then used on an event-by-event basis to correct the DCR parameter for its drift time dependence, resulting in the `dcr_corr` parameter in the skim files. This procedure in general narrows the DCR distribution as shown in Figure 3.9 where the `dcr` parameter used in the previous analysis is compared to `dcr_corr` described herein. The DCR efficiency, including the drift time correction, for each data set is summarized in Table 3.2. The main uncertainties in calculating DCR are the statistical uncertainty, the time-stability systematic and the pulse-shape bias systematic. The process underwent to compute these uncertainties is described in great detail in Section 4.7 and so will not be reproduced here.

The DCR cut used in the analysis is:

```
dcr_corr < 2.326
```

3.9 Multisite Events & AvsE

P-type point contact detector geometries have slower charge collection and more localized weighting potential than traditional coaxial or well-type HPGe detectors do. This allows for the identification and rejection of events that deposit energy in multiple sites within a

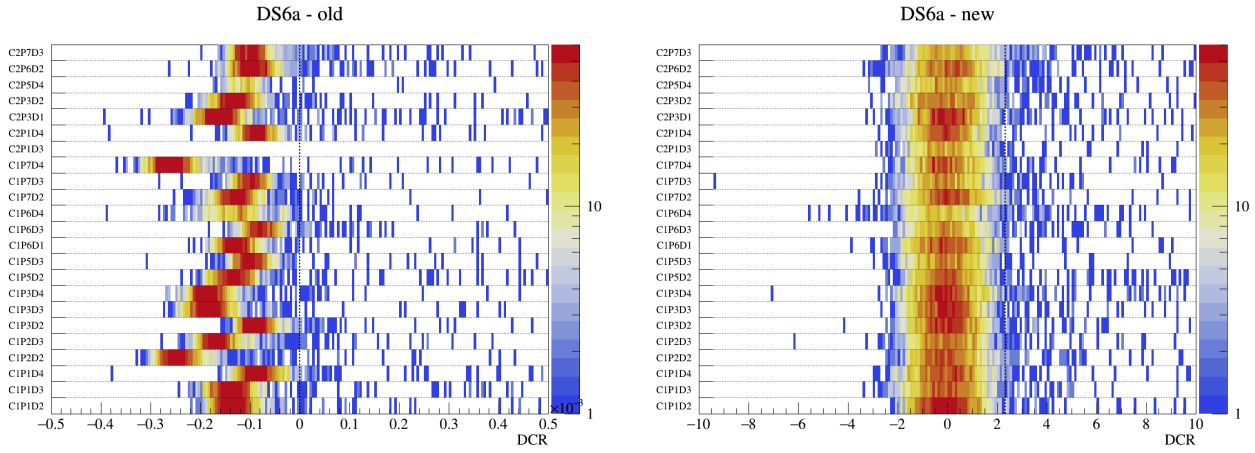

Figure 3.8: The DCR distribution for each detector in DS6a before (left) and after (right) PZ correction and the mean and width correction. The data shown is the background data between 100 and 2650 keV. The cut value for 99% acceptance is at 0 in the previous version of the cut and 2.326 when using the `dcr_corr`. Taken from [60].

detector, such as Compton-scattered gammas. The parameter AvsE represents a comparison between the maximum amplitude (A) of a waveform current pulse and the total energy (E) [1]. Multi-site events usually have lower maximum current amplitude A for a given energy than single-site events, as can be seen in Fig. 3.10, as the overall energy is split into several sites, the electrons of holes of which reach the contacts of the detector at different times. As such we see two peaks in the current amplitude, each of which is lower in A than the equivalent single-site event.

Before getting into technicalities, here is a brief definition of Single Escape Peak (SEP) and Double Escape Peak (DEP). In a detector, if the energy of the incident γ particle is above 1022 keV, pair production can occur, resulting in two 511 keV annihilation gamma-rays. If one of these γ-rays escapes while the other is completely absorbed in the detector, 511 keV will be carried away from the detector by the escaped γ. This results in a separate peak in the spectrum at (in our case) 2614-511 keV, or 2103 keV, called the single escape peak (SEP). If both γ-rays escape the detector, 1022 keV of energy is lost, giving rise to the

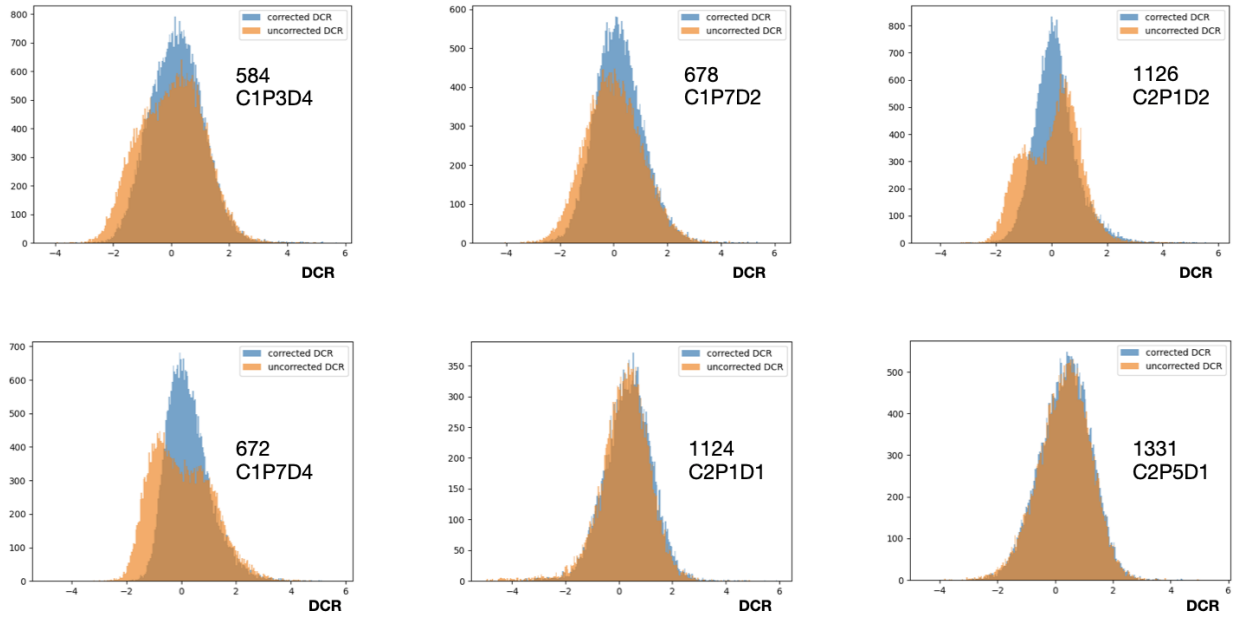

Figure 3.9: The DCR distributions for a sampling of detectors, before and after Charge Trapping Correction, in the last long calibration of DS6a. The filled orange and blue histograms show DCR and the drift-time corrected DCR respectively.

double escape peak (DEP) at 2614-1022 keV, or 1592 keV.

Low AvsE Cut

The AvsE cut is tuned for each detector to an efficiency of 90% based on the double-escape peak (DEP) from the 2614.5 keV ^{208}Tl line. Typical, weekly calibrations lasting $60-90$ min have insufficient statistics for tuning of the AvsE parameters, so 18 hr long dedicated calibrations are used to calibrate the AvsE parameter for each detector such that a cut at the value of -1 can be applied. The more frequent, short calibration runs are then also used to estimate the cut efficiency across the data set in addition to any uncertainty caused by time-variation of detector performance. This is described in detail in [49, 1].

Compared to the analysis in [22], the AvsE parameter has been updated to better correct

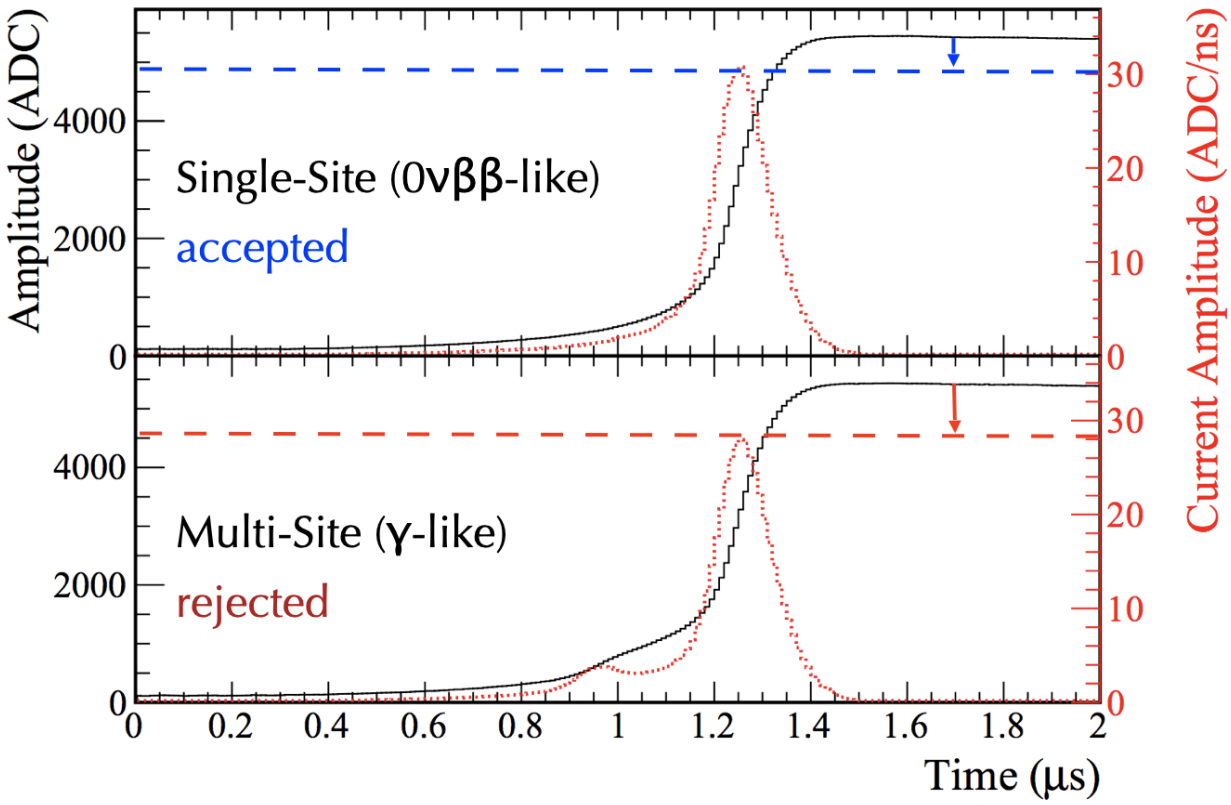

Figure 3.10: For multi-site events, the maximum current amplitude is usually lower than for a single-site event of the same energy. From [1].

for the energy dependence of both the mode and width of the AvsE distribution. This relies on ^{56}Co calibration data which contains multiple double-escape peaks near the $0\nu\beta\beta$ ROI, unlike the ^{228}Th spectrum. Due to the relatively short half-life of ^{56}Co, the logistics of obtaining source and deploying the source limited this data to a single dedicated deployment during DS6.

Figure 3.11 (from [1]) shows the AvsE acceptance for numerous single- and double-escape peaks in the ^{56}Co data in additional to multiple Compton continuum regions in the spectrum. The darker points in each case show the acceptance after correction of the mode and width to remove the energy dependence as described in [1]. The ^{56}Co data shows that pre-

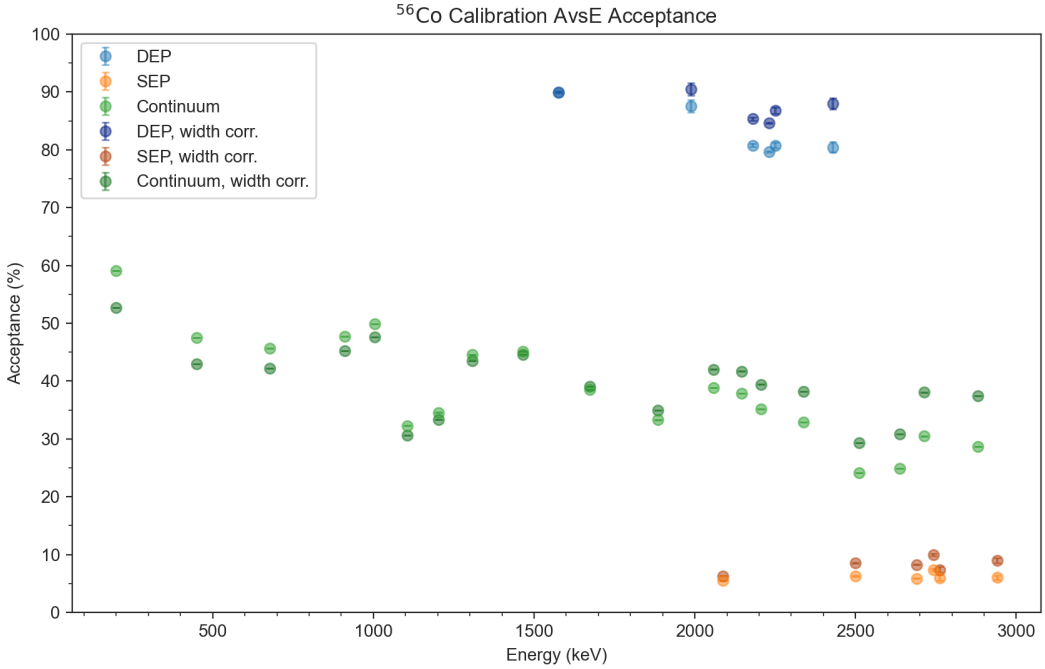

Figure 3.11: From [1], the AvsE acceptance over various energy regions with and without a width-energy dependence correction. The error bars are smaller than the marker size.

vious estimates of the AvsE acceptance at the $0\nu\beta\beta$ ROI were systematically high by 3-5%, depending on how one extrapolates over energy. Although this effect was not measured in previous analyses, systematic uncertainties in the energy dependence of the mode and width of the AvsE distribution were the dominant contributions to the $\approx 3.3\%$ uncertainty in the AvsE acceptance as shown for the earlier data sets in Table 3.2. Using the ^{56}Co data, the systematic uncertainties associated with acceptance of AvsE at the $0\nu\beta\beta$ ROI are reduced, although the overall uncertainty is dominated by stability between calibration runs so the final uncertainty is similar to previous analyses [1].

In previous analyses [67] and [22], the true acceptance of single-site events at the $0\nu\beta\beta$ ROI was closer to 85-87% although the systematic uncertainty on the 90% acceptance quoted

bracketed this possibility. In better accounting for the energy dependence of AvsE based on the ^{56}Co data, the true acceptance of single-site events at the $0\nu\beta\beta$ ROI is expected to be much closer to 90%, and a cut at 90% single-site acceptance is still the optimal cut value as shown in [1]. However, this change in single-site event acceptance will also result in a more background events passing the AvsE cut. The level to which the background will increase is dependent on the shape of the AvsE distribution for background events near the $0\nu\beta\beta$ ROI. Because we are increasing the acceptance of the cut, we also expect an increase in backgrounds; the SEP acceptance increased by $1-2\%$, and the ROI Compton acceptance by an average of 3%, from $\sim 40\%$ to $\sim 43\%$. This produces an expected increase of Compton-continuum backgrounds of $\sim 7.5\%$, compared to a 5% increase in signal acceptance, so we still expect an overall increase in sensitivity.

The effect of the mean stabilization and width correction of AvsE also tends to make the distributions more uniform across detectors. Figure 3.12 shows the AvsE distribution for each detector in DS6a before and after the improvements described in [1] with the cut value indicated by the vertical line. The improved cut also allows for a tighter criterion for removing events due to instabilities in AvsE; this results in a smaller systematic uncertainty on the efficiency, but removes slightly more exposure relative to the 2018 and 2019 results. Detectors which are not present in the left panel of Figure 3.12 but are present in the right panel were previously not used in Majorana Demonstrator analyses for the entirety of DS6a.

As described in [1], the AvsE parameter was found to correlate with drift time as was first noted and corrected for the ORNL analysis [76]. For each detector in each data set, the main body of the two-dimensional AvsE vs drift time distribution from calibration data is first fit to a line. This is then used on an event-by-event basis to rotate the AvsE distribution in the AvsE-drift time plane which results in a narrower AvsE distribution. An example of this procedure is shown in [1]. The drift-time-corrected AvsE is referred to as the parameter `avse_corr` in the skim files.

The AvsE efficiency and uncertainties for each data set is summarized in Table 3.2.

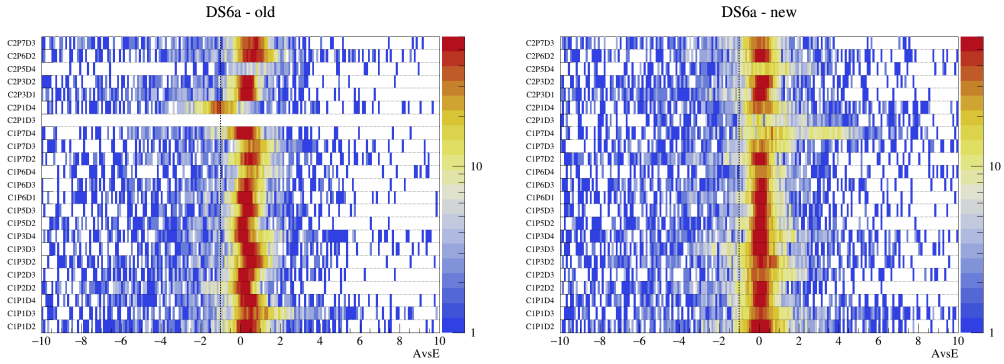

Figure 3.12: The AvsE distribution for each detector in DS6a before (left) and after (right) the improved mode and width energy dependence. The data shown is the background data between 100 and 2650 keV. The cut value in both cases is at -1 which is indicated by the vertical dashed line. Note that the parameter plotted here is AvsE prior to the charge trapping correction. Taken from [60].

This includes the corrections for the energy dependence of the mode and width of the AvsE distribution in addition to the drift time correction. The overall AvsE efficiency is simply the average efficiency of all active detectors, natural and enriched with no livetime or exposure correction made. The same long calibration data used to tune the AvsE parameters are used to assess the efficiencies. The uncertainty is the sum in quadrature of the following components [1]:

- Statistical uncertainty of the DEP survival fraction, calculated channel by channel and then averaged for all long calibrations in each dataset summed together. The uncertainty for each channel is calculated from binomial statistics and the combined error is the standard error on the mean.

- Uncertainty from the AvsE energy dependence, accounting for the difference in the AvsE distribution mode at the DEP vs. the ROI. The cut is varied in both directions by the difference of these two values. Finally, the difference in the efficiency is considered

as the systematic. It is calculated channel by channel and then averaged.

- Uncertainty due to the residual differences between calibration and physics data. The cut is changed by the largest difference between μ_{AvsE} for $2\nu\beta\beta$ (950-1400 keV) and Compton events, for all operating detectors in each dataset. The difference in the efficiency is considered as the systematic.

- Systematic uncertainty due to the difference between the survival fraction of $0\nu\beta\beta$ events and DEP single-site events. It is estimated from pulse shape simulation and the acceptance trend in ^{56}Co DEPs after the AvsE width correction is applied. The difference in the acceptance for the simulated double escape peak and simulated $0\nu\beta\beta$ peak is taken as the uncertainty.

- The width energy dependent correction, which is determined by fitting the width of the various DEPs in a ^{56}Co spectrum channel by channel-by-channel. There are two terms that effect the width of the AvsE distribution: the empirical width function itself and the $cos(\theta)$ scaling term in the drift-time correction. The width function is varied by $\pm 1 - \sigma$ confidence band and then the $cos(\theta)$ is varied within its standard error from the drift-time correction.

- Stability: calibration skim files are chained together and analyzed. The energy window used is 1570 keV - 1610 keV, specifically a DEP window from 1587.5 keV to 1597.5 keV and the sideband. The total number of events, and the number of events passing the cut, is calculated in each region. The uncertainty, σ is computed by standard error propagation.

High AvsE Cut

In addition to multisite events, `avse_corr` is sensitive to near-point contact hits, which have significantly faster rise-times than waveforms in the bulk of the detector. As a result, by

cutting on high value of `avse_corr`, we can remove backgrounds that would escape other cuts and improve our sensitivity. These events may result from alphas incident on the detector surface near the point contact, which suffer less energy degradation. In this way, a high `avse_corr` cut is complementary to the DCR cut (Sec. 3.8). In addition, multisite events with most energy deposited near the point contact may have high, rather than low, values of `avse_corr`. A cut value of 9 was selected by optimizing for sensitivity, balancing signal sacrifice against the improved background rejection, as described in [1]. This cut has a signal acceptance of 98%, measured using Compton continuum events with the same energy range near the ROI used for DCR. The high AvsE efficiency for each data set is also summarized in Table 3.2. It is determined for each dataset by calculating the double escape peak acceptance for all detectors and long calibrations in each dataset. To ensure that the AvsE efficiency and the High AvsE efficiency multiply together to give an overall AvsE efficiency, the $AvsE > -1$ cut is applied before computing the High AvsE efficiency. The systematics for the high AvsE cut efficiency are computed as follows:

- The statistical uncertainty of the DEP survival fraction, calculated for all detectors combined.

- The difference in the DEP survival fraction and the survival fraction of events in the energy window of 950 - 1400 keV in open background after applying a DCR cut. This is used to account for the difference in high AvsE acceptance from calibration data and open background data.

- The stability of the DEP survival fraction over short calibrations, calculated using the same method as the standard low AvsE cut.

.

The AvsE cuts used in the analysis are:

`avse_corr>-1 && avse_corr<9`.

3.10 ORNL Analysis

The ORNL analysis is an important cross-check for the MAJORANA DEMONSTRATOR analysis, done by a subgroup of the MAJORANA Collaboration. Due to time constraints, the MAJORANA DEMONSTRATOR analysis has adapted several parts of the ORNL analysis. Specifically, for the new, larger inverted-coax point contact (ICPC) detectors, the ORNL analysis is used exclusively. Similarly, the ORNL implementation of the Late Charge, or LQ, cut is used (see Sec. 3.10.1). Details in this section are taken from the 2022 MAJORANA DEMONSTRATOR technical document (Unidoc) [60] as well as the ORNL Analysis unidoc [76].

In DS8, Module 2 contains 4 ICPC detectors that use a separate analysis for signal processing. A separate scheme is needed because of the larger size of the detectors, leading to longer drift times, which may require additional corrections. The ICPCs use a separate calculation performed by the ORNL group, for energy, A/E (equivalent in scope to the typical MAJORANA DEMONSTRATOR AvsE), DCR, LQ and waveform data cleaning [76]. These parameters are then used directly in the analysis; the cuts all are applied around a value of zero.

The full set of ICPC background cuts is:

```
!isMuVeto && !(isLNFill2&&C==2) && mHL==1 && isGood && isICPC && wfDCBits==0
&& ORNL_AoverE>0 && ORNL_DCR<0 && ORNL_AoverE<10 && ORNL_LQ<0 && Final_Energy>100
```

For the ICPCs, `Final_Energy` is imported from the ORNL analysis without performing a dedicated calibration. It represents the calibrated energy. The analysis of ICPCs is performed separately from the PPCs, due to both physical differences between the detectors and different cuts used.

`ORNL_AoverE` is a waveform feature used to distinguish between single-site and multi-site events. In the ORNL analysis, this variable is optimized to select single-site events using the Double Escape Peak (DEP) of the 2614.5 keV γ, as that peak is purely single-site. It is the

ratio of the current amplitude (measured using a triangular filter) to waveform amplitude, transformed so that 90% of 1592 keV DEP events fall above zero and 98% of the Compton continuum around $Q_{\beta\beta}$ falls below 10. The distribution is not transformed to be aligned by width or mode as is done with `avse_corr`. A/E values are correlated with drift time and energy, so a drift time correction is applied, similar to `avse_corr`, and an additional energy correction is applied. The cut is then implemented by determining A/E value for which 90% of DEP events are accepted. For the A/E cut applied to the background data between calibrations, a time-wise linear interpolation of the cut value between those calibrations is used for each different run [76]. The same procedure, adjusted for a cut value of 0, for estimating the acceptance efficiency for single-site bulk events is used, with the following exceptions as described in Sec. 3.9. Both low and high cuts are applied, at values of 0 and 10, respectively, tuned for 90% detection efficiency of the DEP at 0 and 2% rejection efficiency of the Compton continuum around $Q_{\beta\beta}$ at 10. First, since no ^{56}Co data was taken with the ICPCs, the energy dependence systematic for the distribution center is measured using the mode of the distribution in calibration runs. Second, we correct down the signal-efficiency by 5.8%, to match the loss of efficiency in PPCs prior to applying the energy correction. The low A/E efficiency is 0.852 ± 0.059 and the high A/E efficiency is 0.995 ± 0.009 [60].

`ORNL_DCR` is calculated by looking at all calibration events above 1500 keV, in order to optimize the cut. The method to establish a background rejection cut arises from the study of the DCR distribution. The cut is chosen to keep an acceptance higher than 99.5%. Since DCR values are correlated with the drift time, the DCR value is first linearly corrected for drift-time effects. The DCR cut value is determined for each calibration independently. It is taken from the FWHM of the DCR peak in the drift-time-corrected DCR distribution. This parameter is aligned so that three-times the FWHM falls at zero, i.e., for a Gaussian distribution, $>$ 99.9% of bulk events would fall below the cut value of zero. The DCR values from the ORNL analysis have been integrated into the MAJORANA DEMONSTRATOR analysis for the ICPC detectors in DS8. The efficiency and systematic uncertainty are measured in the same manner as for the regular `dcr_corr`, just adjusted for the different

cut value. Additionally, the efficiency of the cut after applying the A/E cut is computed to be 0.9957 ± 0.001. This efficiency has been calculated by using event in the DEP region (1590 keV-1595 keV), where 93 events were rejected out of 13910, and a side band of the same side where 53 events out of 1708 were rejected. The uncertainty was calculated by using 2 contributions, a statistical contribution, with $\sigma_\epsilon = \sqrt{\epsilon \times (1-\epsilon) \times N}$, where ϵ is the acceptance and N is the number of events in the DEP region, and a systematic contribution due to the sideband-subtraction, $\sigma_\epsilon(syst.) = \epsilon - \frac{n-\Delta n}{N-\Delta N}$. In the last equation, n corresponds to the accepted events in the DEP region subtracting the side band and the uncertainties of this variable and N are obtained by error propagation of the Poisson error in the DEP region and side band [76].

`ORNL_LQ` is the LQ value as described in Sec. 3.10.1. It applies equally to both ICPC and non-ICPC analyses.

3.10.1 Late Charge (LQ)

The LQ variable identifies events in which the a fraction of the collected charge arrives later in the event than in a typical single-site event. This allows for rejection of multi-site events that pass the A/E cut because they have a large energy deposition close to the p+ contact, thereby generating a high value of A, together with a smaller energy deposition further from the contact. LQ values are correlated with the drift time, so LQ values are corrected for drift-time effects prior to evaluating a cut value. Also, events with a partial charge deposition in the partially-dead layer between the n-type surfaces and the detector bulk experience energy degradation, and have waveforms with an added slow-pulse component. The energy degradation is related to slowness, due to the diffusion of drift charges out of the low-field region. Because of the low-field, the charges reach the contacts of the detector more slowly, which can be seen at the end of the rising edge of the waveform. The slow-pulse component of these waveforms increases LQ, and we cut events that fall $> 5\sigma$ above the center of the parameter distribution [76].

The lq cut is:

`lq < 5`

The late charge (LQ) parameter is calculated by measuring the area above the electronics' pole-zero corrected waveform, starting at the 80% timepoint, which is computed using the trapezoidal filter discussed in Section 3.6, finding the maximum of the waveform and stepping back until the 80% rise time is reached. A visual representation of how LQ is computed in real waveforms is given in Fig. 3.13.

LQ is sensitive to:

- Multi-site events, which tend to have a larger value of LQ; for multisite events with most energy deposited near the point contact, the fast rise time will usually result in a normal or high value of AvsE but in very high values of `LQ`.

- Events with a slow component, due to partial charge deposition in the transition dead layer, which will have an elevated LQ value. Such events can come from high energy betas or multi-site events that deposit charge in the transition layer.

- α events in the passivated layer also show large LQ, due to the charges slowly drifting out of the passivated layer, but are not in the Li+ transition layer. This type of event is discriminated against by LQ as well as DCR, despite LQ not being specifically tuned to catch it.

The LQ parameter is tuned to have the bulk of its distribution centered at 0, with standard deviation of 1. This tuning process also applies a linear correction for 90% rise time, i.e. the time it takes for the waveform to rise to 90% of its maximum. The LQ cut was originally developed in the ORNL code. The signal processing parameter used by the rest of the MAJORANA DEMONSTRATOR analyses differs in that it uses a $[1.5, 0, 1.5]$ μs trap-filter while ORNL uses a $[1, 1, 1]$ μs trap-filter. The MAJORANA DEMONSTRATOR uses the 90% rise time while ORNL uses the ORNL drift time parameter. A cut value of 5σ is used.

To verify that the event populations described above are cut, while bulk single-site events are avoided, we turn to calibration data as in Fig. 3.14. In the BG window of 1950-2350 keV,

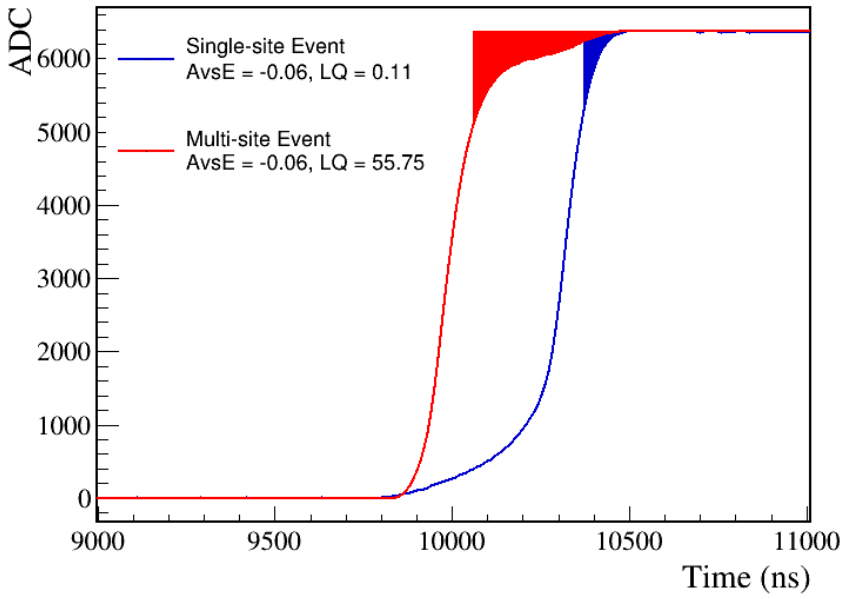

Figure 3.13: Example of an LQ-cut event compared to a single-site bulk event. The area used to calculate the LQ parameter is shaded. In blue we see a normal bulk single-site event, whereas in red we see a multi-site event with high LQ ($LQ \sim 55$). Taken from [60].

which is mostly Compton Continuum (CC) events, and includes the single escape peak (SEP), we cut 2.5% of events. The detection efficiency of the LQ parameter is measured using the 1592 keV double escape peak (DEP) and adjusted for energy dependance based on DEPs in ^{56}Co data.

The continuum between the Compton shoulder and 2614 keV peak is heavily multi-site, resulting in a relatively high fraction (13.9%) of events cut. The Compton shoulder, which is predominantly single-site, does not clearly show up in the LQ spectrum. The tails of the DEP and SEP peaks are much more prominent in the cut-by-LQ spectrum. Those tails are presumably due to partial deposition in the transition layer of the exiting 511 keV γs. The systematic uncertainty in the LQ is calculated using the same method that is used as cross-check in the DCR but applying the DCR cut in the samples too. Using the same regions

used in the DCR calculations, the efficiency of this cut is estimated at 0.9977 ± 0.0011 [76].

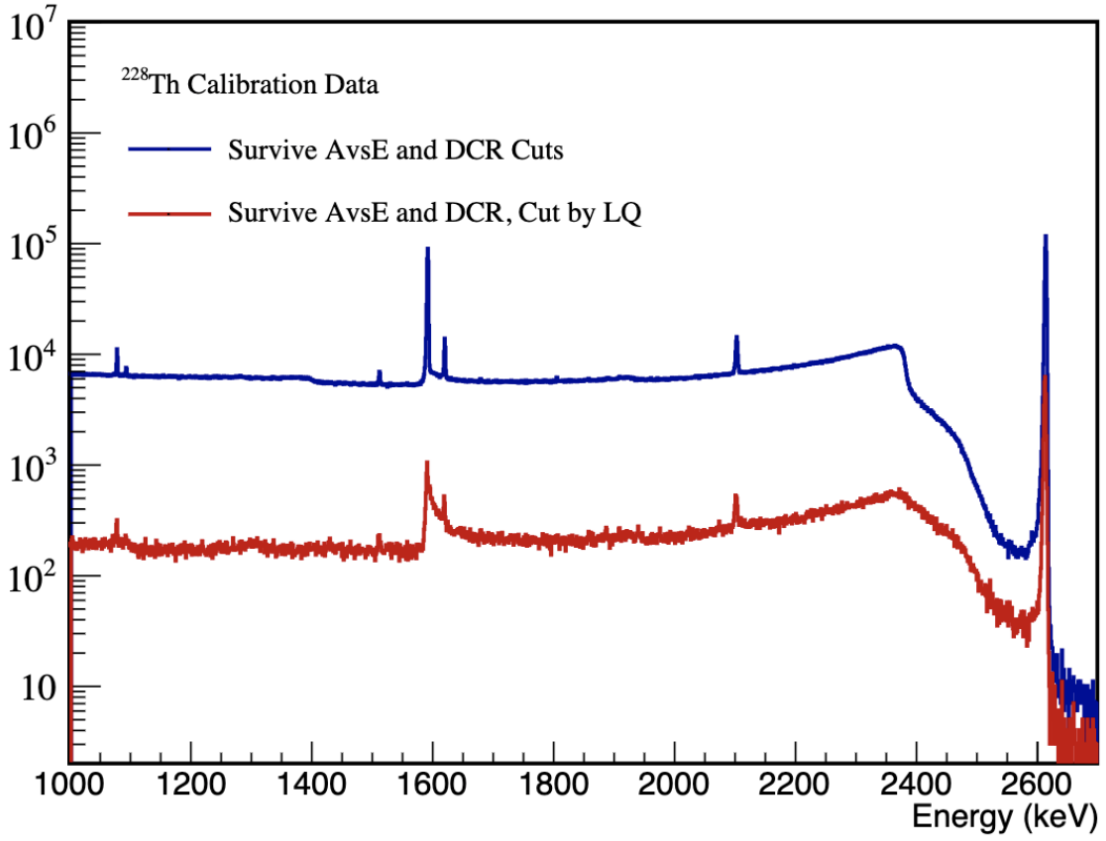

Figure 3.14: Combined spectrum in DS8 for ICPCs, after all other cuts are applied, before the LQ cut is applied, in blue. Events removed by the LQ parameter are presented in the red spectrum. Taken from [60]. The SEP is at 2103 keV and the DEP is at 1592 keV.

3.11 Detection Efficiencies

To calculate the half-life limit of the $0\nu\beta\beta$ process, one important thing to know is the detection efficiency of the MAJORANA DEMONSTRATOR. The $0\nu\beta\beta$ detection efficiencies are summarized in Table 3.2. The order in which the values in this table is important. We assume the cuts are applied in order from left to right, i.e. that the second efficiency is cal-

culated after the first cut is applied, the third efficiency computed after the first and second cuts are applied, and so on. Specifically, the DCR efficiency, as described in Chapter 4 is computed assuming the AvsE cut is already applied. It is also believed that there are possible correlations between the efficiencies from dataset to dataset, i.e. that some of the uncertainties could be split between a fully-correlated part, across datasets and a uncorrelated part, usually smaller in magnitude.

Set	ϵ_{clean}	ϵ_{cont}	ϵ_{lowAE}	ϵ_{highAE}	ϵ_{DCR}	ϵ_{LQ}	ϵ_{tot}	ϵ_{tot} with Resol.
DS0	0.999 ± 0.001	0.907 ± 0.011	$0.898^{+0.029}_{-0.025}$	0.988 ± 0.004	0.985 ± 0.003	0.993 ± 0.007	$0.786^{+0.028}_{-0.025}$	$0.682^{+0.030}_{-0.027}$
DS1	0.999 ± 0.001	0.908 ± 0.011	$0.899^{+0.028}_{-0.026}$	0.982 ± 0.004	0.980 ± 0.006	$0.992^{+0.008}_{-0.006}$	$0.779^{+0.028}_{-0.025}$	$0.672^{+0.030}_{-0.028}$
DS2	0.999 ± 0.001	0.908 ± 0.011	$0.899^{+0.035}_{-0.033}$	0.983 ± 0.015	0.981 ± 0.013	$0.993^{+0.007}_{-0.009}$	$0.780^{+0.036}_{-0.035}$	$0.668^{+0.038}_{-0.037}$
DS3	0.999 ± 0.001	0.908 ± 0.011	$0.899^{+0.038}_{-0.040}$	0.978 ± 0.015	0.983 ± 0.006	$0.994^{+0.006}_{-0.008}$	$0.779^{+0.037}_{-0.039}$	$0.669^{+0.039}_{-0.040}$
DS4	0.999 ± 0.001	0.908 ± 0.010	$0.901^{+0.039}_{-0.042}$	$0.977^{+0.023}_{-0.028}$	0.985 ± 0.007	$0.997^{+0.003}_{-0.009}$	$0.785^{+0.040}_{-0.045}$	$0.691^{+0.043}_{-0.048}$
DS5a	0.999 ± 0.001	0.908 ± 0.013	$0.900^{+0.038}_{-0.043}$	0.978 ± 0.013	0.983 ± 0.008	$0.993^{+0.007}_{-0.012}$	$0.779^{+0.037}_{-0.042}$	$0.677^{+0.038}_{-0.042}$
DS5b	0.999 ± 0.001	0.908 ± 0.013	$0.900^{+0.036}_{-0.041}$	0.978 ± 0.014	0.982 ± 0.007	$0.993^{+0.007}_{-0.009}$	$0.778^{+0.036}_{-0.040}$	$0.676^{+0.037}_{-0.041}$
DS5c	0.999 ± 0.001	0.908 ± 0.013	$0.900^{+0.037}_{-0.043}$	0.978 ± 0.016	0.982 ± 0.006	$0.994^{+0.006}_{-0.010}$	$0.779^{+0.037}_{-0.042}$	$0.672^{+0.039}_{-0.043}$
DS6a	0.999 ± 0.001	0.908 ± 0.013	0.899 ± 0.030	0.980 ± 0.007	0.985 ± 0.007	0.994 ± 0.006	0.783 ± 0.030	0.675 ± 0.032
DS6b	0.999 ± 0.001	0.908 ± 0.013	$0.899^{+0.032}_{-0.030}$	0.977 ± 0.009	0.986 ± 0.008	0.994 ± 0.006	0.781 ± 0.031	0.674 ± 0.033
DS6c	0.999 ± 0.001	0.908 ± 0.013	$0.898^{+0.030}_{-0.028}$	0.978 ± 0.011	0.984 ± 0.006	$0.993^{+0.007}_{-0.005}$	$0.779^{+0.031}_{-0.029}$	$0.672^{+0.033}_{-0.031}$
DS7	0.999 ± 0.001	0.908 ± 0.011	$0.900^{+0.040}_{-0.047}$	0.983 ± 0.009	0.985 ± 0.009	0.992 ± 0.008	$0.784^{+0.038}_{-0.043}$	$0.673^{+0.040}_{-0.045}$
DS8P	0.999 ± 0.001	0.908 ± 0.013	$0.900^{+0.035}_{-0.033}$	0.980 ± 0.017	0.989 ± 0.005	0.993 ± 0.006	$0.786^{+0.036}_{-0.034}$	$0.679^{+0.037}_{-0.036}$
DS8I	0.999 ± 0.001	0.919 ± 0.008	$0.852^{+0.042}_{-0.059}$	0.978 ± 0.014	0.979 ± 0.011	$0.995^{+0.005}_{-0.009}$	$0.745^{+0.040}_{-0.055}$	$0.647^{+0.042}_{-0.056}$

Table 3.2: The efficiencies for the enriched detectors for each data set. The final uncertainty combines all the individual uncertainties in quadrature. The efficiency with resolution is given for reference. The efficiency due to the granularity cut is assumed to be unity across all data sets. The data cleaning efficiency is checked for each data set and has been found to be consistent across all data sets. More details on the efficiencies not computed in this document can be found in [82] for ϵ_{clean}, [1] for ϵ_{AE}, and in Chapter 4 and [79] for ϵ_{DCR}.

.

3.12 Exposure

Referring back to Eq. 3.1, we can see that not only the total efficiency of our cuts, ϵ, is important, but also the product NT, which is called the exposure.

Table 3.3 shows the exposure for the open and blind portions of each data set, as well as the total exposure in datasets 0 through 8. The total exposure for DS0-8 including both open and blind data is 64.5 ± 0.9 kg-yr. The details of the exposure calculation are described below.

3.12.1 Active Mass/Fraction

The dead layer analysis is used to account for inactive mass in each Ge detector [37]. For the ICPC detectors, a characterization measurement at SURF was not performed, so the dead layer analysis from Ortec was used, assuming a 15% uncertainty, which is consistent with the mean uncertainty measured in the PPC detectors. This is done for each detector individually. For the enriched PPC detectors, approximately $91.8 \pm 1.4\%$ of the mass is active; for the enriched ICPCs $90.9 \pm 1.3\%$; for the natural BEGes, $90 \pm 2\%$. For the enriched detectors, this uncertainty dominates the total uncertainty on the exposure. Only the enriched detectors are considered for the $0\nu\beta\beta$ decay limit. As a first estimate of the systematic error in the dead-layer thickness, the RMS of difference of the different dead layer measurements are taken. An estimated uncertainty for the thickness of the passivated surface and the point-contact dead-layer of 0.5 mm is used. To compute the error, uncertainty is propagated in the standard manner. The dead layer study allows the computation of the active fraction, geometrically. It can also be used to compute the active mass of the detector by multiplying the total mass by the active fraction. The error in the measurement of the total mass is assumed to be minimal [37].

Data Set	Access	M1 Exposure (kg-y)	M2 Exposure (kg-y)	$NT\epsilon_{tot}\epsilon_{ROI}$ (10^{24} A-y)
DS0	open	1.12 ± 0.02	0.00 ± 0.00	$5.33^{+0.04}_{-0.04}$
DS1	open	1.81 ± 0.03	0.00 ± 0.00	$8.47^{+0.04}_{-0.04}$
DS1	blind	0.45 ± 0.01	0.00 ± 0.00	$2.08^{+0.04}_{-0.04}$
DS2	open	0.27 ± 0.00	0.00 ± 0.00	$1.25^{+0.05}_{-0.05}$
DS2	blind	0.85 ± 0.01	0.00 ± 0.00	$3.96^{+0.05}_{-0.05}$
DS3	open	0.97 ± 0.01	0.00 ± 0.00	$4.51^{+0.05}_{-0.05}$
DS4	open	0.00 ± 0.00	0.26 ± 0.00	$1.24^{+0.06}_{-0.05}$
DS5a	open	2.15 ± 0.03	0.56 ± 0.01	$12.79^{+0.05}_{-0.06}$
DS5b	open	1.32 ± 0.02	0.45 ± 0.01	$8.32^{+0.05}_{-0.06}$
DS5c	open	0.40 ± 0.01	0.14 ± 0.00	$2.51^{+0.05}_{-0.06}$
DS5c	blind	1.21 ± 0.02	0.42 ± 0.01	$7.62^{+0.05}_{-0.06}$
DS6a	open	3.01 ± 0.04	1.32 ± 0.02	$20.32^{+0.04}_{-0.04}$
DS6a	blind	7.35 ± 0.11	3.07 ± 0.04	$48.93^{+0.04}_{-0.04}$
DS6b	open	1.81 ± 0.03	0.78 ± 0.01	$12.14^{+0.04}_{-0.04}$
DS6b	blind	4.98 ± 0.07	2.21 ± 0.03	$33.70^{+0.04}_{-0.04}$
DS6c	open	3.35 ± 0.05	1.40 ± 0.02	$22.18^{+0.04}_{-0.04}$
DS6c	blind	5.96 ± 0.09	2.55 ± 0.04	$39.75^{+0.04}_{-0.04}$
DS7	open	1.53 ± 0.02	0.00 ± 0.00	$7.15^{+0.05}_{-0.06}$
DS7	blind	2.95 ± 0.04	0.00 ± 0.00	$13.80^{+0.05}_{-0.06}$
DS8P	open	1.47 ± 0.02	0.76 ± 0.01	$10.57^{+0.05}_{-0.05}$
DS8P	blind	3.19 ± 0.05	1.56 ± 0.02	$22.43^{+0.05}_{-0.05}$
DS8I	open	0.00 ± 0.00	0.89 ± 0.01	$4.05^{+0.06}_{-0.08}$
DS8I	blind	0.00 ± 0.00	1.93 ± 0.03	$8.73^{+0.06}_{-0.08}$
DS0-8	all	46.14 ± 0.68	18.32 ± 0.26	$301.81^{+1.13}_{-1.18}$

Table 3.3: The active, enriched mass exposure for each data set [50]. The uncertainty for the totals accounts for correlated uncertainties such as the active mass uncertainty. The total exposure for datasets 0-8 is presented in the last line.

3.12.2 Number of Atoms

The number of target atoms is related to other measured quantities by

$$N = \frac{N_A M_D f_{76}}{M_{mol}}, \tag{3.5}$$

where N_A is Avogadro's Number, M_D is the total detector mass in grams, f_{76} is the isotopic fraction of ^{76}Ge. f_{76} is found in [52] to be $87.4 \pm 0.5\%$ for the PPCs and $88.0 \pm 1.0\%$ for the ICPCs. The molar mass M_{mol} of the enriched material and is 75.668 ± 0.010 g/mol for the PPCs and 75.681 ± 0.020 g/mol for the ICPCs [52]. To estimate the true unknown isotopic fraction of different species in a pure sample based on direct measurements of those abundances, one minimizes the χ^2 statistic under the assumption that the isotopic fractions sum to 1. The uncertainty in each of the isotopic fraction f_i can be computed by standard error propagation. The same method can be used to evaluate uncertainties in functions of f_i, such as the molar weight M [52].

3.12.3 Live Time

In MJD the live time differs from the run time due to the muon veto, the LN fill vetos and electronics dead time [50]. The exposure for each dataset is listed in Table 5.2. Furthermore, we eliminate certain runs and channels at certain times according to our channel and run selection procedures. Dead time is induced through removing periods where an event is within some coincidence window of other events external to the detector (pulsers, LN fills, muon veto, etc) as summarized in the following sub-sections. These effects are included in the exposures listed in Table 5.2.

Muon Veto

Despite the MAJORANA DEMONSTRATOR being located deep underground, cosmic muons can still penetrate the Earth and reach our detector. Therefore we have implemented a muon cut, briefly described below.

The

```
!isMuVeto
```

cut removes any muons, as follows. The muon veto system triggers when at least two veto panels surpass an energy threshold, specifically 500 QDC units using CAEN 792 digitizers, chosen based on how much charge muons deposit in the veto panels underground. It has a detection efficiency of 96.71% [121]. A separate muon analysis then identifies muon candidate events from this data and identifies time periods to cut around muon events. For 20 ms before and 1 s after this event, all HPGe events are tagged using `muVeto` and removed from the analysis. The duration of this period is subtracted from the livetime for each detector. We generally expect that the majority of muon-induced physics events will hit the Ge array nearly instantaneously after hitting the veto system, such as direct muons and gamma events in a muon shower. On a slightly longer timescale are the majority of excited isotopes in Pb, Cu, and Ge that can be produced by the muon event, the vast majority of which decay in less than 1 second [121]. While there are longer-lived decays, we decision was made to not keep the array dark for these extremely few events. The livetime uncertainty associated with this cut is negligible [120].

LN Fill Cut

The

```
!(isLNFill1&&C==1) && !(isLNFill2&&C==2)
```

cut removes any periods in which the MAJORANA DEMONSTRATOR liquid nitrogen (LN) dewars were re-filled. These periods are associated with higher microphonic noise.

LN fills take place in each module approximately every 35 h or so. Data is rejected during each fill for the module in which a fill is underway. Through DS6a, a 20 minute period surrounding the fills was identified using the LN level sensors. After DS6a, we identify a 12-15 min period around LN fills using the ami valve state. This defines the isLNFill flag. No evidence of significant microphonics in the module which is not filling has been

observed during a fill of the other module. The livetime uncertainty associated with this cut is negligible [119].

3.12.4 Run Selection and Channel Selection

The

`isGood`

cut is a data quality cut, which includes only data passing the run selection and channel selection procedures.

Run selection summaries are provided in [86]. As a first pass on run selection, the run bits are used to designate a run rank. A list of run bits and rankings is provided in Fig. 3.15. If any of the bits from the "Bad" column are set, for example, that run is automatically ranked as bad regardless of what other bits are set. Runs with no run quality bits set are ranked as gold (equivalent to 'good' runs). Based on ongoing activities, runs can be manually promoted or demoted from their initial automatic ranking. Additional checks are performed for each data set to look for breakdown events or any potential issues with the initialization of GRETINA cards.

In some cases, issues result in particular channels that must be excluded for a particular run range. A list of reasons for channel selection is provided below

- Calibration: Channels that could not be calibrated, or sufficiently large gain drifts between calibrations.
- Builder: Problems in building events for a specific channel.
- Data Cleaning: Runs where the data cleaning cut rate in a channel is anomalously high.
- Run Selection: Multiple checks, such as pulser timing, as well as whether detectors are biased down.

Bit	Silver	Bronze	Bad	Cal	Definition
0			1		Run length: shorter than 3 minutes (1 min for cal).
1			1		Run end: due to ORCA crash -StopTime must be non-zero.
2		1			Radon purge: sufficient - average below 5 slpm during run.
3	1				Pulser monitor channels are not running (after run 4549).
4			1		Run initialization: issues at run start.
5		1			Veto not running: "It does not match expected rate".
6			1		Builder: events not able to be built.
7	1				Rate: too high and may affect live-time.
8				1	Source present: thus a calibration run - check Th and Co bits (5,6).
9			1		Trigger Card: Master Trigger card reported lost lock.
10		1			Maintenance/disruptive work: (check run bits 12-machine shop and 13-disruptive work).
11	1				Shield: not present (Partial shield bit (7) set).
12			1		Slow control DB is not "available for run".
13			1		Data production logs - bad errors.
14		1			One or more Orca bits set: 0 (DAQExpertMode),16 (Transition), 17 (NonStandard), 18 (PulserCal), 19 (ElectronicsCal)

Figure 3.15: Definition of run rankings and run bits. If any of the bits from the "Bad" column are set, for example, that run is automatically ranked as bad regardless of what other bits are set. Runs with no run quality bits set are ranked as gold [86].

- AvsE: Channels with poor AvsE efficiency or stability.
- DCR: Channels with poor DCR performance or stability.

These cases are handled via channel selection files, which are used by in the production of skim files and in the calculation of the exposure. We flag detectors as veto-only or bad (i.e. completely ignored), depending on the nature of the issue [86].

Additionally, if energy or pulse-shape parameter calibration is unsuccessful for a particular channel in a set of runs, channels can be flagged as veto-only or bad and then accounted for in the exposure calculation. Finally, if a channel has abnormally low data cleaning efficiency, it may be removed from the physics data set for some number of runs. This is summarized in [82].

A campaign to recover DS0-6a exposure lost by run and channel selection in this analysis was undertaken since the previously released analyses. This campaign includes improvements in channel selection for the `avse_corr` and `dcr_corr`, described in [1] and [79], respectively. In addition, the run selection was expanded to include runs that were previously rejected

due to the presence of potentially disruptive work, but were manually inspected to check that the data quality is on par with the automatically accepted runs. Run selection also expanded to include runs rejected due to the "partial shield" bit, when the missing shield components were judged to have negligible impact on backgrounds. These improved run and channel selection criteria were also applied to DS6b-8. This procedure is described in [85].

3.12.5 Pulser and Granularity Cuts

The MAJORANA DEMONSTRATOR analysis uses pulsers to monitor detectors and estimate the livetime for the experiment. The pulsers in principle have fixed energies and fixed time stamps. However, because of the crosstalk and the instability of electronics, the pulser amplitude can occasionally be lower or a pulser event can be dropped. Pulser events are removed by making an energy cut on events in coincidence with the pulser event flag [40]. This pulser tag is a global event cut, which rejects all operating channels within 8 μs window of the event. The pulser amplitudes are tuned to be near 600 keV, so that no events are lost at the ROI. However, there is a small dead time due to the cut using the coincidence with the pulser tag. In order to calculate this re-trigger deadtime, we need to calculate the pulser count for each channel run by run, because the livetime estimation depends on the channel and run selection. There is an algorithm described in [40] to tag, and therefore count, pulser events. Basically, one finds where most events happen in the pulser spectrum, consider an energy range of 10 keV on either side of the peak, calculate the mean and RMS of this region and consider the uncertainty as $\frac{RMS}{\sqrt{N}}$, where N is the number of counts.

$0\nu\beta\beta$ events in the ROI are assumed to occur as single-site events. Events with hits in multiple detectors are automatically discarded by the multiplicity analysis cut "mHL==1". If one channel triggers, e.g. due to a noise event, any channel which triggers within the event-building coincidence resolving time of 4 μs is rejected, since it appears as a multiplicity (mHL) > 1 event. All detectors then have an effective dead time for a trigger in any detector above the energy for contributing to the granularity (5 keV). These multiple-detector events are most frequently pulsers (whose deadtime is already accounted for) or multiple-scatter γ

events, which have a very low rate during background data taking. The combined pulser and granularity dead time can be estimated using the total number of triggers multiplied by twice the event-building coincidence resolving time of 4 μs. The granularity cut typically introduces a correction of less than 1 second per dataset, and has been declared negligible in the final calculation and is described in further detail in [50].

3.12.6 Hardware Dead Time

For very brief periods, the digitizer hardware is at times unable to trigger and record an event. Due to the low data rate, it is also possible that the DAQ could stop taking data for a period of time without an obvious effect. To calculate the extent of this deadtime, for run subsets defined between instances of running the threshold finder, the pulser period of each channel is measured. The expected number of pulser events is compared with the actual number found in each channel. Events are "lost" when they occur within a digitizer dead period, as when the GRETINA card has already triggered - most likely on a noise event or "phantom trigger" [50]. Another source of hardware dead time results in the repeated trigger of a GRETINA channel on a polarity which is not recorded. Based on the value of the internal baseline parameter when the channel is initialized, there will be some unknown rate of negative triggers. In cases where the opposite polarity is triggering at a very high rate without being read out, there is an induced dead time from the effect described above. Specifically, due to the way the triggering logic is implemented in the FPGA, if a channel triggers on a negative-going pulse (or noise fluctuation), it can't trigger on a positive-going pulse until a specified time later, dependent on the on-board trap filter settings. To minimize this dead time in later data taking, the threshold finding algorithm sets the thresholds for each channel based on the rate on each channel for positive polarity triggers. The ORNL analysis [76] package is used to monitor and correct for this dead time based on the fraction of missed pulser events on each channel. This is the "hardware deadtime" which is described in [50].

An important aspect of the Majorana Demonstrator analysis is the "H or L" mode,

also known as hit swapping, i.e. substituting in a low-gain (LG) channel if the high-gain (HG) channel is not available. This can be applied globally (if e.g., a channel is dead) via channel selection, or at the run level. The livetime code then tracks the exposure totals so that the correct total from each detector is computed. HG hits are usually kept unless they don't exist, are marked as bad or veto-only, or are saturated. There exists a routine that counts the number of pulsers in each channel, for each data subset in the good run lists. The hardware dead time is calculated in 3 modes, including the OR (H or L) mode, which finds the deadtime of the final merged spectrum. The dead fractions for a detector is calculated by taking $\frac{n_{exp}-n_{pulsers}}{n_{exp}}$. n_{exp} is the expected number of pulsers, and $n_{pulsers}$ the measured value thereof. In the event that HorL deadtime is bad, the HG hardware deadtime is preferentially substituted over the LG. The hardware deadtime varies by detector, ranging from 0.01% in DS6b to 3.99% in DS1 [50].

To compute the final exposure and exposure uncertainty for each detector and each DS, we consider three main sources of uncertainty: the active mass uncertainty, which is dominant, and the livetime uncertainties σ_{HG} and σ_{LG}. The total exposure uncertainty for the full dataset is then expressed as a linear sum of the individual detector uncertainties, as the uncertainties in the active mass and in the runtime are highly correlated [50].

3.13 Cuts applied

With the updated analysis routines, `Final_Energy`, `lq`, `avse_corr` and `dcr_corr` are the final parameters used for the analysis. These parameters each use additional corrections that were not applied in previous analyses. The following string includes all cuts which are applied to the background data for PPC detectors:

```
!isMuVeto && !(isLNFill1&&C==1) && !(isLNFill2&&C==2) && mHL==1 && isGood
&& isEnr && !isICPC && wfDCBits==0 && avse_corr>-1 && dcr_corr<2.326 &&
avse_corr<9 && lq<5 && Final_Energy>100
```

The `isEnr && !isICPC` cuts admit only data from enriched PPC detectors, and not ICPC

detectors, which are normally analyzed separately. Other cuts were described in the sections above.

3.14 Summary

The design of the MAJORANA DEMONSTRATOR and the analysis methods were developed in order allow the experiment to achieve great performance in searching for $0\nu\beta\beta$. Background rejection methods allow us to achieve a very low background index in our region of interest. The MAJORANA DEMONSTRATOR is able to tag surface events and discriminate between such events and bulk events. Surface events originate outside the active mass of our detector and therefore are unlikely to be $0\nu\beta\beta$ events. DCR and LQ are designed to catch events with high delayed charge and slow pulses, which indicate surface events. High AvsE is able to cut events close to the point contact, acting as a volumetric cut. As such, this cut also handles surface events that DCR and LQ might miss. The experiment can also tag multi-site events, which cannot be $0\nu\beta\beta$ events, as the latter occur in one place in our detector. The low AvsE cut is our primary multi-site rejection factor. The MAJORANA DEMONSTRATOR's excellent energy resolution of 2.5 keV FWHM at $Q_{\beta\beta}$ allows us to probe a very narrow region of interest around our 2039 keV peak, and get a more accurate count of events which could originate from $0\nu\beta\beta$ events. All of these techniques improve (and in fact, make possible) the performance of MAJORANA DEMONSTRATOR, allowing the experiment to set limits on $0\nu\beta\beta$ half-life and the mass $m_{\beta\beta}$, as discussed in Chapter 5. The spectrum shown in Fig. 3.16 shows the importance of our main analysis cuts, DCR, AvsE and LQ.

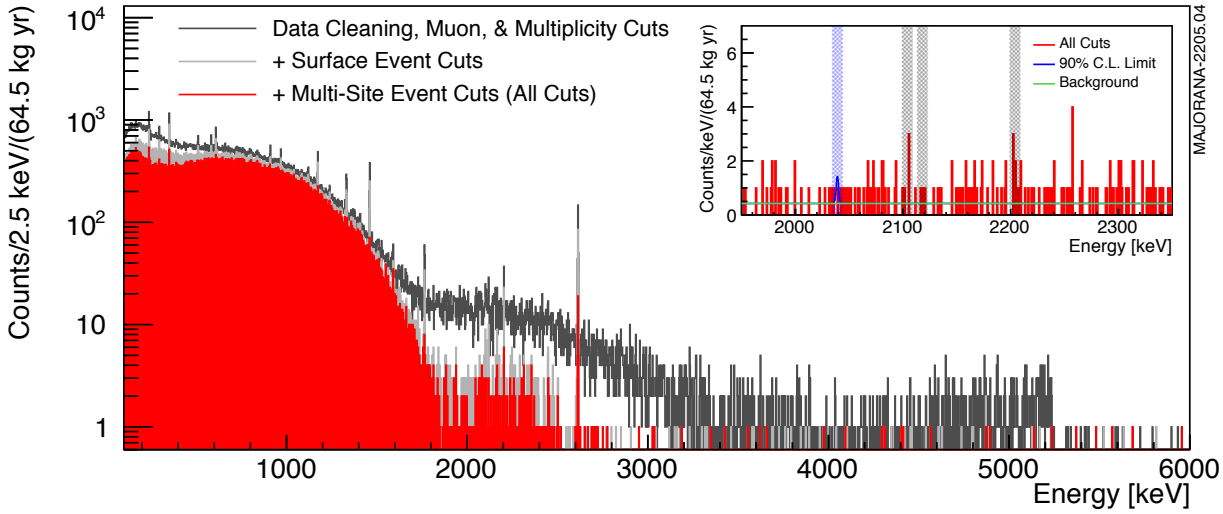

Figure 3.16: Spectrum after all cuts are applied. Currently from the 2022 MJD release [26]. The data cleaning cut is wfDCBits==0 in Section 3.13, whereas the muon cut is $!ismuVeto$ and the multiplicity cut is mHL==1. The surface event cuts are dcr_corr < 2.326 & avse_corr < 9 && lq < 5, whereas the multi-site cut is avse_corr > -1.

Chapter 4

DCR

A detailed description of the delayed charge recovery (DCR) pulse-shape discrimination in the MAJORANA DEMONSTRATOR is provided in this chapter. DCR is used to identify background events originating from radioactive decays resulting in alpha particles which deposit their energy on the passivated surface of PPC detectors,. These problematic background events, which can be highly degraded in energy and therefore populate the $0\nu\beta\beta$ region of interest (ROI), exhibit distinctive pulse shapes due to slow charge collection. This chapter describes the implementation, efficiency, and uncertainties associated with the DCR discriminator. Specifically, the variable name is dcr_corr and the cut value is 2.326, i.e. the cut as implemented means rejecting events with

$$\texttt{dcr\_corr} > 2.326$$

This variable includes a drift time correction to DCR, as will be discussed in the following sections.

4.1 Potential Origins of the Delayed Charge Recovery Effect

The response of PPC detectors to events near the passivated surface is difficult to predict. High charge trapping has been observed on similar surfaces in BEGe-type and segmented Ge detectors [13] [10], but the charge collection properties near this surface can vary for different types of detectors [27, 75, 102]. In the MAJORANA DEMONSTRATOR, events have been observed in which alphas originating on the passivated surface are significantly degraded in energy. This loss of energy leads to a potential background contribution in the ROI for $0\nu\beta\beta$. It is also observed that charge mobility is drastically reduced on or near the passivated

surface. The trapped charge is slowly re-released on the timescale of waveform digitization, leading to a measurable increase in the slope of the tail of a recorded pole-zero-corrected pulse.

The alphas observed are due to ^{210}Po exposure, ultimately originating from ^{222}Rn, which is a very common isotope. This isotope decays to ^{206}Pb, releasing an alpha particle with an energy of roughly 5.4 MeV, impinging on the passivated surface of the MAJORANA DEMONSTRATOR detectors. That said, as discussed in Sec. 4.9, the energies that we collect for most of these alphas are degraded, with the high-energy alpha peak around 2000-2500 keV. While great efforts to minimize radioactivity have been done, such as by electroforming copper underground, not all parts are entirely radiopure. It is suspected that the plastic parts used in assembling detector holders, as high-voltage insulators, have some level of radioactive contamination.

There are two potential models for charge collection in the passivated surface region. One possibility is that this effect is due to surface propagation of the electron contribution to the signal, with the holes being collected normally. This matches the model developed in [97]. In an alternative model, some fraction of the holes are trapped when they originate in a few-micron-thick region at the passivated surface, and are then slowly re-released.

In both cases, part of the the energy of the event appears as a normal, fast pulse, and the remainder of the charge is collected slowly; see Fig. 4.1 for sample waveforms, simulated using the `siggen` software package [106]. These waveforms are generated using the electron surface drift model, but the waveforms appear similar regardless of the cause of the delayed charge.

For the purpose of identifying alpha events, the cause of the delayed charge is irrelevant. By using a parameter that can identify the occurrence of this delayed charge recovery (DCR), surface alpha events can be tagged, allowing for their efficient rejection in offline analysis. The goal of such a filter is to detect the presence of slow charge collection occurring after the bulk charge collection has been completed. In a waveform that has been fully corrected for the electronic response function, also known as pole-zero (PZ) corrected waveforms, this

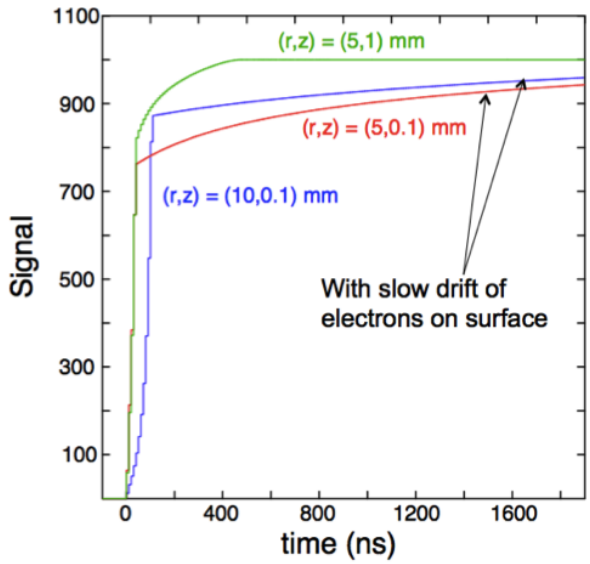

Figure 4.1: Bulk (shown in green) and passivated-surface (shown in blue and red, for varying radial positions of charge deposition) waveforms, created using `siggen`, a dedicated software package used to simulate signals in semiconductor detectors. Surface charge transport of electrons is induced by incorporating an arbitrary small amount of passivated-surface charges in the model of the detector's electric field, which leads to field lines that carry charge to the passivated surface. Figure courtesy of David Radford.

appears as a positive slope of the tail. Normal events should have a DCR distribution centered at zero. This implies that a cut on the high side of this distribution will remove alphas. However, the mean can be subtly different from zero, and it could potentially change over time, and vary with energy. Similarly, the width of the distribution can also be different for different calibrations (time) and is observed to vary with energy. Charge trapping also occurs in the bulk, and in this case it is proportional to the drift time of the charges to the contacts of the germanium detector. This effect is detector dependent but widens the DCR distribution in a significant proportion of detectors.

4.2 *Implementation in the* Majorana Demonstrator

For the DCR analysis cut, the relevant parameters are found using calibration data, since these runs contain a negligible fraction of surface events, if any at all. A description of the processing steps and the efficiency and uncertainty results are given in this chapter. Currently, pile-up and pulser cuts are implemented, alongside a hard energy cut of trapENFCal > 100 keV. Charge trapping in the bulk of the detector is corrected for, by using the drift time parameter. The low AvsE cut is applied when calculating the parameters for the charge-trapping corrected DCR, since we desire to compute the average DCR distribution of single-site events, as these are the important events for the $0\nu\beta\beta$ analysis. As such, it is important to note that, e.g., the calculated DCR efficiency is always after applying the AvsE cut.

The strategy for calculating DCR is described briefly below, with details left for the following subsections. Ideally, DCR would have a Gaussian distribution, with a mean of approximately 0 and a width, or standard deviation, of approximately 1. This is helpful because then any choice of cut value would give a similar acceptance for all channels and runs, without the need to re-tune, unlike in the previous version of DCR used in the 2019 & 2018 Majorana Demonstrator releases [21, 2]. An example DCR distribution for calibration data is shown later, in Fig. 4.5. As intended, the main feature of this distribution is a Gaussian peak centered around 0, with a standard deviation close to 1.

4.3 *Pole-Zero Correction and Slope Calculation*

We start by applying a pole-zero correction to the baseline-subtracted, ADC nonlinearity-corrected waveforms, as justified below. Using such a correction in the DCR algorithm, the stability uncertainty of the DCR parameters could be reduced by re-tuning the decay constant for each detector with each weekly energy calibration of the Majorana Demonstrator. In the previous version of DCR, multi-hour calibration runs were needed to re-tune the parameters. Therefore, physics live time considerations prevented regular re-tuning –

that is, at collaboration level, it was decided to have limited calibration time in order to maximize the time used for physics runs. On the other hand, true pole-zero correction can be determined using very few (less than 500) events, allowing DCR to be re-tuned using the already-established weekly energy calibrations. An overshoot appears in the waveforms, and must also be corrected for when applying the PZ correction.

First, using waveforms from each calibration (short and long), we compute the pole zero correction parameters, including

- tau, which represents the best fit to the long time decay of circa 70 μs,
- tau2, a secondary decay time, of the overshoot, fixed to 2.1 μs, and
- frac2, the fractional amplitude corresponding to the overshoot.

The pole-zero parameter calculator, developed by one of the ORNL Collaborators, David Radford was tuned to make the ORNL version of DCR flat, and the differences between algorithms is responsible for much of the energy dependence seen in DCR. As such the slope is not actually clustered around zero, and a further correction is needed, which is described in the following section (it is the shift of the mean of the distribution to zero). The parameters are calculated using raw files, chained together for each calibration, and are stored in the database. The pole-zero correction is then applied to the non-linearity-corrected baseline-removed waveforms (nlcblrwf), to get the transfer-function-deconvolved waveforms (tfdwf).

4.4 DCR Calculation

We compute DCR as described in Sec. 4.2, except we now use the pole-zero corrected waveforms, unlike for the previous MAJORANA DEMONSTRATOR releases [21, 2]. This means that the average slope of the waveform (which defines DCR) is close to 0. The DCR parameter is found by first calculating the slope of the waveform tail for each event in a channel. 1 μs of the waveform is averaged, corresponding to 100 waveform samples in non-multisampled data, at each of two points on the waveform. First, the 97% rise point of the waveform

is determined by finding the maximum height of the pulse before pole-zero correction, and stepping forward from the start of the digitization window until 97% of that maximum value is reached. The first region lies between 2 and 3 μs after this 97% rise point, and the second is the final μs of the waveform (in the non-multisampled data, this is between 19 and 20 μs), as in Fig. 4.2. These values were chosen to avoid introducing unwanted sensitivity to the shape of the waveform turnover, i.e. a starting value of 2 microseconds was chosen as it was deemed far enough away from the rising edge of the waveform, and to decrease sensitivity to noise, leading to a choice of 1 μs, which is a significant, but not large, proportion of the waveform on which to average. 97% was chosen due to possible noise effects when considering the maximum of a given waveform.

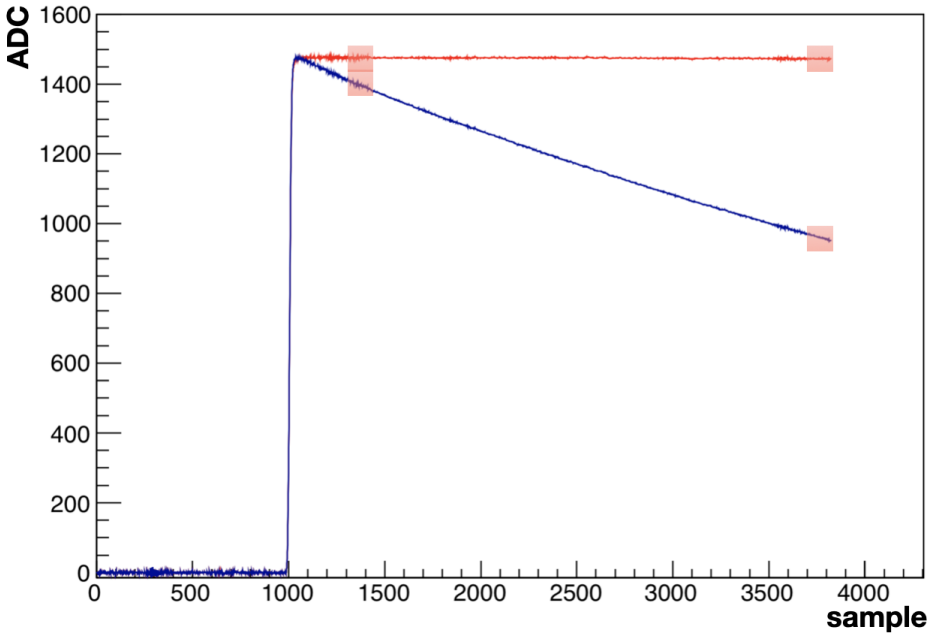

Figure 4.2: A sample MJD waveform, not pole-zero corrected in blue, and pole-zero corrected in red. The ADC values are averaged in each of the two shaded regions. The slope between those average points is taken to be δ. Each sample corresponds to 10 ns.

We proceed to calculate the parameters α, β and γ, which are used to determine DCR.

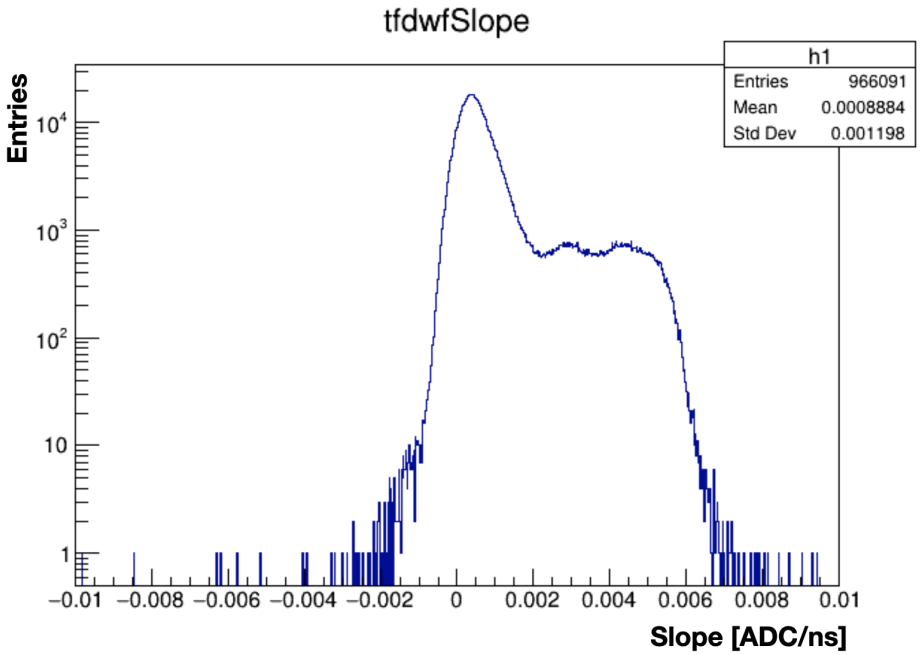

Figure 4.3: Slope distribution for channel 626 in the first long calibration in Dataset DS6b.

Plotting the mean and width of the distribution of the slope of the transfer function deconvolved waveforms (tfdwf) (Fig. 4.4) versus the calibrated energy (trapENFCal), we find a linear correction of its mean with energy, of form $\mu = \alpha E$ (see Fig. 4.7). Applying a linear fit appears to be well-motivated with the residuals appearing not to have any particular functional form. The performance with this choice was good, with most DCR distributions having a mean of roughly 0; see Fig. 4.5. Also note that a further mean correction is applied during the drift-time correction process, if the mean is found to diverge from zero. Linear and constant terms are used in the width σ of the distribution of form $\sigma = \beta E + \gamma$ (see Fig. 4.6). Thus, the DCR parameter is defined as:

$$\mathrm{DCR} = \frac{\delta - \alpha E}{\beta E + \gamma} \tag{4.1}$$

Correcting for μ and σ is done (by obtaining the best-fit values of α, β and γ) such that the new DCR is, ideally, Gaussian, with a mean of approximately 0 and a width, or standard

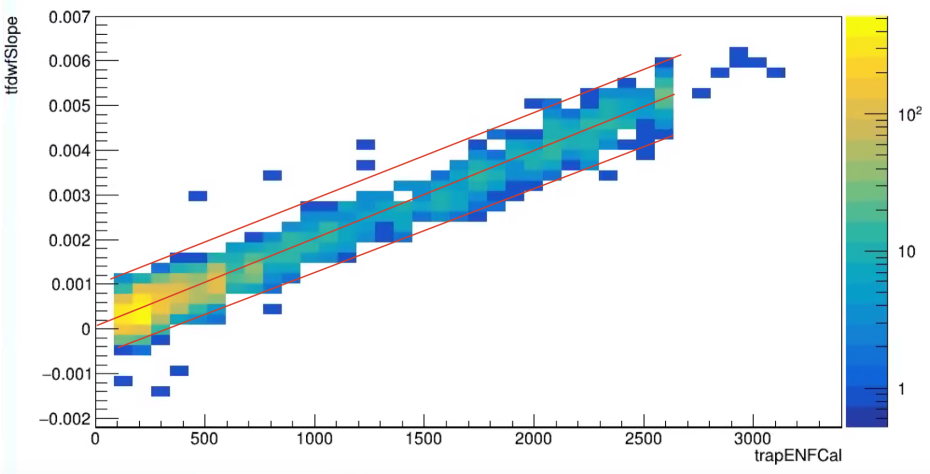

Figure 4.4: Slope of the pole-zero corrected waveform vs calibrated energy, for channel 626 in the first long calibration in Dataset DS6b. The red lines guide the eye to show a linear correction for the mean and a linear and constant correction for the width are appropriate.

deviation, of approximately 1. This is helpful because then any choice of cut value would give a similar acceptance for all channels and runs, without the need to re-tune, unlike in the previous version of DCR used in the 2019 & 2018 MAJORANA DEMONSTRATOR releases [21, 2]. An example DCR distribution for calibration data is given in Fig. 4.5. For each calibration, we determine the value of the parameters α, β and γ by applying a linear fit to the mean vs energy and sigma vs energy distributions in the 6 energy bins (see Fig. 4.6 and Fig. 4.7). These values are then saved into the MAJORANA DEMONSTRATOR parameter database (MJD DB) and used in the next stage of data processing, for each waveform, given its features (slope and energy). As intended, the main feature of this distribution is a Gaussian peak centered around 0, with a standard deviation close to 1. While there are low and high DCR tails, these are quite small. The high DCR tailing could occur from, e.g., mistagged pile-up events, which if missed are expected to have high DCR. Such events do not dominate the distributions and pose no actual problems, as we expect them to be high in DCR.

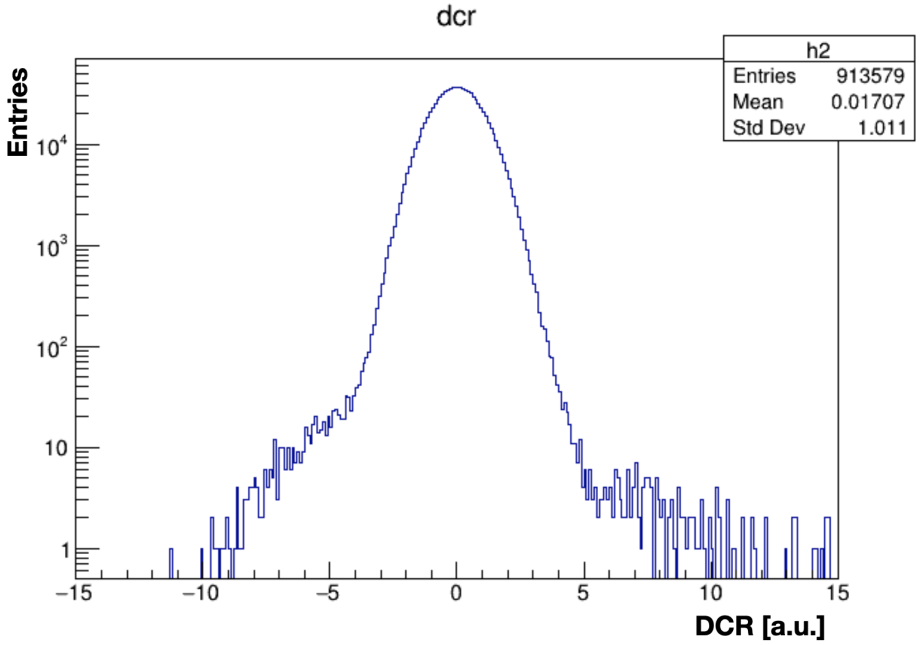

Figure 4.5: DCR distribution showing its mean and sigma compared to 0 and 1 respectively, for test channel 626, for the first long calibration in DS6b.

4.5 Drift Time Correction

Work has been done on a waveform-by-waveform charge-trapping correction, which uses the drift time of the pulse to calculate the expected amount of charge lost in the bulk. The idea behind this correction can be visualized in Fig. 4.8. Since there is, for many channels, some correlation between DCR and drift time, we can use the relationship between them to define a corrected parameter, for which the distribution is narrowest. Figure 4.9 shows how the drift time relates to DCR in a typical calibration. A first look at this correction indicates a circa 25% narrower distribution of DCR (see Fig. 4.10) and a more Gaussian shape for many but not all channels. The difference in performance is due to intrinsic properties of the individual detectors – some of which are known to have a significant charge trapping effect. Note that the spread out, double peaked distribution in DCR has been transformed into a single peak,

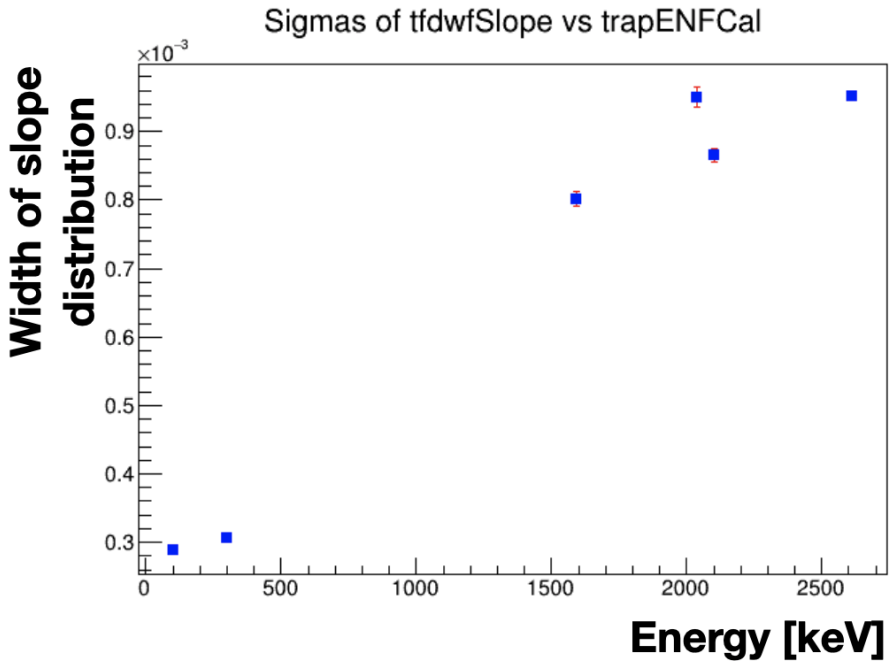

Figure 4.6: Example (ch.672, 2nd long calibration in DS6b) relationship between the sigma of the tfdwfSlope (δ) distributions and energy. The error bars are drawn but not visible for the leftmost two points due to better statistics.

quasi-Gaussian, albeit with a long high-DCR tail, distribution in $\mathrm{DCR_{corr}}$. Similar results are observed for other channels with significant charge trapping effects. The effects of applying a drift time correction range from 0 to somewhere around a 25% narrower DCR distribution, depending on the channel, as can be seen in Fig. 3.9. This correction provides more effective surface alpha discrimination for many of the MAJORANA DEMONSTRATOR detectors.

The drift time correction is computed as follows. For each channel and calibration subset, we look at all events that pass certain cuts, such as energy (trapENFCal) between 100 and 2700 keV, single-site events ($avse > -1$), and the pulser and pileup cuts. Furthermore, to look at the bulk of the dcr vs drift time distribution, we focus on events with $-5 < \mathrm{DCR} < 5$ and the drift time, here taken to be the 0 to 90% rise time, t_{0-90}, between 2000 and

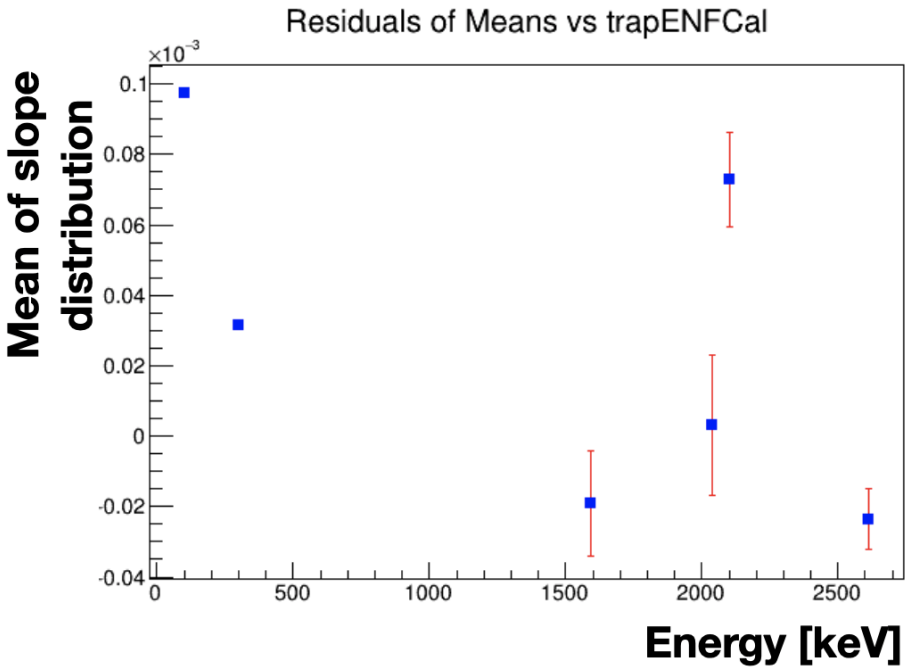

Figure 4.7: Example (ch.672, 2nd long calibration in DS6b) relationship between the residuals of the means of the tfdwfSlope (δ) distributions and energy. The error bars are drawn but not visible for the leftmost two points due to better statistics.

5000 ns. Then, we subtract the mean of the DCR distribution and mean of the drift time distribution (as calculated by numpy) from DCR and drift time respectively. This yields a distribution centered at $(0, 0)$. We calculate the rotation in DCR-drift time space needed to have as narrow a $\mathrm{DCR_{corr}}$ distribution as possible, using Principal Component Analysis (PCA), implemented in numpy and scipy. The output is the vector representing $\mathrm{DCR_{corr}}$, corresponding to a slope in DCR-drift time space. A simple visual representation of what we aim to achieve is given in Fig. 4.8.

The DCR mean (μ_{DCR}), drift time mean ($\mu_{t_{drift}}$), and slope (m) parameters are calculated for each channel and for each calibration subset, and saved and input into the

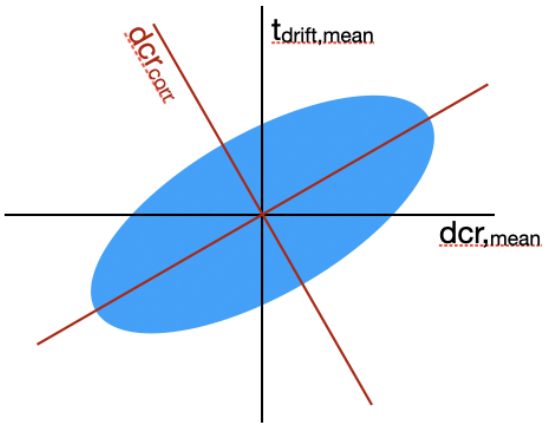

Figure 4.8: A representation of the relationship between the original DCR and the corrected DCR. The vector defining DCR_{corr} is calculated such that the distribution is as narrow as possible.

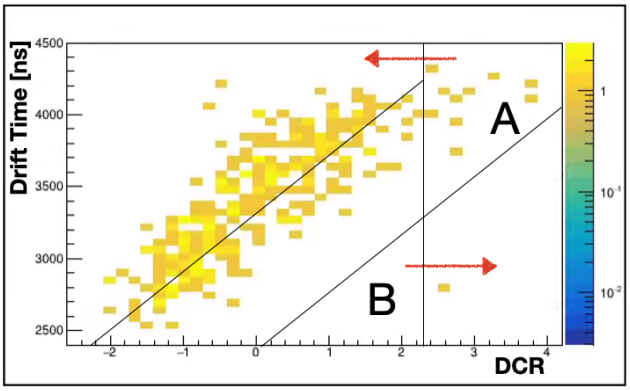

Figure 4.9: The relationship between the drift time and the DCR in one calibration. We are no longer cutting events in area A, with normal DCR values for their drift time, but we are cutting events in triangle B with high DCR values for their drift time, which more likely to be alphas.

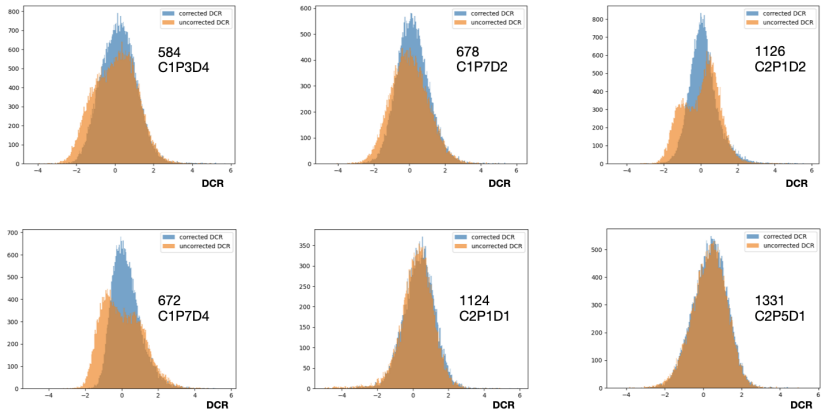

Figure 4.10: A comparison of DCR and DCR_{corr} distributions for a number of detectors for a long calibration in DS6b. The y-axes are number of counts, as these are histograms.

DCRv2DriftTime Database. The preliminary applied correction $\text{DCR}'_{\text{corr}}$ is then

$$\text{DCR}'_{\text{corr}} = (\text{DCR} - \mu_{\text{DCR}}) - \frac{t_{\text{drift}} - \mu_{t_{\text{drift}}}}{m} \tag{4.2}$$

The $\text{DCR}'_{\text{corr}}$ distribution is again fit to a Gaussian to extract final corrections for the mean μ_c and width σ_c. The final value for DCR used in the analysis, DCR_{corr}, is then

$$\text{DCR}_{\text{corr}} = (\text{DCR}'_{\text{corr}} - \mu_c)/\sigma_c \tag{4.3}$$

4.6 Channel Selection

DCR channel selection is applied in two separate procedures. One is performed via analysis of calibration data, and the other looks for large instabilities in DCR using the parameters in the database. Channel selection is required as there can be cases in which we do not obtain parameters, they are unstable (vary between consecutive calibrations), or have poor performance, as will be described in detail in the following subsections. If any of the above happens, DCR is considered not well tuned for a given range of our data for some channel.

Then, this range and channel are not used for physics analyses, due to the existence of potential problems. This is referred to as channel selection. It is needed to not use problematic data for our physics analyses.

4.6.1 Calibration-based channel selection: parameter validity

This aspect of DCR channel selection is implemented at the same step as computing the drift time correction parameters. We have implemented the following process. If, for some channel, the acceptance for a given calibration fluctuates from the average, by more than 4σ upward, then the calibration is added to the DCR Channel Selection (CS) file, for that particular channel. Since this analysis is performed on the same calibration data that was used to tune DCR, this channel selection catches problems where the DCR parameters have failed to sufficiently capture the shape of the DCR distribution. This happens very rarely due to the 4σ requirement, but could happen due to, e.g. very low statistics in one of our short calibrations.

The output file is a .json file which is then used in the official Channel Selection. This process is done for each channel in each dataset, using the exact same cuts as in the drift time correction analysis. Channels which are known to have poor performance overall are also added to the CS file, manually. The total exposure removed by DCR that is not removed by other channel selection is on the order of 0.02 kg years for enriched detectors, and 0.08 kg years for natural detectors.

4.6.2 Database-based channel selection: DCR instability

Since DCR is tuned using each short calibration, we cannot use the calibration data to assess the stability of DCR without reprocessing all calibration data with previous-calibration DCR parameters. Since this is not only computationally intensive but would also require significant disk space, it was decided instead to use the parameters in the database describing the DCR distributions and evaluate how the DCR distribution changes from parameter set to parameter set based on the parametric dependence of DCR on each parameter. To model and

quantify the shifts, we use a Gaussian approximation for the DCR distribution: parameter shifts are modeled as resulting in changes in the mean and width of the DCR distribution. If the distribution for a particular channel shifts by too much from one calibration to the next, we select out that channel for the corresponding data subset prior to the shift.

The relevant parameters are those defined in Eqs. 4.1, 4.2, and 4.3. All of the parameters in these equations appear explicitly in database records, however for DCR v2, the value of δ computed for each event depends additionally on the long PZ time constant τ. Mathematically, δ is the slope of the PZ-corrected flattop, in units of ADC/ns. If τ shifts by a small amount $\Delta\tau$, the value of δ will shift by an amount

$$\Delta\delta = \frac{\Delta\tau}{\tau^2} E_{\text{uncal}} \tag{4.4}$$

We thus track shifts in τ in addition to the other parameters α, β, γ, μ_{DCR}, $\mu_{t_{\text{drift}}}$, m, μ_c, and σ_c. Since we require the parameters in order to compute the shifts, we also require the parameters to pass a series of validity tests. Periods with invalid parameters are also channel selected out for DCR.

Database parameters

The database parameters for a particular channel are subjected to the following tests [79]:

- The channel has parameters (the channel maps are used to determine which channels should be present for a given run range).

- PZ τ is within 10 μs of the default value 72.5 and the PZ fraction is within ± 0.005 of the default value 0.007.

- PZ τ and the PZ fraction are not both equal to their default values (indicates PZ estimation failure).

- Parameter records do not cross a DS boundary when detectors were biased down

(e.g. parameters should change between DS0 and DS1; however, some overlap in DS5 and DS6 is unavoidable).

- When PZ τ parameters change, the DCR parameters α, β, and γ should also change.

- When the DCR parameters change, the $\mathrm{DCR_{corr}}$ parameters μ_{DCR}, $\mu_{t_{\mathrm{drift}}}$, m, μ_c, and σ_c should also change.

Most of the invalid periods discovered via these checks in the initial pass were recovered by reprocessing. In the final pass, the only channels failing these tests were known bad veto-only channels or other channels which were simply impossible to calibrate. The removed enriched livetime was only 0.2 kg y, with overlap with other analyses.

Sometimes, a shift in one parameter (e.g. τ) can be compensated directly by a shift in another parameter (e.g. α). What matters in the end is how all the shifts sum together to move the cut position on the distribution, i.e. the total shift $\Delta D_{c,\mathrm{tot}}$. Figure 4.12 shows the distribution of $\Delta D_{c,\mathrm{tot}}$ for all channels for all run periods in the range $|\Delta D_{c,\mathrm{tot}}| < D_c$. The typical random shift is on the order of 0.6.

We apply channel selection based on $\Delta D_{c,\mathrm{tot}}$. A data subset for a channel is selected out if it is succeeded by an upward or downward shift by more than D_c itself (2.326). An upward shift by more than this amount would correspond to the center of the distribution moving up to the cut value, giving a loss of efficiency by more than 50%, for the affected channel during the given run period. For channels passing the cut, such efficiency losses are accounted for in the systematic uncertainty of the DCR cut efficiency, as described below. Downward shifts, on the other hand, do not amount to efficiency losses but begin to let in more background. The choice of cut was made to balance this heightened background risk against the exposure loss due to the channel selection (at the percent level). The open data contains only a handful of counts in the range $D_c < \mathrm{DCR} < 2D_c$ (all of which are presumably alphas), but the chance these events occur randomly during one of the downward-shift periods is the same as the percent-level efficiency loss. Since there are tens of background counts total, the increased

background risk thus remains negligible. Applying these criteria to all channels in all runs removes about 3 kg yr, 1.94 kg yr of which is not channel-selected out by other analysts.

This channel selection was not applied for a small set of runs in DS6b, namely the data subsets starting from runs 38757, 39165, 39687, and 40262. Waveform analysis was originally performed for DS6b runs in 2019, but in Fall 2021 a few periods in this region were marked for reprocessing due to parameter instabilities and exposure recovery efforts. However between 2019 and 2021, the estimation method for the pole-zero correction parameter τ had been modified. This results in disjoint sets of parameters being used in neighboring data subsets. Unsurprisingly, we find that the $\Delta D_{c,\mathrm{tot}}$ values computed at these boundaries often exceed the instability threshold for many channels. However, since these discontinuities can be attributed to differing parameter sets rather than detector instabilities, we ignore big shifts at these run boundaries.

There is one other data subset that was not marked for removal despite being followed by big parameter shifts in many channels. This is the data subset starting with run 44926, which preceded the ^{56}Co special calibration that ran over an extended period. The DCR parameters that follow this calibration are expected to be somewhat larger than typical since a longer than usual time had passed since the previous DCR parameter tuning.

4.7 Uncertainties Associated with DCR

For DCR with drift time (charge trapping) correction, the efficiencies are presented in Tables 4.2 through 4.4. First, we show the efficiencies and uncertainties for all active detectors, then just the enriched detectors, and finally just the natural detectors.

4.7.1 Statistical Uncertainty

For the DCR parameter, we calculate the statistical uncertainty as follows. In each channel, the statistical uncertainty of the cut efficiency is given by the expression $\sigma_{stat} = \sqrt{\frac{(\epsilon)(1-\epsilon)}{N}}$, where N is the number of events rejected by the DCR cut in the set of all calibration runs used to set the cut, and in the energy window being used – the $0\nu\beta\beta$ efficiency window of

2028 keV - 2050 keV. We use exposure weighting to get the average statistical uncertainty across every high-gain channel, and further divide by $\sqrt{k}$ where k is the number of channels. Thus, the total statistical uncertainty is:

$$\sigma_{stat} = \frac{\sum_{i=1}^{k} \left(\sqrt{\frac{(\epsilon_i)(1-\epsilon_i)}{N_i}} e_i \right)}{\sqrt{k}\ e_{total}} \tag{4.5}$$

where e_i is the individual detector's exposure and e_{total} the total exposure.

Generally, the statistical uncertainty is on the order of $\pm 0.1\%$, usually negligible compared to other sources of uncertainty. See Table 4.2.

4.7.2 Stability

To study DCR stability over time, calibration skim files are chained together and analyzed. Only calibrations from the Module 1 source are used in the DS0-DS3 & DS7 analysis, and only calibrations from the Module 2 source are used in the DS4 analysis. Calibrations from both sources are used in the DS5, 6 & 8 analysis. All single-site (as determined by AvsE) data that is not rejected by channel selection, and passes the pile-up cut and a pulser cut is used. We use an energy window of 2028 keV - 2050 keV.

We chain together all calibration runs, for each calibration separately. We compute the total number of events passing the basic cuts described above, as well as the number of events passing the basic cuts and the DCR cut. The efficiency and its uncertainty is computed with ROOT's TEfficiency class using Clopper-Pearson intervals. The efficiencies are plotted as a function of the initial run number of each calibration. The results are presented in Figs. A.2 - A.17. A flat line was fit to the data to compute the weighted average efficiency. We conservatively estimate the time stability systematic uncertainty as the weighted standard deviation of the deployment-by-deployment efficiencies about the weighted mean, as computed by the flat-line fit.

A second aspect of stability based on variations in the DCR parameters in the database is described in Sec. 4.6.2. The ΔD_c values can be directly translated into DCR efficiency

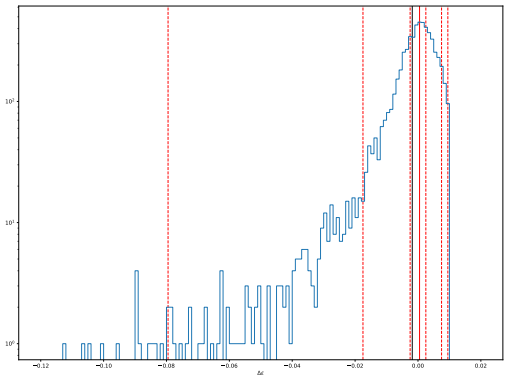

Figure 4.11: The $\Delta\epsilon$ values computed for all detector periods in all datasets using Eq. 4.6 with the ΔD_c values in Fig. 4.12. The vertical black line is the mean of the distribution, while the solid red line is the median, and the dashed red lines show the 1, 2, and 3σ quantiles about the median. Figure courtesy of Jason Detwiler.

uncertainty contributions under the approximation of a predominantly Gaussian DCR distribution. Figure 4.12 shows the distribution of ΔD_c for all channels for all run periods. The computed ΔD_c values there represent the total shifts observed for a channel from the start to the end of a data taking period. Assuming an approximately linear drift model, the time-averaged deviations over the various run periods will thus be only half of the values plotted in this histogram. To compute the efficiency contributions, we divide the ΔD_c by dataset and detector type, and then compute the change in efficiency for each ΔD_c as

$$\Delta\epsilon = -(\mathrm{erfc}((2.326 - D_c/2)/\sqrt{2})/2 - \mathrm{erfc}(2.326/\sqrt{2})/2) \tag{4.6}$$

The distribution of $\Delta\epsilon$, computed from the ΔD_c values in Fig. 4.12, is shown in Fig. 4.11. We then compute the mean $\overline{\Delta\epsilon}$, its uncertainty $\sigma_{\overline{\Delta\epsilon}}$, and the quadrature sum of the two, $\sigma_{\Delta D_c}$. The latter values are tabulated in Tab. 4.1 and are added as a time stability contribution to the total DCR efficiency uncertainty.

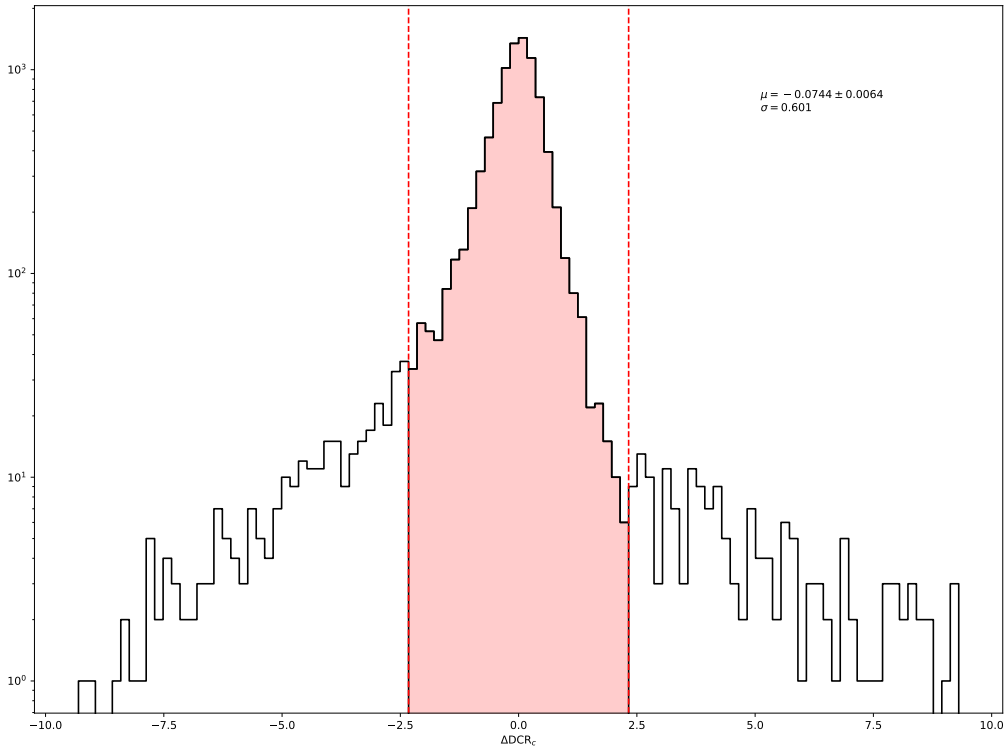

Figure 4.12: D_c shifts for all channels for all run periods. Periods outside the red shaded region are marked for channel selection. The distribution shows an asymmetry that is not understood in detail, but is accounted for by conservatively including the mean of the corresponding change in efficiency as a contribution to the systematic uncertainty. Figure courtesy of Jason Detwiler.

Data Set	Enriched	Natural	All
DS0	0.17%	0.05%	0.12%
DS1	0.14%	0.14%	0.14%
DS2	0.07%	0.14%	0.06%
DS3	0.14%	0.18%	0.14%
DS4	0.14%	0.18%	0.15%
DS5a	0.58%	0.19%	0.42%
DS5b	0.21%	0.19%	0.19%
DS5c	0.21%	0.11%	0.17%
DS6a	0.12%	0.06%	0.09%
DS6b	0.39%	0.24%	0.33%
DS6c	0.14%	0.15%	0.14%
DS7	0.22%	0.21%	0.21%
DS8 PPC	0.15%	0.09%	0.12%
DS8 ICPC	0.23%	N/A	0.23%

Table 4.1: DCR cut efficiency uncertainty contribution from DCR parameter time instabilities, as evaluated from database records.

4.7.3 Pulse Shape Bias

Irregularities in the pulse shapes of events, particularly charge-trapping re-release and multi-site effects, can have an effect on the calculated DCR acceptance. Events with a transition layer multi-site component, multi-site events occurring very near the point contact, and events with high charge trapping, like those seen in Fig. 4.13, may be untagged by A vs. E and data cleaning analyses, but will be accepted by the DCR cut with lower-than-average efficiency. In Fig. 4.13, the event in blue is thought to be a near-point-contact multi-site event, where the late charge arrival from the second site gives it an incorrect energy for its tail slope. The one in red is thought to be a transition-layer multisite event or an event with high bulk charge trapping, either of which would contribute the additional slow component that changes its tail slope. All events are normalized by energy. We calculate a Pulse Shape systematic uncertainty estimation. We compare the acceptance near the ^{228}Th double-escape peak (DEP), specifically in the 1580 - 1604 keV range, to that in the ROI energy window, 2028 keV - 2050 keV. The difference is the pulse shape-dependent bias, σ_{PS}. As before, the overall bias is given by an exposure-weighted average of the bias in each channel, since σ_{PS} is assumed (conservatively) to be fully correlated among the detectors. For context, in dataset DS6b and a cut value of $DCR_{corr} < 2.326$, the systematic PS bias is 0.34%. More details can be seen in Table 4.2. The pulse shape bias indicates that the DCR analysis is effectively tagging some fraction of the multisite gamma events that go untagged by the A vs. E analysis. This implies that the DCR analysis has power for rejecting background gamma events in addition to alpha events. The inclusion of this bias as a systematic uncertainty also accounts for the possible fraction of calibration events that exhibit true passivated-surface delayed charge recovery. Since these events are degraded in energy by definition, they do not appear in the DEP, but do appear in the Compton continuum. The acceptance in the DEP therefore provides a measure of the acceptance of true bulk events.

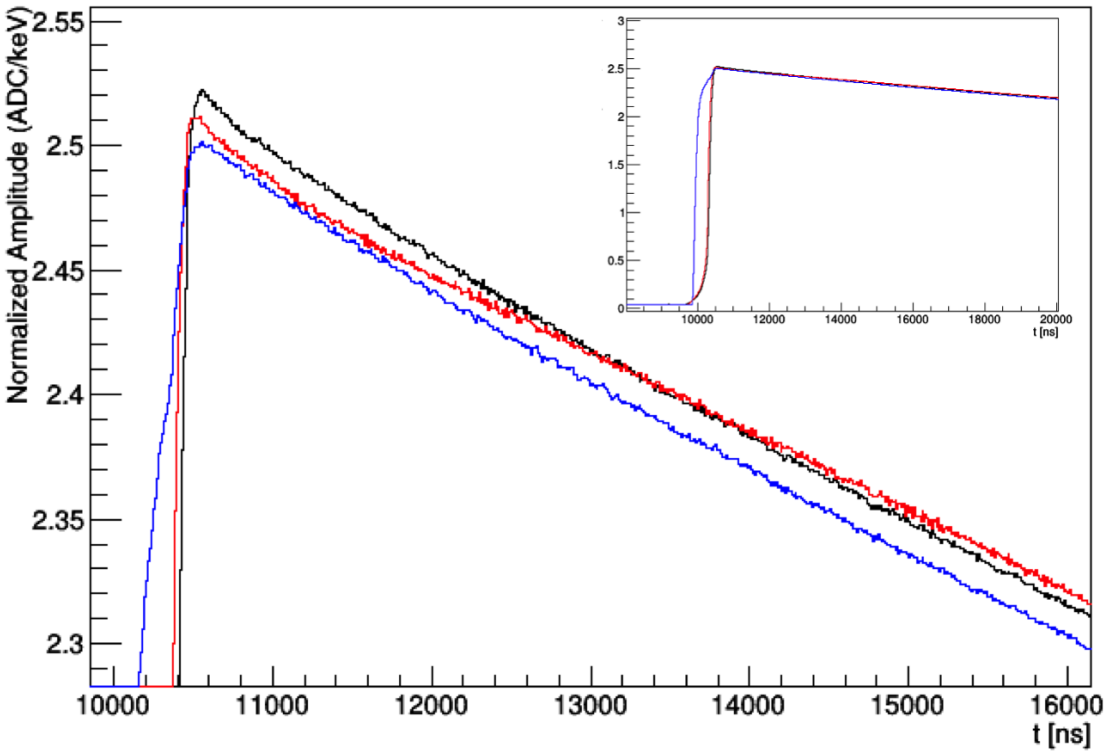

Figure 4.13: The effect of pulse-shape on DCR, in one channel, in the previous version of DCR, presented strictly as motivation for a pulse shape bias investigation. The event drawn in black passes the DCR cut, those in blue and red fail the cut. Figure courtesy of Julieta Gruszko.

4.7.4 Other aspects of DCR systematic uncertainty

Other sources of systematic uncertainty can be considered. One source is residual ADC nonlinearities that cause a saw-tooth behavior in the DCR efficiency as a function of energy. Another is a possible energy dependence in DCR due to mis-tunings / time instabilities of some of DCR's parameters that tune its energy dependence. Both of these source of uncertainties are already accounted for in σ_{PS} since it compares the DCR acceptance for Th DEP events with events near $Q_{\beta\beta}$, and those differences would already incorporate both of these effects.

Another aspect worth noting is that the Th compton-continuum events near $Q_{\beta\beta}$ used to computed the central value for the efficiency contain some fraction of events with significant energy deposition in the transition layer, that then lead to a high DCR tail that is not expected for $0\nu\beta\beta$ events. A study by the ORNL code reveals this to give a $\sim$1% downward bias in the efficiency estimate used by the GAT analysis. This difference again is also captured in σ_{PS}, but one might argue that this aspect should give an asymmetric "upward" contribution to the uncertainty. However, we also note that a similar argument could be made for an asymmetric "downward" contribution due to the large negative-$\Delta\epsilon$ tail seen in Fig. 4.11 that is accounted for in the σ_{stab} by the inclusion of the mean $\overline{\Delta\epsilon}$ as a direct contribution. The combination of these effects thus results in an uncertainty that to first order is roughly symmetric, and is conservatively already accommodated by the quadrature sum of σ_{PS} and σ_{stab}.

4.8 Performance of DCR

4.8.1 Efficiencies

We set the DCR cut to 2.326, determined by wanting a roughly 99% acceptance of calibration events. The value of 2.326 arises from a one-sided 99% cut in a Gaussian distribution. Events with DCR $>$ 2.326 are rejected. There are often DCR tailing effects, both high and low DCR, depending on the channel. The reason for which these appear are not fully

understood. However, specifically the high-DCR tails are likely to decrease the efficiency of the 99% cut; indeed, with the drift-time correction included, the acceptance is circa 98.5%. An optimization study was done, which set the value of DCR to 2.326 for alpha rejection – a further cut, $AvsE < 9$, was determined to be desired and is used in the analysis [1].

The DCR efficiencies are presented starting in Table 4.3. These are simply the fraction of events which pass the DCR cut for calibrations in each dataset; other cuts applied when computing the efficiency (in both numerator and denominator) are the low AvsE, isGood, event and waveform-level data cleaning cuts. Table 4.3 presents the final enriched numbers for DCR to be used in the main analysis. The DEP efficiencies presented have been calculated using background subtraction, using sidebands at 1568-1580 keV and 1604-1616 keV.

Table 4.3 logs the efficiencies and the uncertainties of DCR for all our distinct datasets. The overall efficiency is 98.48% and the overall uncertainty (statistical and systematic, dominated by the latter) is 0.7%. The calculation of the PS systematic itself assumes the bias is one-sided, i.e. it is strictly the number obtained when calculating the difference between the efficiency at the DEP and that around 2039 keV, usually positive. However, unlike the original version of DCR, when the previous iteration of the numbers (using new DCR, in 2021) was calculated, we saw a few instances where the PS bias had the opposite sign. As such, it was decided we should report the bias as double sided in all cases. The fact that the PS bias is smaller than in the original version of DCR is probably due to the fact that new DCR is calculated such that the distribution and acceptance are relatively flat with energy, at least between 1000-2615 keV. The acceptance is not perfectly flat, but it helps explain why the PS bias values are smaller. This also explains why in a small percentage of cases, the PS bias might have a different sign than observed in the previous implementation of DCR. The improvement in pile-up removal is also thought to have helped. Applying the low AvsE cut before computing the final dcr_{corr} means there are fewer multi-site events which get through the cut than before, also potentially reducing the PS bias.

Data Set	$\epsilon_{2028-2050}$ (%)	ϵ_{DEP} (%)	σ_{PS} (%)	$\sigma_{DB_{stab}}$ (%)	$\sigma_{cal_{stab}}$ (%)	σ_{stat} (%)	σ_{tot} (%)
DS0	98.46	98.79	±0.33	±0.12	±0.23	±0.06	±0.43
DS1	98.07	98.56	±0.50	±0.14	±0.24	±0.06	±0.58
DS2	98.15	99.33	±1.18	±0.06	±0.34	±0.12	±1.23
DS3	98.21	98.83	±0.63	±0.14	±0.30	±0.09	±0.72
DS4	98.46	98.30	±0.16	±0.15	±0.33	±0.11	±0.41
DS5a	98.26	98.71	±0.45	±0.42	±0.36	±0.13	±0.72
DS5b	97.99	98.64	±0.65	±0.19	±0.34	±0.11	±0.76
DS5c	98.27	98.49	±0.22	±0.17	±0.03	±0.11	±0.30
DS6a	98.56	99.24	±0.68	±0.09	±0.04	±0.03	±0.69
DS6b	98.64	99.23	±0.59	±0.33	±0.19	±0.04	±0.71
DS6c	98.53	99.12	±0.58	±0.14	±0.01	±0.04	±0.60
DS7	98.57	99.33	±0.76	±0.21	±0.27	±0.07	±0.83
DS8 PPC	98.79	99.35	±0.57	±0.12	±0.23	±0.05	±0.63
DS8 ICPC	97.85	98.80	±0.95	±0.23	±0.36	±0.13	±1.05

Table 4.2: $0\nu\beta\beta$ efficiency and uncertainties, for dcr_{corr} (i.e. with the drift time correction included) given for high gain channels in each data set, including both enriched and natural detectors.

Data Set	$\epsilon_{2028-2050}$ (%)	ϵ_{DEP} (%)	σ_{PS} (%)	$\sigma_{DB_{stab}}$ (%)	$\sigma_{cal_{stab}}$ (%)	σ_{stat} (%)	σ_{tot} (%)
DS0	98.51	98.79	±0.28	±0.17	±0.02	±0.06	±0.34
DS1	98.02	98.50	±0.49	±0.14	±0.24	±0.06	±0.57
DS2	98.05	99.27	±1.22	±0.07	±0.35	±0.13	±1.28
DS3	98.28	98.74	±0.46	±0.14	±0.31	±0.10	±0.58
DS4	98.54	97.97	±0.57	±0.14	±0.38	±0.14	±0.71
DS5a	98.26	98.61	±0.34	±0.58	±0.40	±0.16	±0.80
DS5b	98.21	98.73	±0.53	±0.21	±0.37	±0.14	±0.69
DS5c	98.22	98.55	±0.33	±0.21	±0.37	±0.14	±0.56
DS6a	98.50	99.18	±0.68	±0.12	±0.04	±0.03	±0.69
DS6b	98.57	99.21	±0.64	±0.39	±0.22	±0.05	±0.78
DS6c	98.43	99.02	±0.59	±0.14	±0.21	±0.05	±0.64
DS7	98.51	99.30	±0.79	±0.22	±0.30	±0.09	±0.88
DS8 PPC	98.90	99.36	±0.47	±0.15	±0.24	±0.06	±0.55
DS8 ICPC	97.85	98.80	±0.95	±0.23	±0.36	±0.13	±1.05

Table 4.3: DCR efficiency and uncertainties, for dcr_{corr} (i.e. with the drift time correction included) given for high gain channels of enriched detectors in each data set.

Data Set	$\epsilon_{2028-2050}$ (%)	ϵ_{DEP} (%)	σ_{PS} (%)	$\sigma_{DB_{stab}}$ (%)	$\sigma_{cal_{stab}}$ (%)	σ_{stat} (%)	σ_{tot} (%)
DS0	98.35	98.91	±0.46	±0.05	±0.07	±0.12	±0.49
DS1	98.54	99.15	±0.61	±0.14	±0.43	±0.19	±0.78
DS2	99.14	99.90	±0.76	±0.14	±0.52	±0.27	±0.97
DS3	97.91	99.27	±1.36	±0.18	±0.48	±0.23	±1.47
DS4	98.38	98.65	±0.26	±0.18	±0.40	±0.16	±0.54
DS5a	98.25	98.88	±0.63	±0.19	±0.45	±0.21	±0.82
DS5b	97.63	98.46	±0.85	±0.19	±0.44	±0.19	±1.00
DS5c	98.34	98.37	±0.02	±0.11	±0.43	±0.18	±0.48
DS6a	98.69	99.37	±0.67	±0.06	±0.22	±0.05	±0.71
DS6b	98.80	99.27	±0.47	±0.24	±0.26	±0.07	±0.59
DS6c	98.75	99.33	±0.58	±0.15	±0.18	±0.06	±0.63
DS7	98.78	99.43	±0.65	±0.21	±0.36	±0.13	±0.78
DS8	98.59	99.32	±0.73	±0.09	±0.32	±0.10	±0.81

Table 4.4: $0\nu\beta\beta$ efficiency and uncertainties, for dcr_{corr} (i.e. with the drift time correction included) given for high gain channels of natural detectors in each data set.

4.9 Alpha Distribution Study

In this section, the main goal is to try to understand where the α particles occur versus energy, as well as what distribution they have in DCR. This is not an alpha calibration study. Rather, it is a study that endeavors to use the DCR parameter, as well as LQ and high AvsE, to determine the properties of the alpha distribution in physics data. We also investigate the parameters' performance in discriminating alpha particles.

It is important to see the effect of applying DCR on the MAJORANA DEMONSTRATOR spectrum, and the effect of using DCR in combination with similar surface cuts. It is also important to check that these parameters are working as expected when applied to the MAJORANA DEMONSTRATOR data. Furthermore, it is also crucial to explore the shape of the alpha spectrum and its distributions in the various analysis cuts in order to understand the residual background after all cuts and to build strategies for improved cuts in the future. Below, in Fig. 4.14 through 4.32, is shown the effect of applying DCR as well as other pulse shape discrimination cuts, LQ and AvsE. First, a cut was applied in the full spectrum for each of DCR and LQ, in the energy range of 1950-2350 keV. The number of events around the peak of the distributions was counted, in both physics data and calibrations, in the range from -1 to 1 units. This allowed the two sets of distributions to be compared, by normalizing the number of events in the peak to the same value. The results are shown in Fig. 4.14, Fig. 4.15 and Fig. 4.17. This most directly shows the different populations with high DCR, high LQ and high AvsE respectively.

In the first figure, Fig. 4.14, one can compare DCR in physics data vs DCR in calibration skims, in the energy region between 1950-2350 keV. The distribution shows that, above a few units of DCR, the distribution in the physics data is strongly dominated by events not present, or rarely present, in the calibration, which are expected to be alpha particles on the surface of our detector. We can also conclude that DCR distinguishes alphas and bulk events well in this energy range. For a typical alpha DCR of about 20, there are on the order of a few thousand more alpha events than bulk events, as can be seen from the plot – we expect

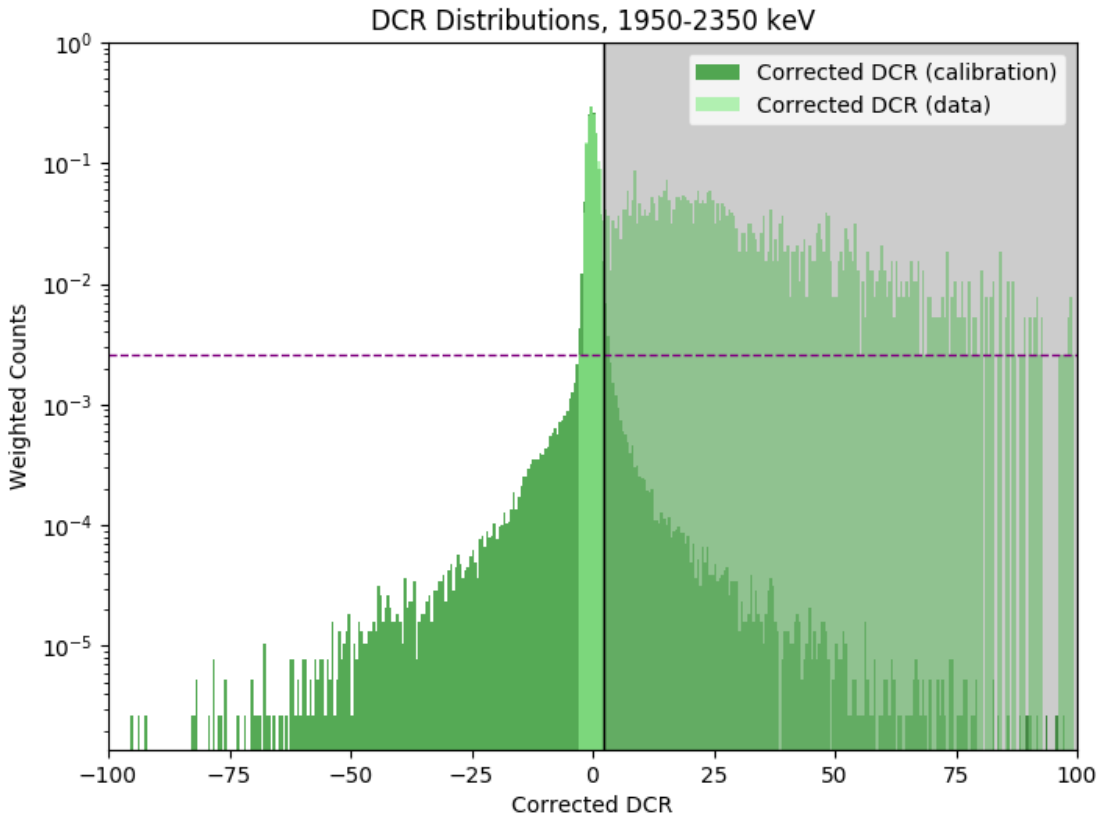

Figure 4.14: The corrected DCR distribution, for the full dataset (light green) and for one long calibration (dark green), overlaid, in the region between 1950-2350 keV. The shaded grey indicates the region where the DCR parameter cuts events. The dashed purple line represents one event in the data – differences below this line should be ignored due to statistical effects.

O(1000) more alphas than bulk events to be tagged by DCR. This indicates that DCR is working as expected, i.e. there are plenty of expected alpha events in full $0\nu\beta\beta$ dataset, which, in comparison, have a much higher DCR than the signal-like events in calibrations.

The second plot, Fig. 4.15, shows the performance of the LQ discriminator in the same energy range as above, 1950-2350 keV. This spectrum shows that LQ is also effective at discriminating non-bulk events. Alphas are also expected to have high LQ alongside high DCR. As such, there seems to be a synergy between DCR and LQ, as will be explored in the rest of this subsection.

In Fig. 4.16, we see the differences in DCR and LQ distributions between data and ^{228}Th calibration. After comparing the normalized distributions, they can also be subtracted. This allows a look at the residual of these distributions, i.e. in effect, a comparison of where there are more events in the data, which correspond to alphas and possibly other surface events, or more events in the calibrations, which can simply be an effect of statistics. Note that the alpha distribution in Fig. 4.14 curves down as it meets the DCR=0 peak. This indicates that most alpha events have DCR well above 0, but since the distributions meet before the stats get particularly low, we can expect that some alphas are passing the cut. We accept those alphas because lowering the cut further rapidly degrades efficiency for unknown additional background removal. Due to the region in which the calibrations are normalized, the sums of events in the residual distributions' bins do not necessarily add to zero. As DCR is a surface alpha cut, the high-DCR residual distribution indicates what DCR distribution alpha particles have. Specifically, it peaks around 20 DCR units, and extends down to a couple of units of DCR. The LQ distribution is broader, somewhat less peaked, but has similar features.

Although the AvsE cuts are not the mainstay of this thesis, a look at the AvsE spectra of $0\nu\beta\beta$ and calibration data in the region between 1950-2350 keV is provided. We see relatively fewer high AvsE events in calibration than in physics data, indicating some high AvsE events which were not present in calibrations. Indeed, the high AvsE cut previously discussed, which tags mainly events close to the point contacts of the HPGe detectors, was implemented as

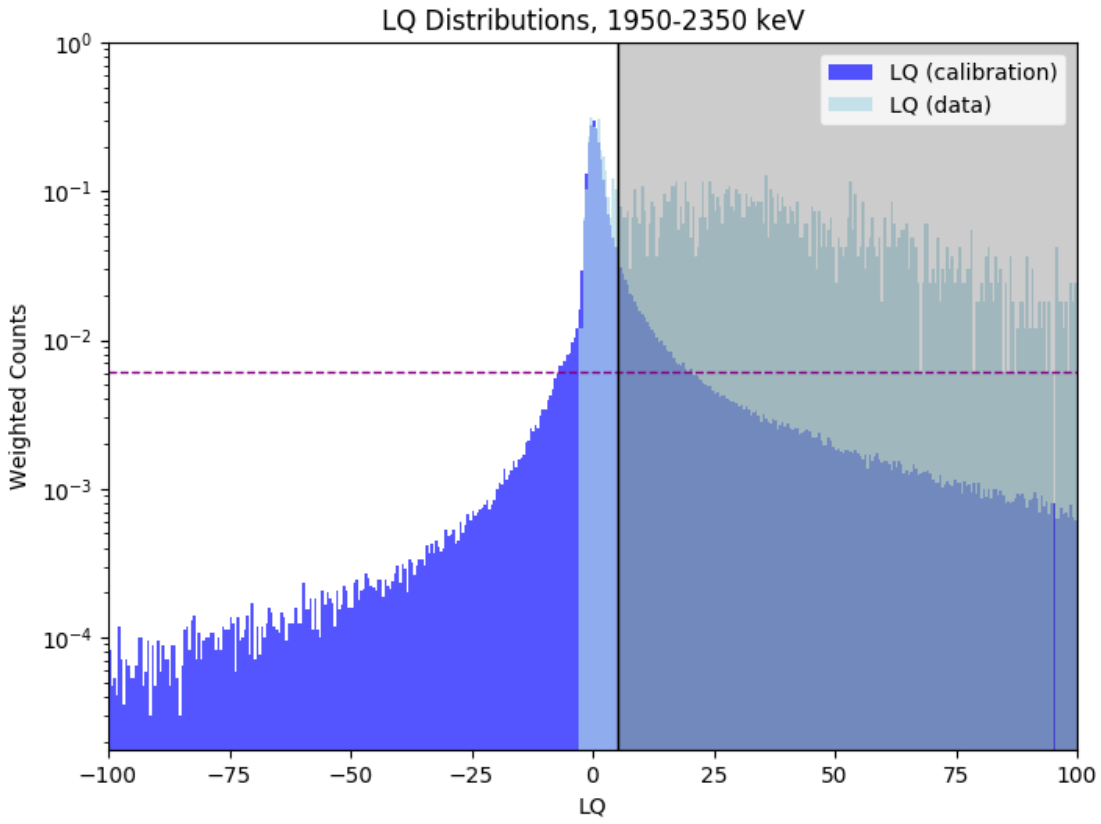

Figure 4.15: The LQ distribution, for the full dataset (light blue) and for one long calibration (dark blue), overlaid, in the region between 1950-2350 keV. The shaded grey indicates the region where the LQ parameter cuts events. The dashed purple line represents one event in the data – differences below this line should be ignored due to statistical effects.

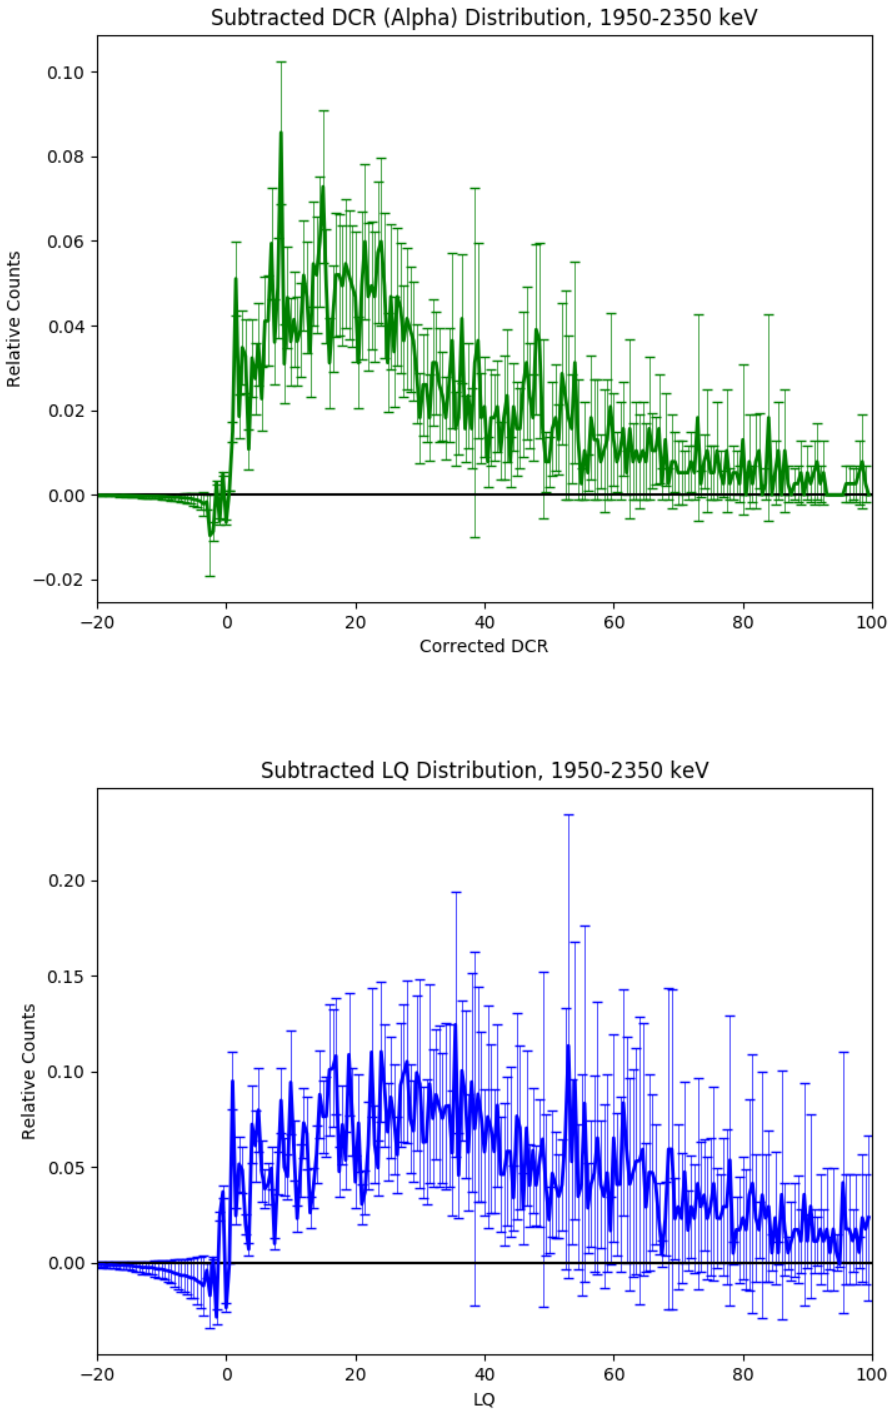

Figure 4.16: Estimates of the DCR and LQ distributions of surface alpha events in MAJORANA DEMONSTRATOR, i.e. subtracting the long calibration distributions from the full physics dataset distributions, in the region between 1950-2350 keV. There is a black horizontal line at 0 to guide the eye.

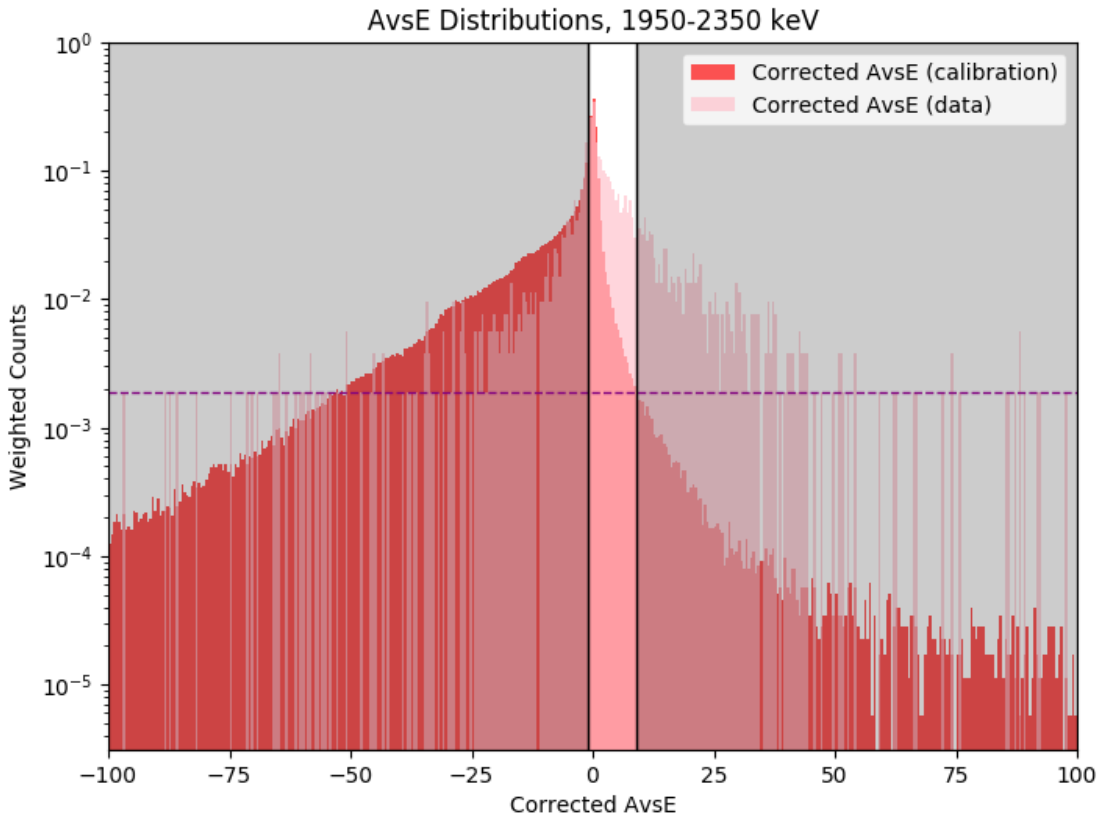

Figure 4.17: The corrected AvsE distribution, for the full physics dataset and for one long calibration, overlaid, in the region between 1950-2350 keV. The shaded grey indicates the regions where the AvsE parameter cuts events. The dashed purple line represents one event in the data – differences below this line should be ignored due to statistical effects.

an aide to surface alpha discrimination,, which this plot further justifies. According to Nick Ruof's studies [1] [109], while AvsE is a reliable cut for bulk events, the short drift time and non-linear weighting potential around the point contact lead to events with one site near the point contact. This type of multi-site events often passes the AvsE cut. A high AvsE cut would remove these events as well as degraded alphas near the point contact which escape the DCR cut. However, it would also effectively remove all events which deposit energy near the point contact, acting as a volumetric cut. For Ruof's thesis [109], an optimization was done to determine where to set the high AvsE cut. In ^{208}Tl calibration background events, the cut only decreases the figure of merit (FOM), i.e. the maximum FOM occurs with no high AvsE cut at all. The FOM scales with efficiency over the uncertainty on the signal as a function of the background. For high background the latter would scale with the square root of the background, but we are at such low background that we need to use a version that accounts for the non-Gaussian nature of the Poisson distribution at low stats. More details can be found in Ruof's thesis [109]. However, high AvsE removes αs on the detector surface near the point contact, which are not removed by DCR. α events do not occur in calibration data with any appreciable frequency, if at all. Selecting a high AvsE < 9, circa 7% of events in the background estimation window in open data (not in ^{208}Tl calibration!) are removed, while decreasing the overall signal acceptance in calibration data by roughly 2%. In the full dataset, some events are removed only by the high AvsE cut. The discovery sensitivity is thus increased, and the number of counts in the background estimation window is reduced. Fig. 4.20 shows events which are not cut by DCR, but are cut by AvsE (note: some of these events are also cut by LQ). The main idea was to use open data, to make sure events with values of AvsE and DCR values, which rarely appear in calibration but which appear in open data, can be removed. According to Fig. 4.17, the amount of events cut by high AvsE is confirmed. There is an open question whether the AvsE < 9 cut is optimal. Looking at the respective slopes of the bin counts, we can conclude that going out to a higher cut would cut fewer events in background but could also miss a fair amount of potential alphas. The idea is to try to tag events which rarely appear in calibration but appear in data. There are several

shoulders in the pink data AvsE distribution, with the largest one corresponding to a value of AvsE of roughly 8-12; these are exactly the events we would like to be able to cut for the best performance. As such, the high AvsE cut appears close to optimal, although a slightly lower cut could be investigated, but any such investigation would probably be marred by statistical fluctuations in our data,, due to our very low backgrounds.

In Fig. 4.18, one can compare DCR in physics data in the energy region between 1500-1800 keV to DCR in calibration skims, around the Double Escape Peak (DEP). Again, this figure shows that, above a few units of DCR, the distribution in the physics data is strongly dominated by events not present, or rarely present, in the calibration data, which are expected to be alpha particles on the surface of our detector. We can also conclude that DCR distinguishes alphas and bulk events well, also in this energy range. For a typical alpha DCR of about 20, there are on the order of a few hundred more alpha events than bulk events, as can be seen from the plot – we expect O(100) more alphas than bulk events to be tagged by DCR. Again, this indicates that DCR is working as expected in this lower energy range, i.e. there are plenty of expected alpha events in full $0\nu\beta\beta$ dataset around 1500-1800 keV, which, in comparison, have a much higher DCR than the signal-like events in calibrations. However, the discriminatory power appears to be higher around the ROI. Similar conclusions can be drawn for LQ as well, looking at Fig. 4.19. The residual plots can be looked at in the appendix, Fig. B.2, and are generally similar to before.

Fig. 4.20 through Fig. 4.26 are the N-2 plots of the DCR, LQ, and AvsE parameters, i.e. plots showing how the distributions of the parameters relate to one another, both for an energy range of 1950-2350 keV and for a larger energy swath (E $>$ 100 keV). In Fig. 4.20, the 2-D distribution of DCR vs AvsE in the ROI range, we note that the most events cut by DCR have relatively high values of AvsE; this supports the implementation of a high AvsE cut, and shows that this cut further restricts the phase-space in order to better reject background. Interestingly, there are some events in the calibration, which don't appear at all in data, with relatively high AvsE and very low DCR, i.e. in the lower right quadrant of the calibration plot. A scatter plot, Fig. 4.21, shows that these events are correlated

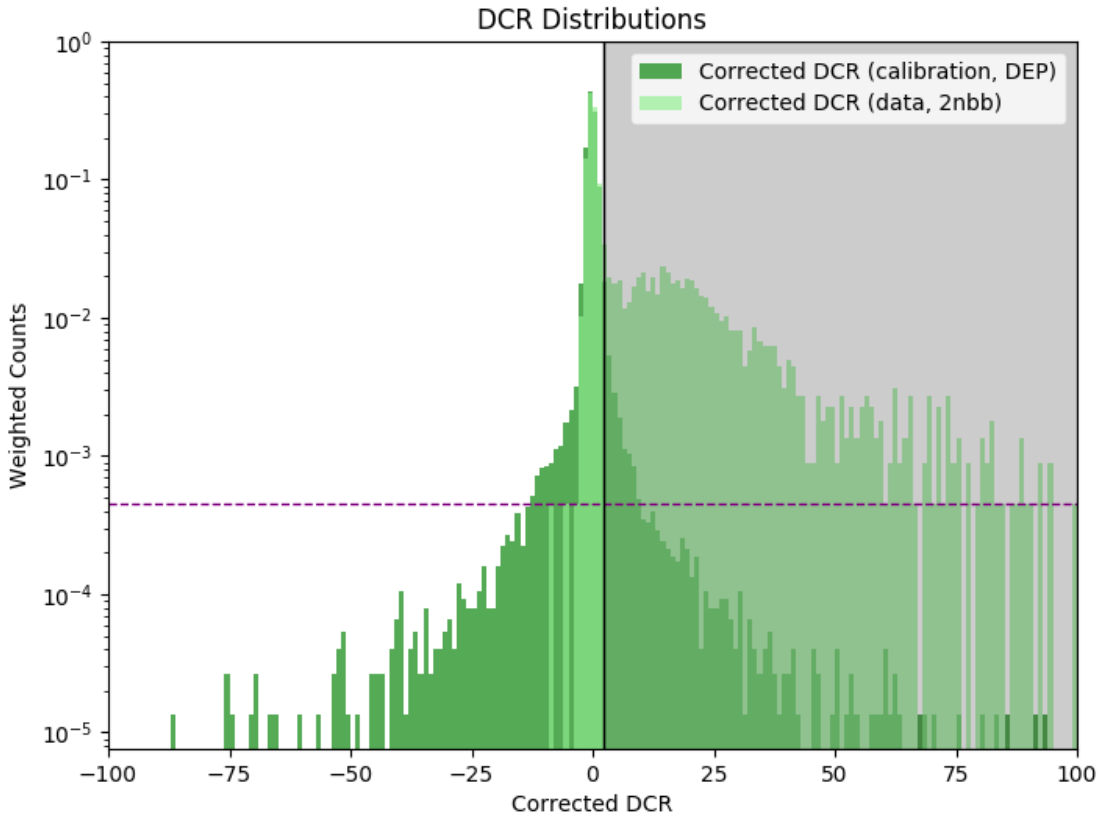

Figure 4.18: The corrected DCR distribution, for the full physics dataset and for one long calibration, overlaid, in the region between 1500-1800 keV and around the DEP respectively. The shaded grey indicates the region where the DCR parameter cuts events. The dashed purple line represents one event in the data – differences below this line should be ignored due to statistical effects.

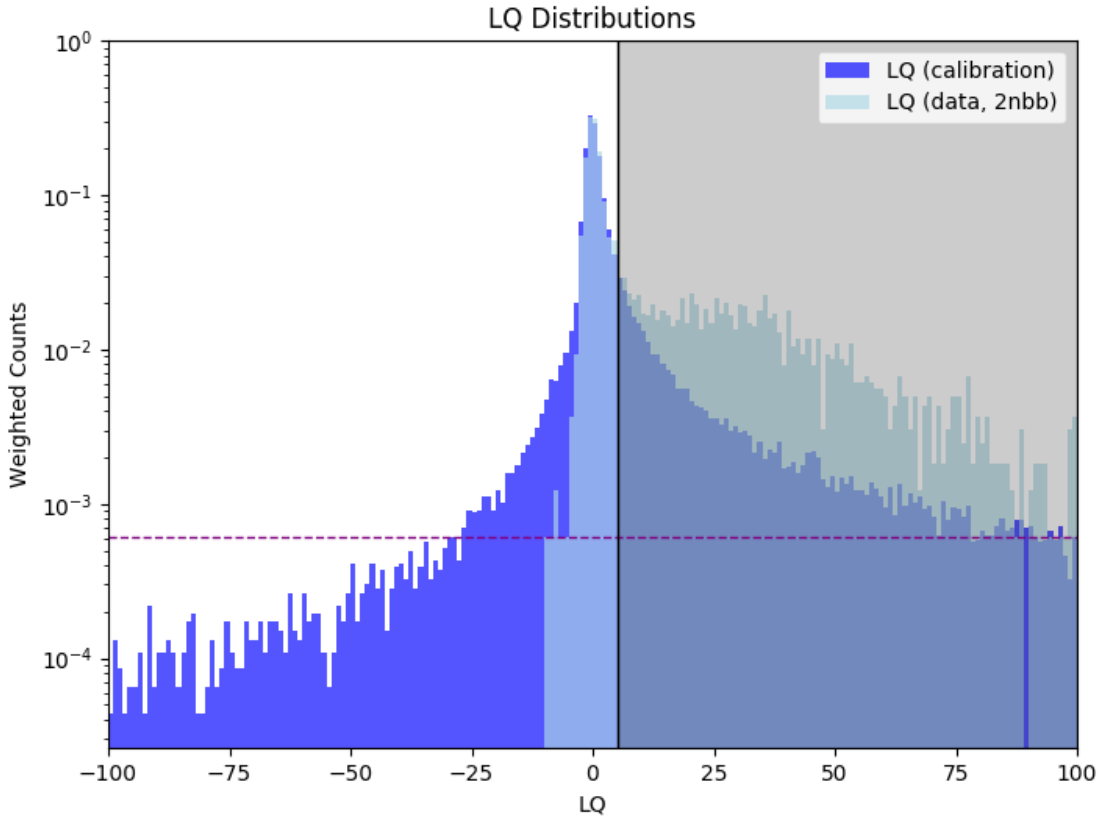

Figure 4.19: The LQ distribution, for the full physics dataset and for one long calibration, overlaid, in the region between 1500-1800 keV. The shaded grey indicates the region where the LQ parameter cuts events. The dashed purple line represents one event in the data – differences below this line should be ignored due to statistical effects.

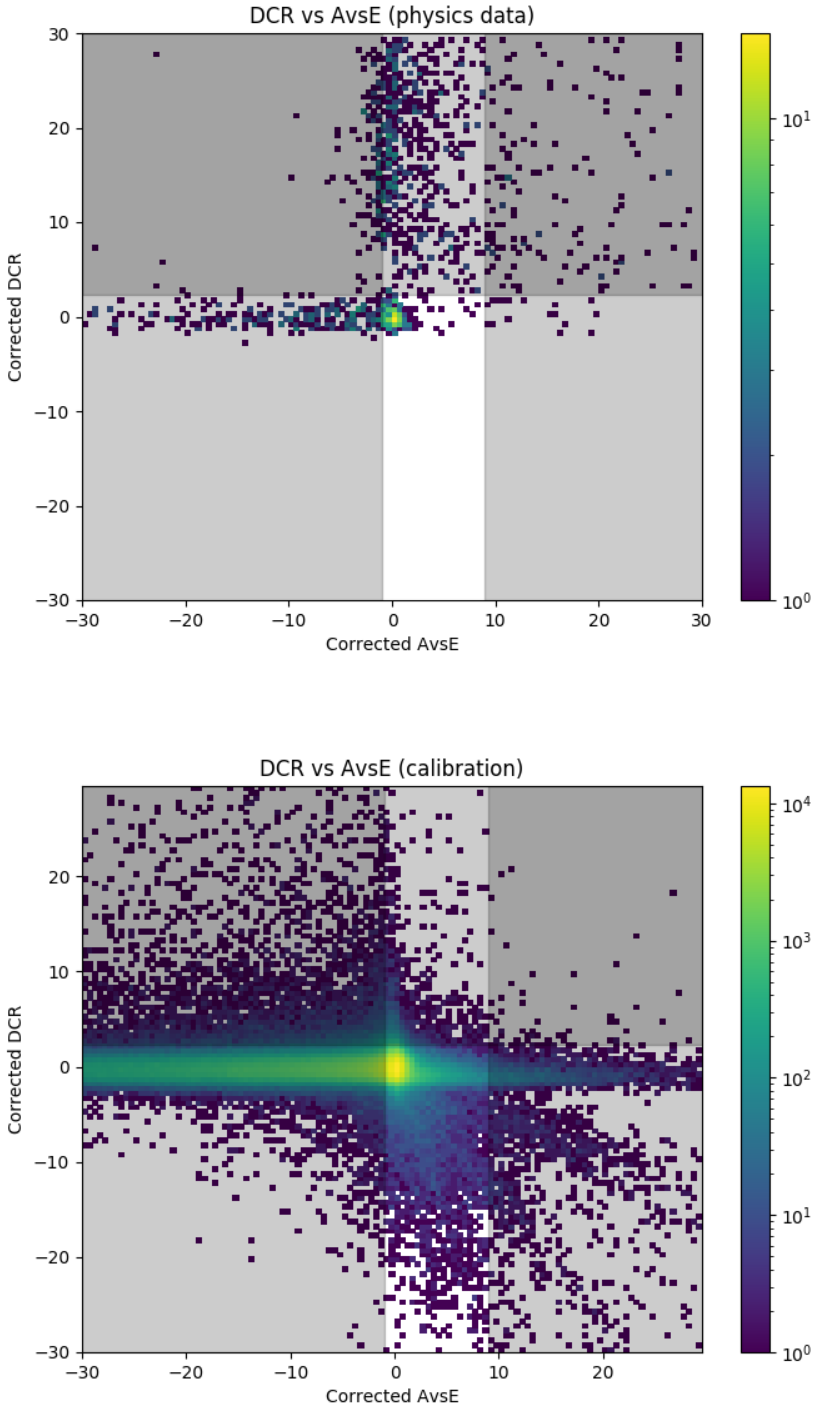

Figure 4.20: N-2 plots, i.e. Corrected DCR vs corrected AvsE, for the full physics dataset and for one long calibration, in the region between 1950-2350 keV. The shaded grey indicates the regions where the DCR and AvsE parameters cut events.

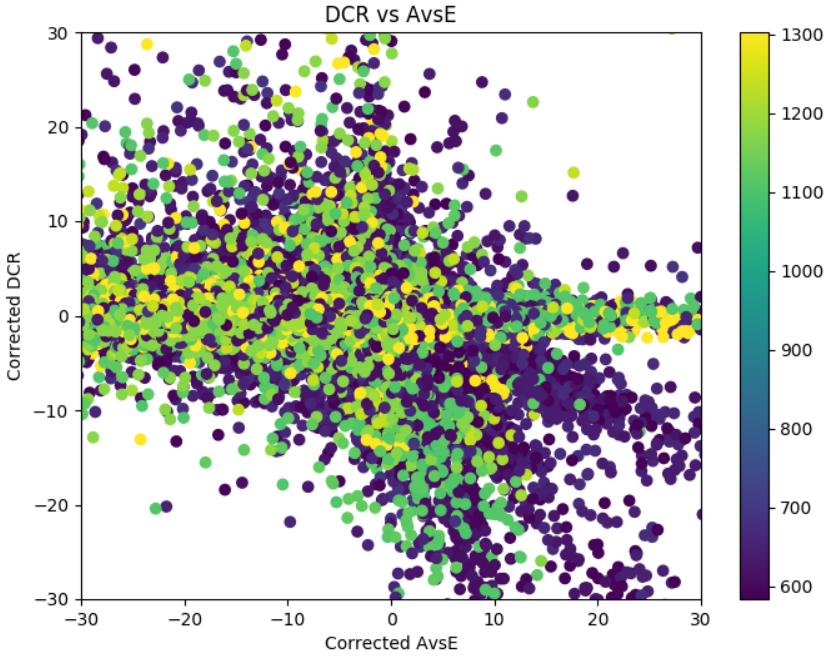

Figure 4.21: N-2 plot, i.e. Corrected DCR vs corrected AvsE, for one long calibration: a scatter plot vs channel on the z-axis, in the region between 1950-2350 keV. We note the diagonals are associated with a given subset of detectors.

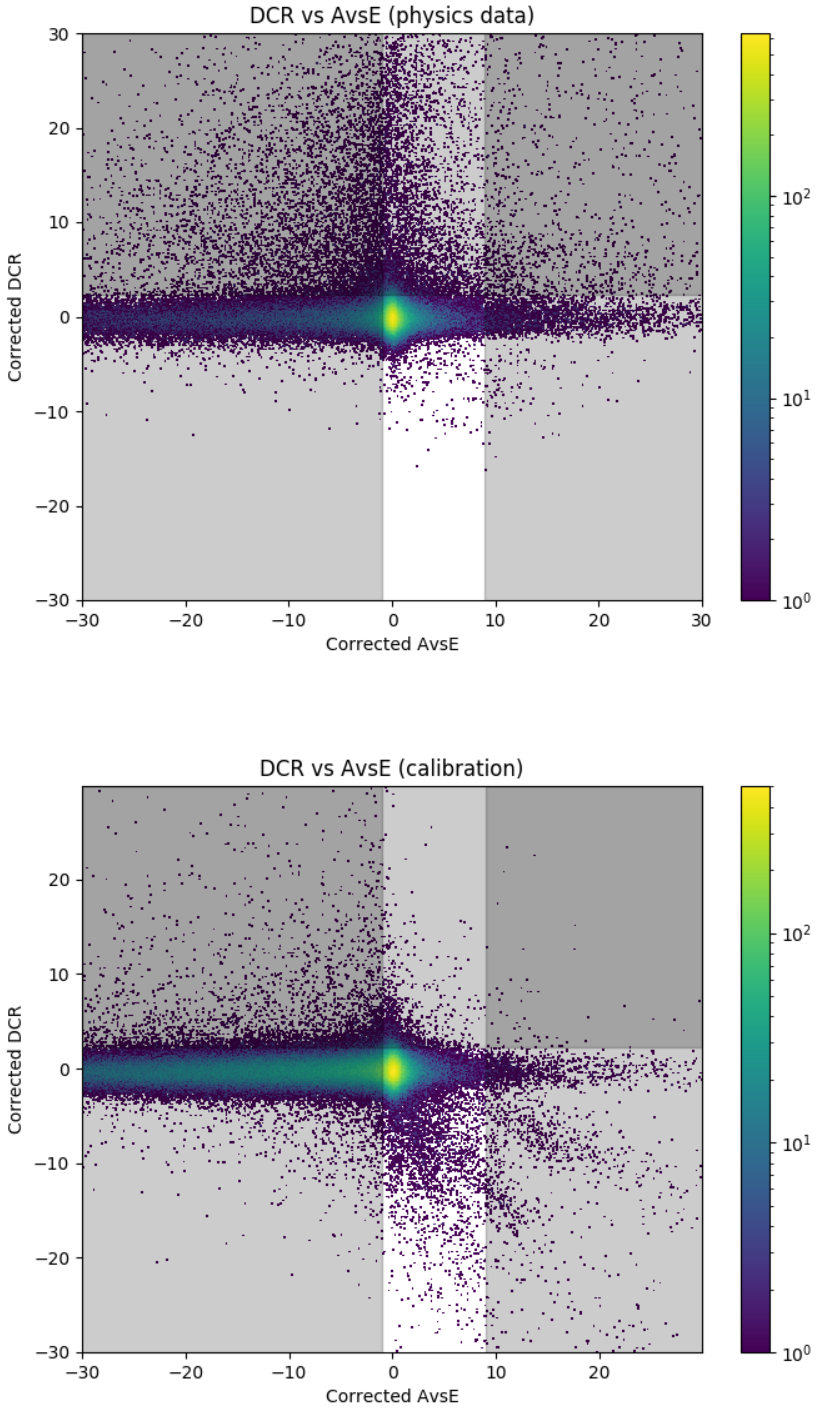

Figure 4.22: N-2 plots, i.e. Corrected DCR vs corrected AvsE, for the full physics dataset and for one long calibration, with energy greater than 100 keV. The shaded grey indicates the region where the DCR and AvsE parameters cut events.

with mostly Module 1 detectors and a couple of Module 2 detectors. While interesting, the appearance of these events is at the sub-percent level. In the full energy plot, Fig. 4.22, the 2-D distribution of DCR vs AvsE in the full energy range, these events are still visible but appear to make up even less of the distribution, indicating some energy dependence.

A similar feature to that in Fig. 4.20 appears in Fig. 4.23, the 2-D distribution of DCR vs LQ in the ROI range. There are a couple of diagonal lines with few events with a higher than expected DCR for a given LQ. These are strongly visible around 1950-2350 keV, but are barely visible in Fig. 4.24, the 2-D distribution of DCR vs LQ drawn for the full energy range, indicating some energy dependence. Fig. 4.26 shows the 2-D distribution of LQ vs AvsE for the full energy. The most interesting feature is the slope which correlates low AvsE with slightly higher LQ. Fig. 4.25, the 2-D distribution of LQ vs AvsE in the ROI range, is presented for reference, but does not show anything other than its full energy counterpart.

In Fig. 4.27, we plot the DCR vs Energy distribution. This is arguably the most interesting plot for DCR, as it shows how the DCR distribution varies with energy. The flat line extending up to roughly 2600 keV in calibration indicates DCR's performance upon tuning. Seeing a similar distribution in physics data adds to the conclusion that DCR performs well across energy. Just above 5000 keV, we notice what appears to be the ^{210}Po alpha peak, around 5400 keV. Towards the upper left, we see alphas, more and more energy-degraded as DCR increases, i.e. these events appear to have more and more trapped and re-released charge, and therefore higher DCR, as we decrease in energy. This feature indicates DCR works as a delayed charge recovery cut, and works as expected, with more and more charge recovered with energies further from the aforementioned peak. It can be seen, however, that DCR does not succeed in tagging all the alpha events with energies near 5400 keV. The AvsE vs Energy plot, Fig. 4.28 shows that all calibration events are below a certain energy around 3500 keV whereas the physics distribution extends further out in energy. Some of these high-E events are cut by both high and low AvsE cuts. The LQ vs Energy plot, Fig. 4.29 shows LQ can tag these events so there is a measure of complementarity between these cuts. LQ, as seen in the 1D distributions as well, has a longer high-LQ tail than DCR's high-DCR

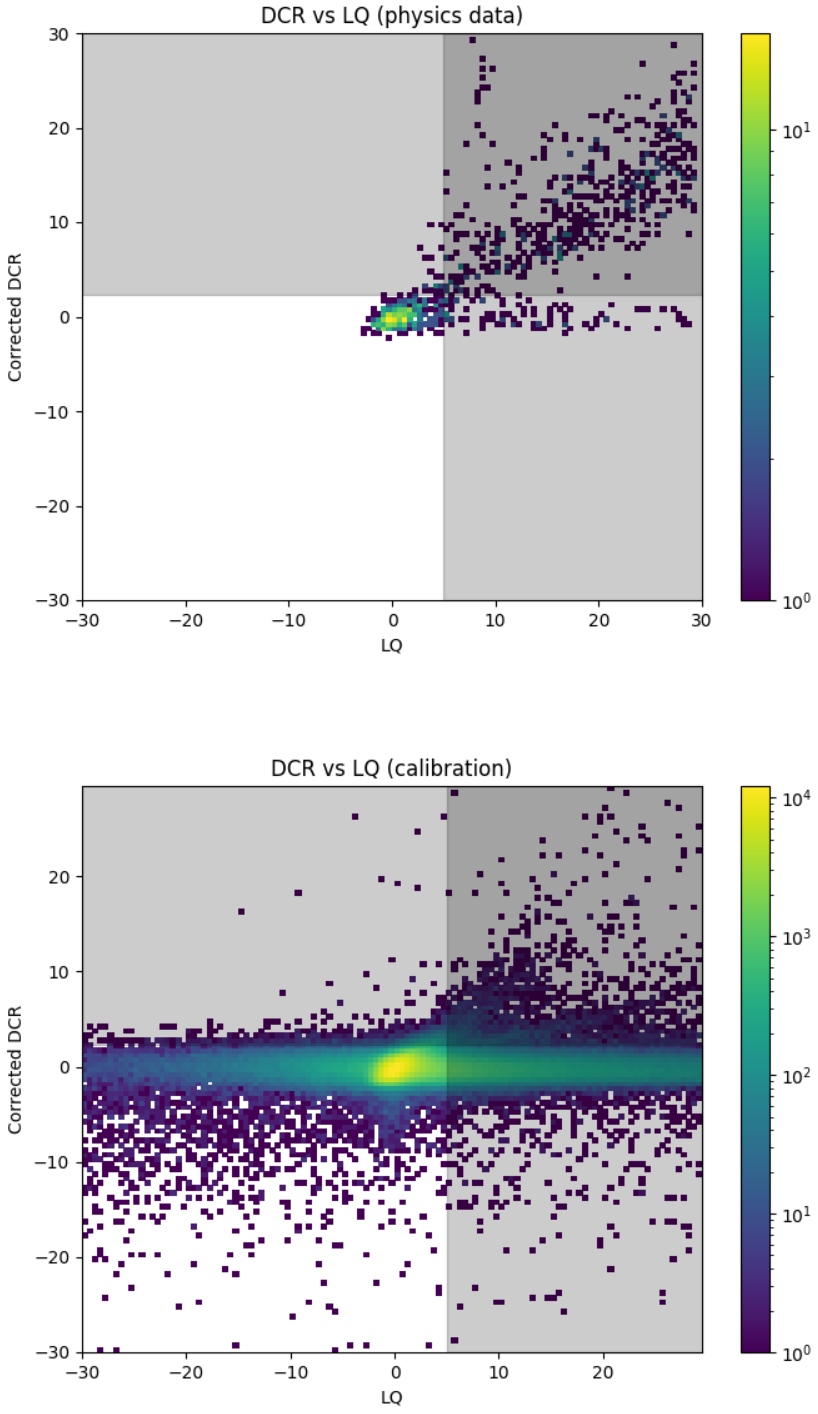

Figure 4.23: N-2 plots, i.e. Corrected DCR vs LQ, for the full physics dataset and for one long calibration, in the region between 1950-2350 keV. The shaded grey indicates the region where the DCR and LQ parameters cut events.

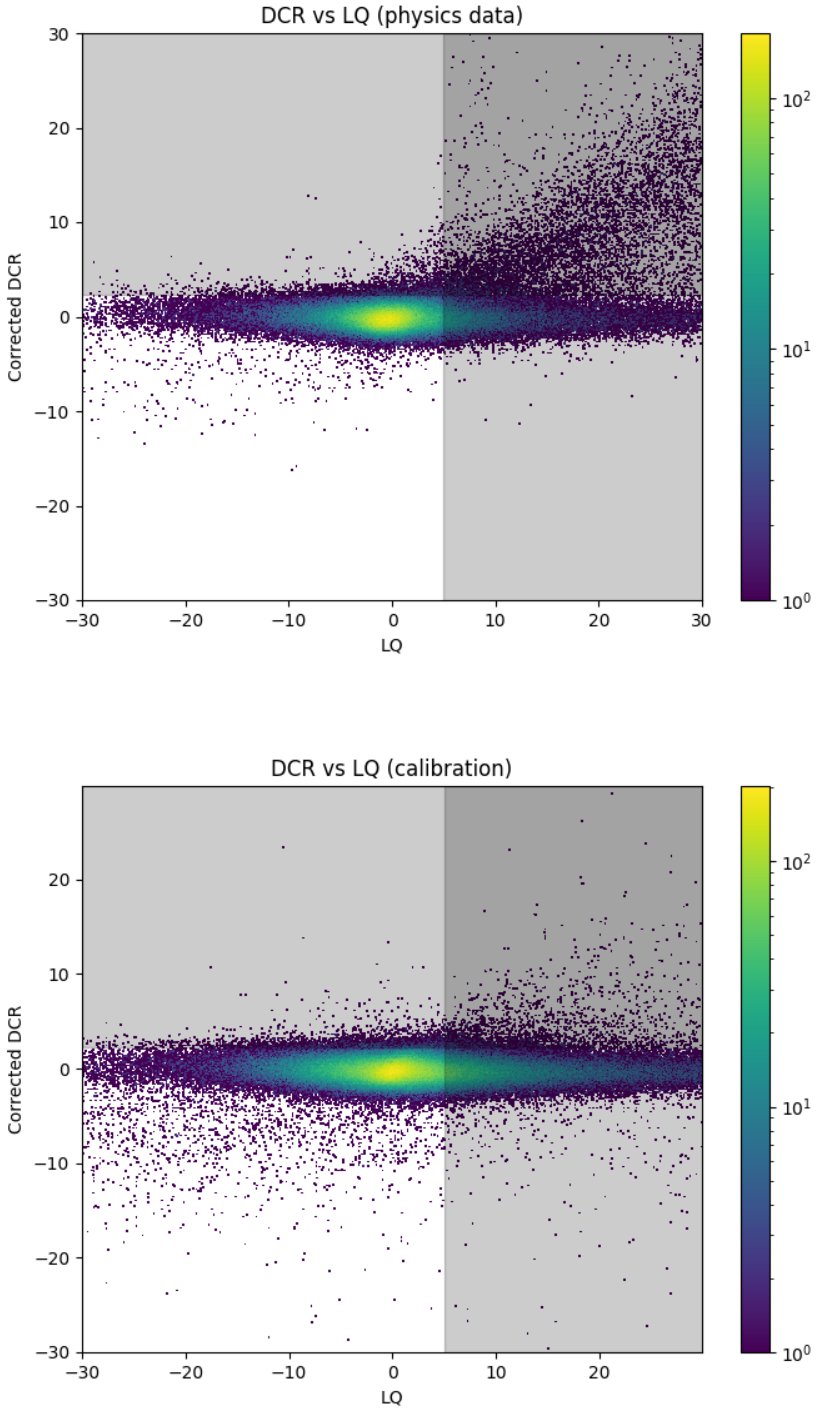

Figure 4.24: N-2 plots, i.e. Corrected DCR vs LQ, for the full physics dataset and for one long calibration, with energy greater than 100 keV. The shaded grey indicates the region where the DCR and LQ parameters cut events.

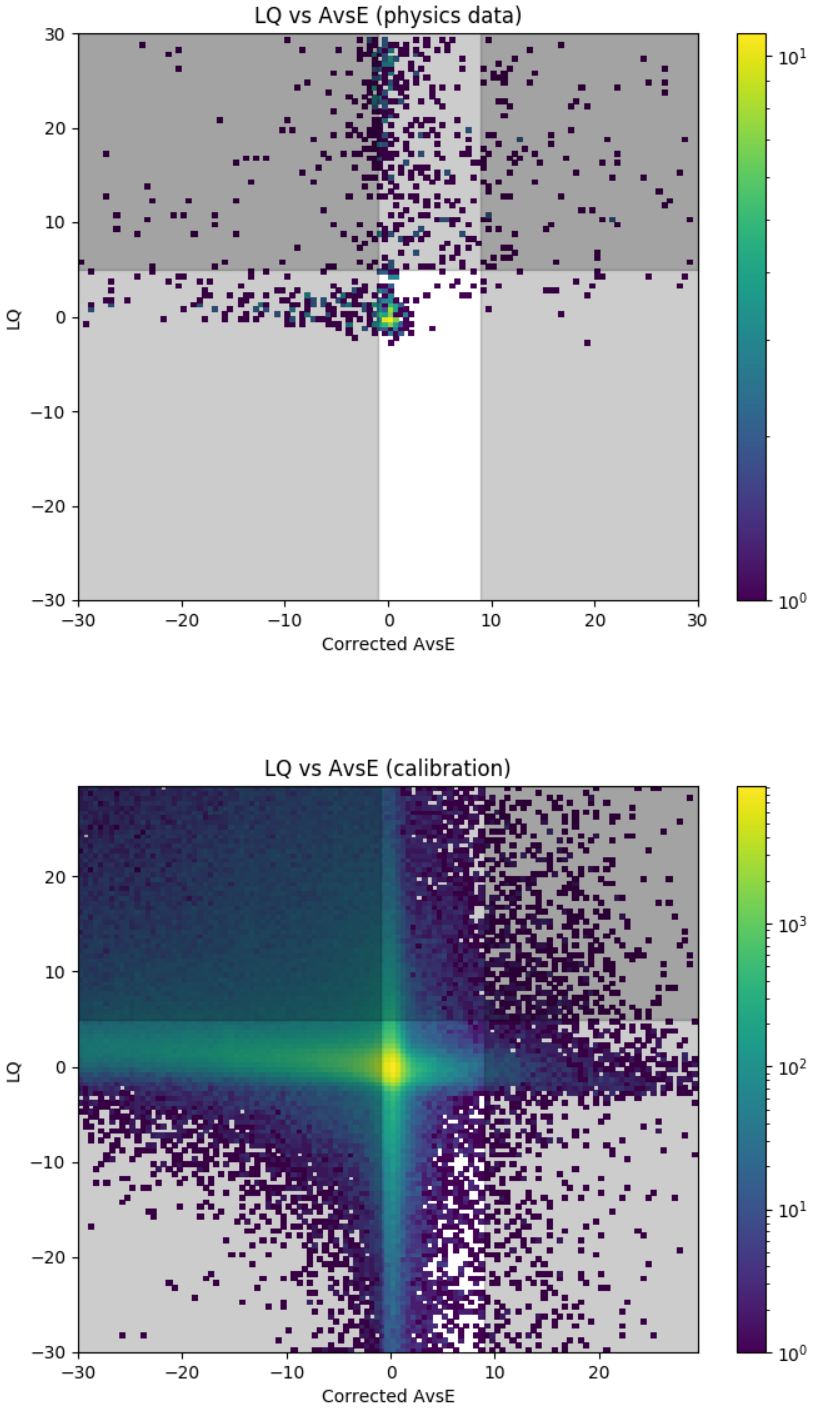

Figure 4.25: N-2 plots, i.e. LQ vs corrected AvsE, for the full physics dataset and for one long calibration, in the region between 1950-2350 keV. The shaded grey indicates the region where the AvsE and LQ parameters cut events.

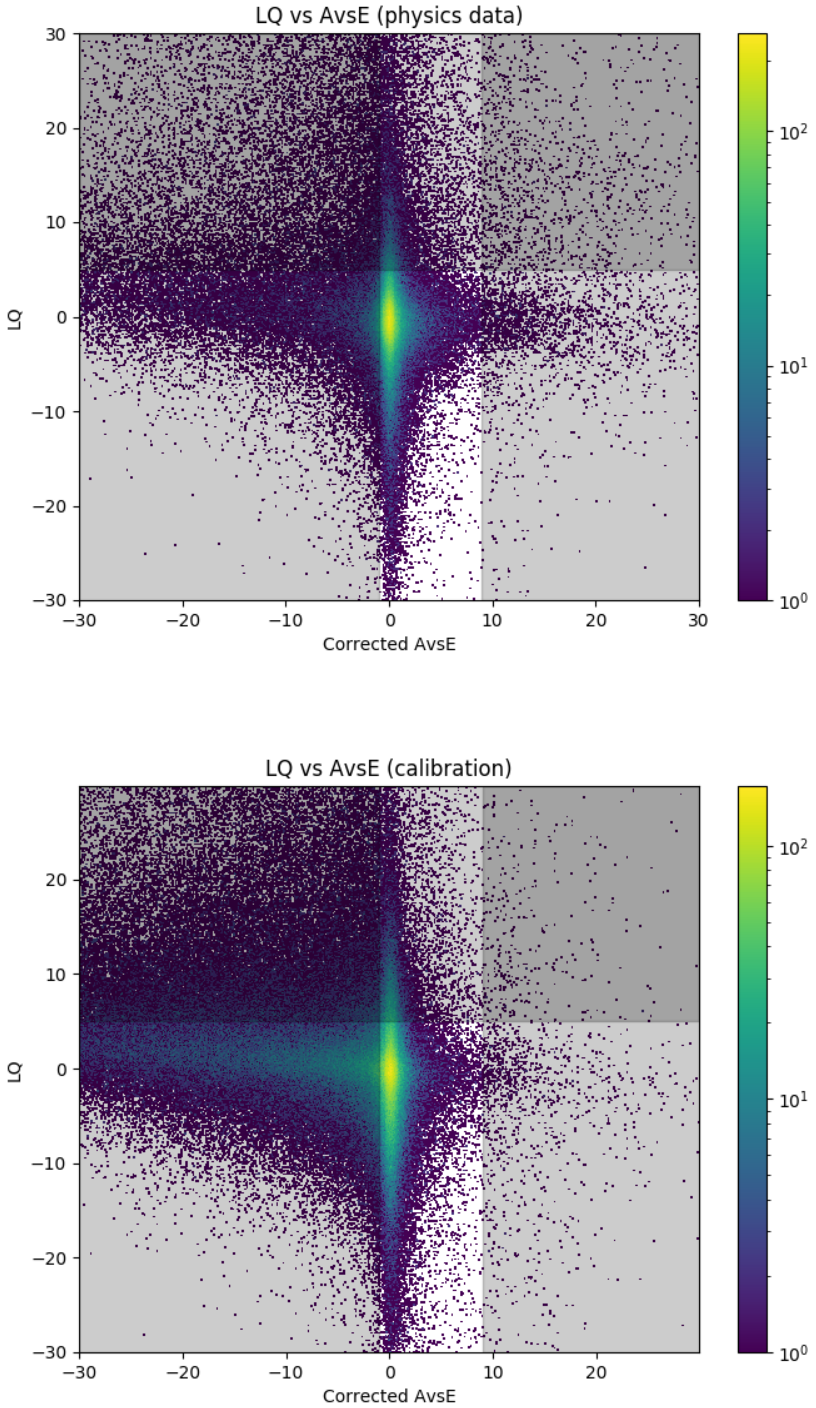

Figure 4.26: N-2 plots, i.e. LQ vs corrected AvsE, for the full physics dataset and for one long calibration, with energy greater than 100 keV. The shaded grey indicates the region where the AvsE and LQ parameters cut events.

tail.

DCR's main use in MAJORANA DEMONSTRATOR is rejecting surface alpha particles. Therefore, an important spectrum to show is that which shows the events which fail the DCR cut, i.e. supposedly alpha events. This can be seen in Fig. 4.30. Also shown are the equivalent plots for the LQ and AvsE parameters, 4.31 and 4.32.

In Fig. 4.30, we note that the anti-DCR cut spectrum extends up to roughly 5 MeV, i.e. to close but not at the alpha energy of the ^{210}Po decay to ^{206}Pb. This is, again, consistent with the expectation that the DCR parameter tags alphas with a degraded energy. The left-hand side plot of physics data tells us an interesting story. Going from the anti-DCR (dark blue) cut only to adding the anti-LQ cut (purple), we lose virtually no events above 2000-4800 keV, i.e. most events tagged by DCR are also tagged by LQ. The relative performance of LQ is better at higher energy, and much better above 4800 keV. DCR then appears to do better towards lower energy. This remains consistent when comparing the two combinations in addition to the anti-high-AvsE cut as well. The latter cut completely removes the very high energy events and shifts the higher energy peak to around 2500 keV, as opposed to 2000 keV. Therefore, high AvsE appears quite good in removing events around 2000 keV. The alpha, or at least high-DCR, high-LQ, high-AvsE peak extends down to about 1000 keV, peaking at slightly above 2000 keV. Interestingly, the alphas are right in the middle of the ROI. This speaks to how crucial a surface alpha cut, or DCR, is to the functioning of the MAJORANA DEMONSTRATOR.

Overall, the story remains similar when looking at Fig. 4.31. LQ and DCR have very similar performance at $E > 2000$ keV. There are quite a lot of lower energy events (100-300 keV) not tagged by LQ but tagged by DCR. As already stated, the AvsE cut completely removes the very high energy events and shifts the mode of the distribution to around 2500 keV, as opposed to 2000 keV.

Looking at Fig. 4.32, we note that once anti-high-AvsE is applied, there is not much left to cut with DCR and LQ above roughly 1500 keV. Applying anti-LQ removes some lower E events, and applying DCR removes more; at medium energies (1300-1600 keV), perhaps LQ

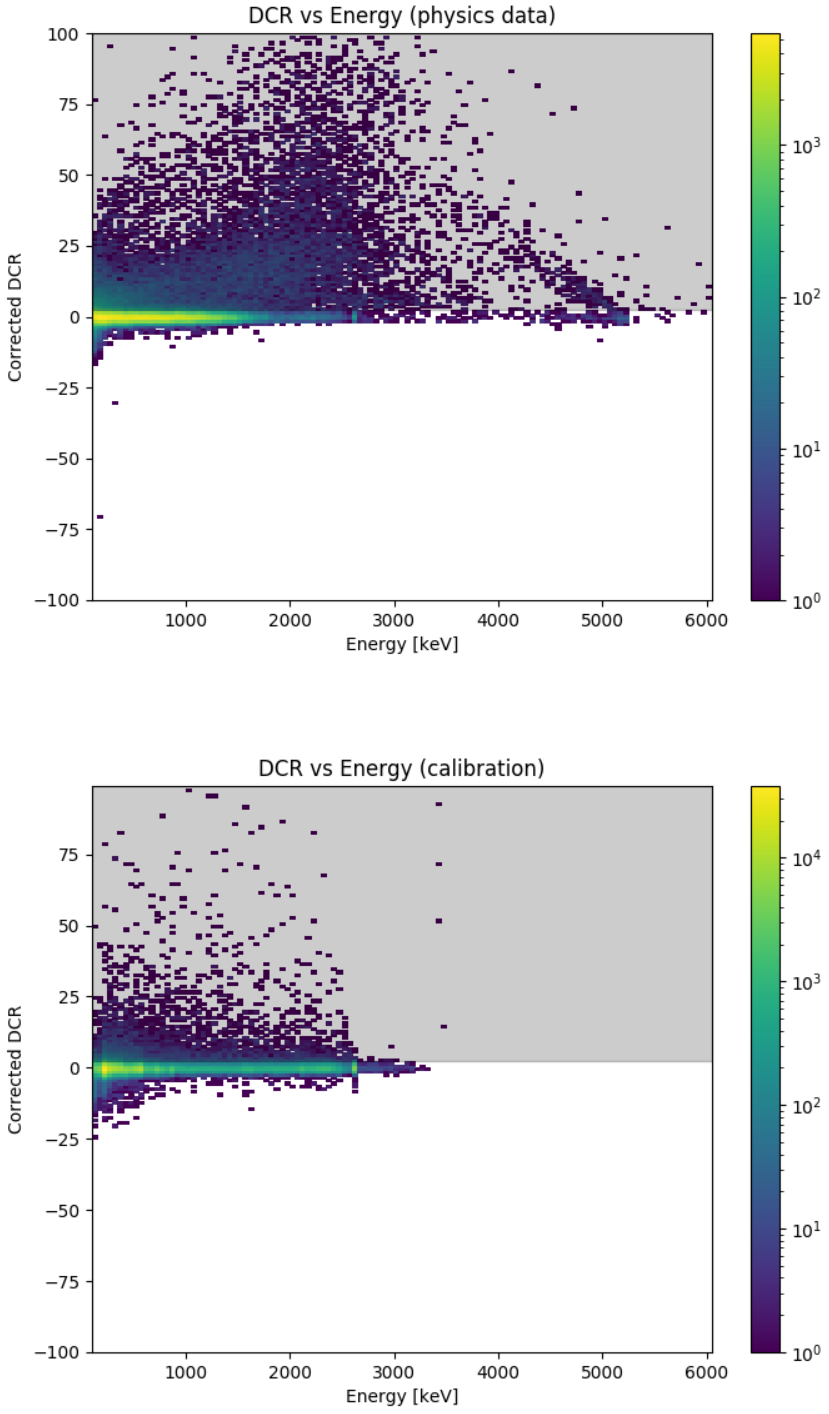

Figure 4.27: Corrected DCR vs E, for the full physics dataset and for one long calibration, with energy greater than 100 keV. The shaded grey indicates the region where the DCR parameter cuts events.

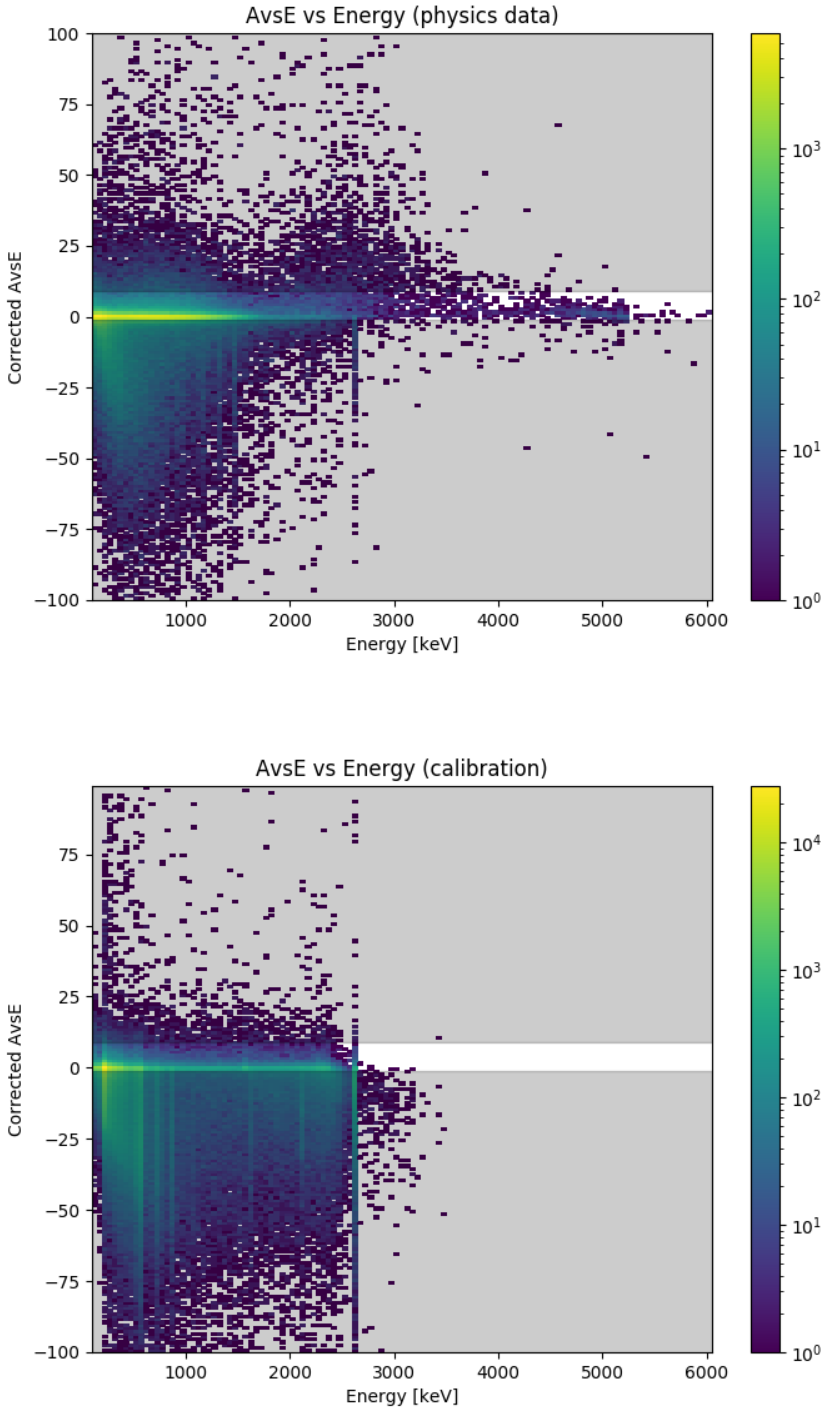

Figure 4.28: Corrected AvsE vs E, for the full physics dataset and for one long calibration, with energy greater than 100 keV. The shaded grey indicates the region where the AvsE parameter cuts events.

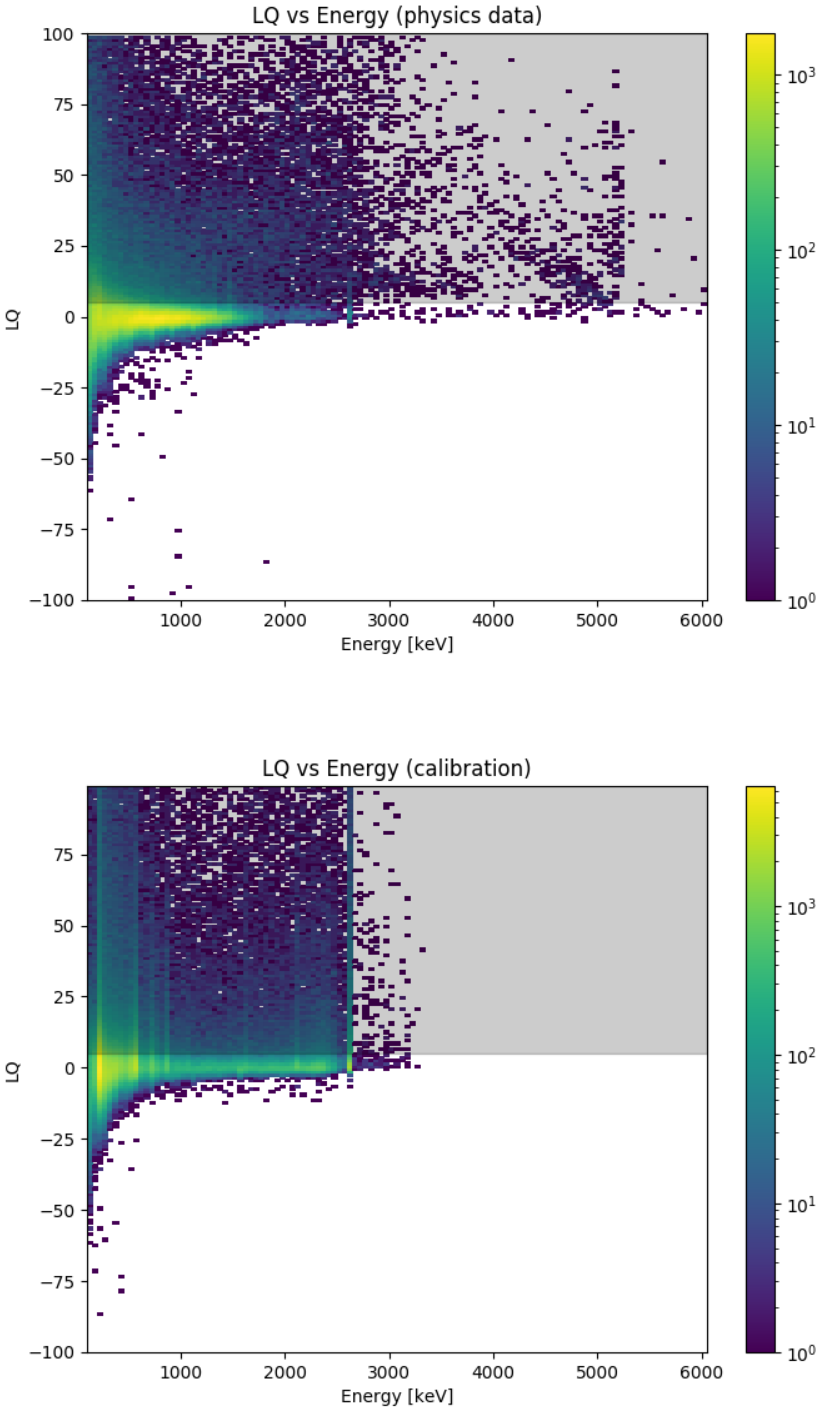

Figure 4.29: LQ vs Energy, for the full physics dataset and for one long calibration, with energy greater than 100 keV. The shaded grey indicates the region where the LQ parameter cuts events.

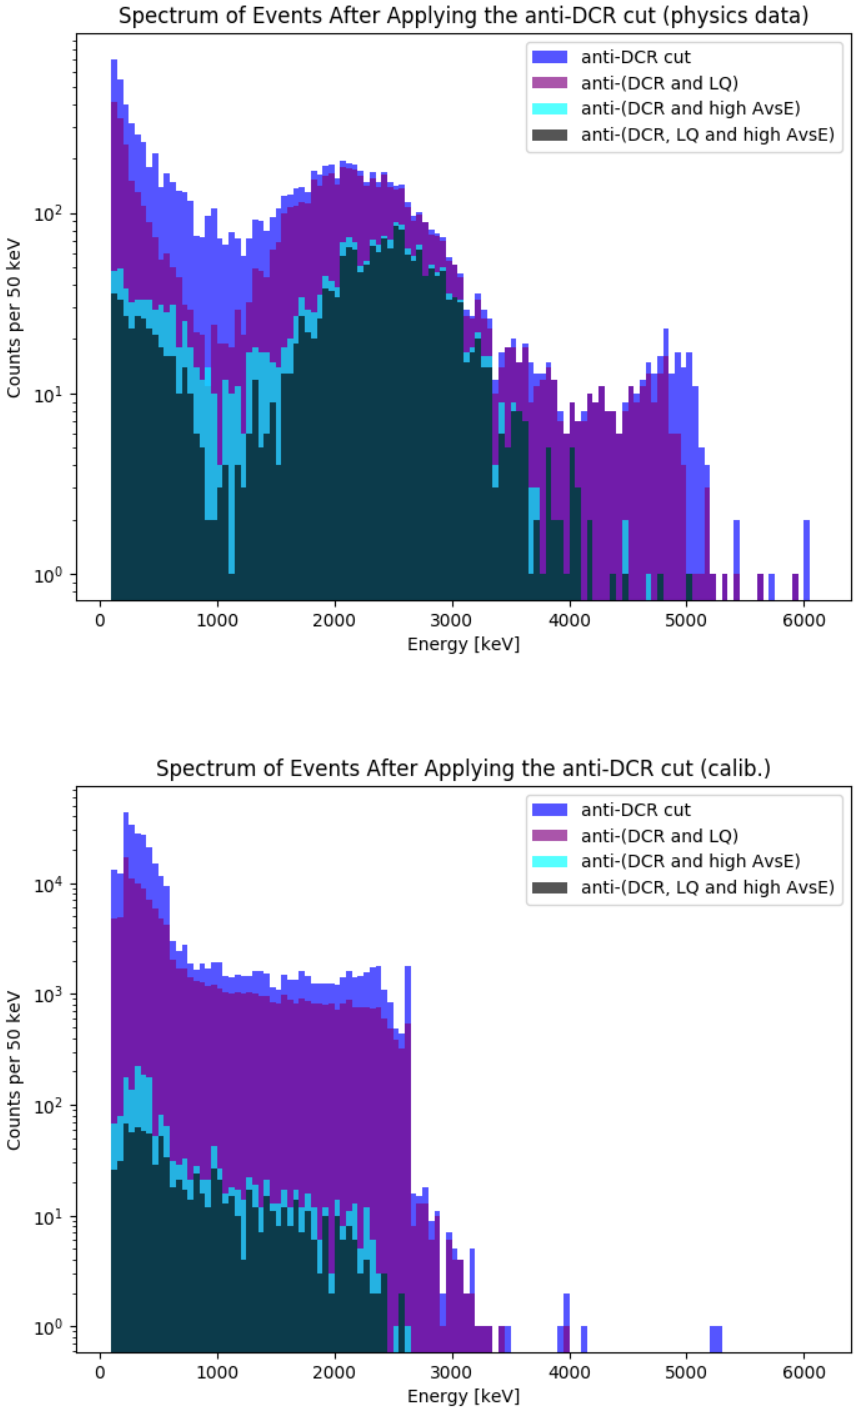

Figure 4.30: Events failing the DCR cut, or passing the anti-DCR cut, i.e. with $DCR > 2.326$. Also shown, events also failing AvsE, LQ, or both. Shown for both data and one calibration.

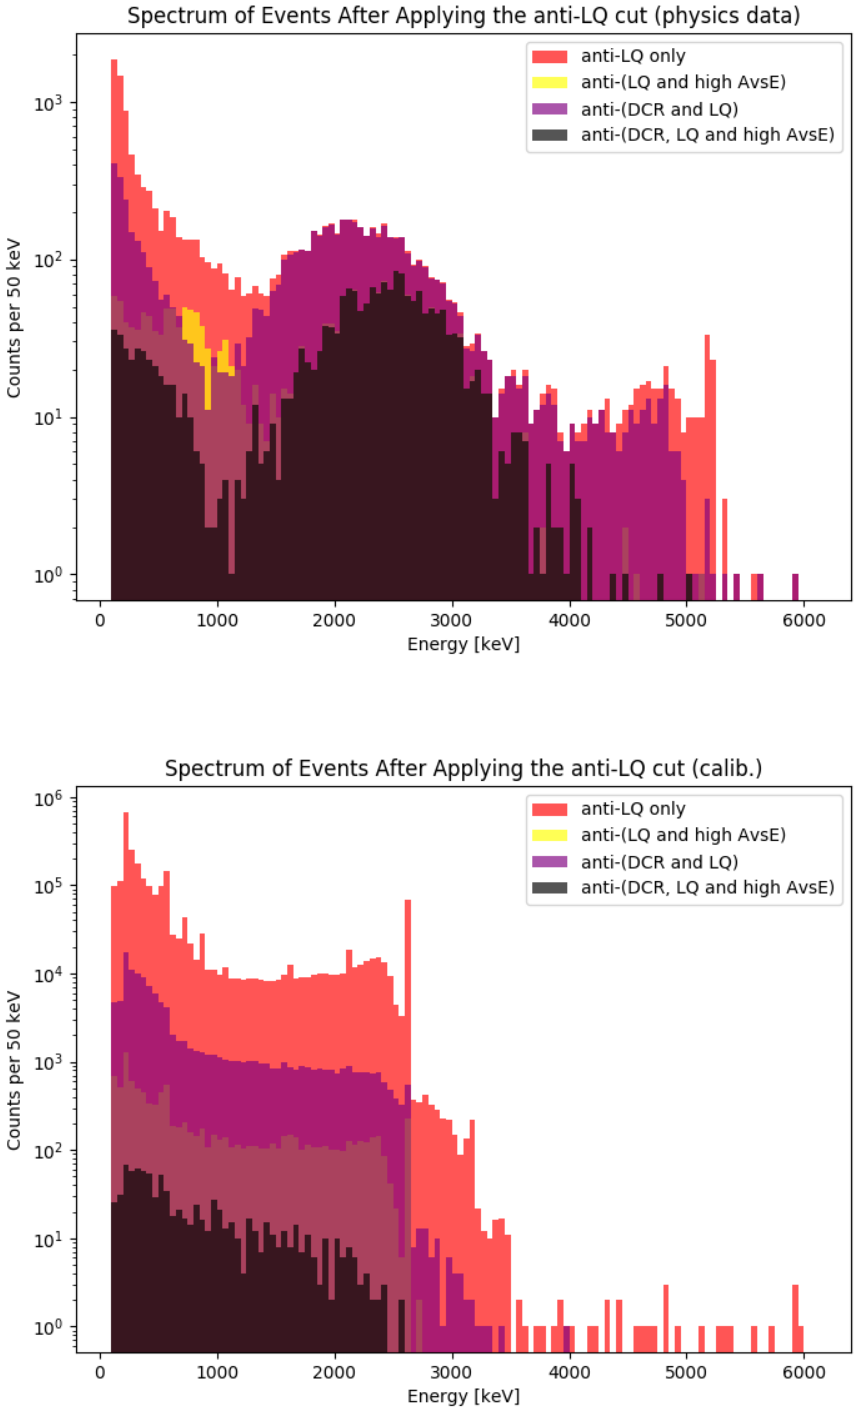

Figure 4.31: Events failing the LQ cut, or passing the anti-LQ cut, i.e. with $LQ > 10$. Also shown, events also failing AvsE, DCR, or both. Shown for both data and one calibration.

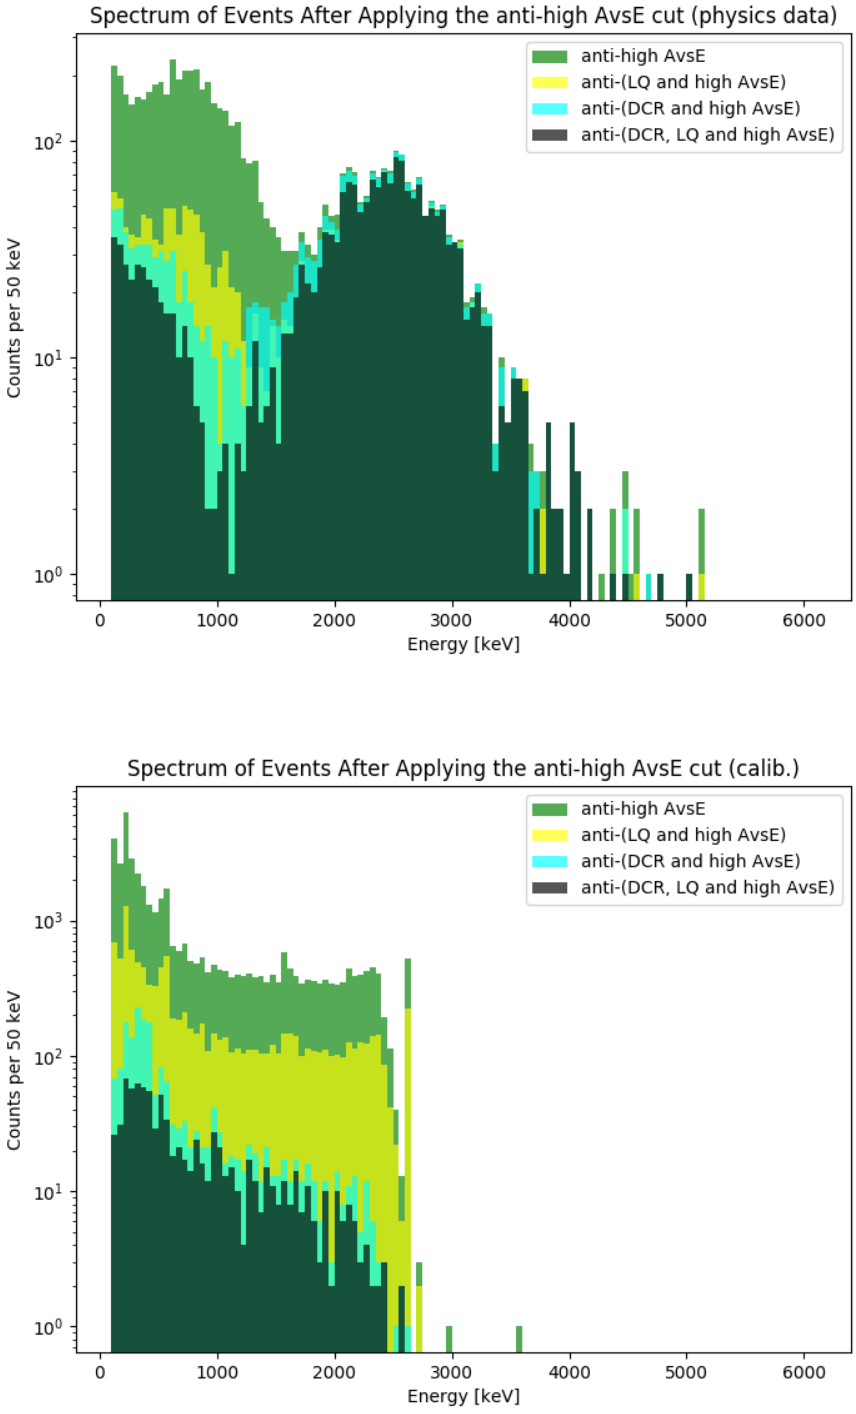

Figure 4.32: Events failing the high AvsE cut, or passing the anti-high AvsE cut, i.e. with $AvsE > 9$. Also shown, events also failing LQ, DCR, or both. Shown for both data and one calibration.

is better. As a cross check, the combination yields the same answer.

In summary, the above cross-checks and resulting plots show that

- DCR works as expected in reducing backgrounds from alphas on the passivated surface of our detector
- DCR has the expected statistical distribution and its narrow distribution for bulk events allows for excellent alpha discrimination
- DCR is crucial to MAJORANA DEMONSTRATOR as degraded alphas are most strongly peaked around 2000-2200 keV, our ROI

As such, DCR is critical to the final MAJORANA DEMONSTRATOR result, and setting a limit on the $0\nu\beta\beta$-decay process, which is also the topic of the next chapter.

Chapter 5

SEARCH FOR NEUTRINOLESS DOUBLE BETA DECAY

This chapter contains strictly original work done in order to run a parallel analysis of the MAJORANA DEMONSTRATOR data for neutrinoless double-beta decay. This work has been undertaken as a cross-check of the official MAJORANA DEMONSTRATOR result. As a reminder, the lower limit on the ^{76}Ge half-life $T^{0\nu}_{1/2}$ is estimated as:

$$T^{0\nu}_{1/2} > \frac{ln(2)NT\epsilon}{S} \tag{5.1}$$

Here, N is the number of ^{76}Ge atoms, T is the measurement time (their product NT being the exposure), and ϵ is the overall detection efficiency of the experiment. $S(C, B)$ is the number of signal events presumed for C observed events in the region of interest, for B expected background events. Different methods can be used to calculate this half-life. We will specifically describe the Feldman-Cousins (FC) Method as well as the Frequentist Extended Unbinned Negative Log Likelihood (EUNLL) Method implementations.

To calculate the half-life limit, we first need to find the events within our ROI that pass all the cuts described in Chapter 3. The energies of these events are needed when computing the result with the EUNLL method, but not when using the FC method.

5.1 What's Included

As a reminder, the following string includes all cuts which are applied to the background data for PPC detectors:

```
!isMuVeto && !(isLNFill1&&C==1) && !(isLNFill2&&C==2) && mHL==1 && isGood
&& isEnr && !isICPC && wfDCBits==0 && avse_corr>-1 && dcr_corr<2.326 &&
avse_corr<9 && lq<5 && Final_Energy>100
```

Data		Counts in Window		Background Index (10^{-3} c/(keV kg y))			ROI	Expected ROI
Set	Access	M1	M2	M1	M2	comb.	(keV)	Background Counts
DS0	open	11	0	$27.20^{+9.06}_{-7.41}$	0	$27.20^{+9.06}_{-7.41}$	3.75	0.11
DS1	comb.	4	0	$4.92^{+2.89}_{-2.07}$	0	$4.92^{+2.89}_{-2.07}$	3.79	0.04
DS2	comb.	1	0	$2.48^{+3.37}_{-1.73}$	0	$2.48^{+3.37}_{-1.73}$	3.75	0.01
DS3	open	0	0	< 2.73	0	< 2.73	3.75	0.00
DS4	open	0	0	0	< 10.21	< 10.21	3.53	0.00
DS5a	open	9	0	$11.63^{+4.32}_{-3.46}$	< 4.69	$9.22^{+3.43}_{-2.74}$	4.08	0.10
DS5b	open	0	0	< 2.00	< 5.86	< 1.49	3.77	0.00
DS5c	comb.	5	0	$8.65^{+4.47}_{-3.32}$	< 4.69	$6.41^{+3.31}_{-2.46}$	3.70	0.05
DS6a	comb.	22	7	$5.90^{+1.35}_{-1.17}$	$4.42^{+1.89}_{-1.47}$	$5.46^{+1.08}_{-0.95}$	3.63	0.29
DS6b	comb.	23	6	$9.41^{+2.11}_{-1.83}$	$5.57^{+2.60}_{-1.98}$	$8.24^{+1.63}_{-1.44}$	3.62	0.29
DS6c	comb.	20	0	$5.97^{+1.44}_{-1.24}$	< 0.67	$4.19^{+1.01}_{-0.87}$	3.62	0.20
DS7	comb.	17	0	$10.56^{+2.78}_{-2.36}$	0	$10.56^{+2.78}_{-2.36}$	4.04	0.19
DS8P	comb.	19	5	$11.32^{+2.81}_{-2.41}$	$5.97^{+3.08}_{-2.29}$	$9.54^{+2.09}_{-1.82}$	3.46	0.23
DS8I	comb.	0	4	0	$3.94^{+2.31}_{-1.66}$	$3.94^{+2.31}_{-1.66}$	3.57	0.04
DS0-8	open	58	9	$8.39^{+1.16}_{-1.06}$	$3.81^{+1.42}_{-1.13}$	$7.22^{+0.93}_{-0.85}$	3.57	0.66
DS0-8	blind	73	13	$7.53^{+0.92}_{-0.85}$	$3.07^{+0.93}_{-0.78}$	$6.18^{+0.70}_{-0.65}$	3.57	0.85
DS0-8	comb.	131	22	$7.89^{+0.72}_{-0.68}$	$3.34^{+0.76}_{-0.66}$	$6.59^{+0.56}_{-0.53}$	3.57	1.52
DS1-4,5b-8	open	38	9	$6.62^{+1.14}_{-1.02}$	$4.16^{+1.55}_{-1.24}$	$5.95^{+0.92}_{-0.83}$	3.57	0.47
DS1-4,5b-8	comb.	111	22	$7.19^{+0.71}_{-0.67}$	$3.44^{+0.79}_{-0.68}$	$6.09^{+0.55}_{-0.52}$	3.57	1.32
DS1-4,5b-6	comb.	39				$4.65^{+0.79}_{-0.71}$	4.32	0.02

Table 5.1: The background rates in the enriched detectors for each data set using a 360 keV window, from 1950 to 2350 keV, excluding 5 keV on each side of the main peaks seen in the region of interest (the SEP at 2103 keV, 2118 keV and 2204 keV), and excluding 5 keV on either side around the Q-value of 2039 keV. M1 and M2 are the 2 modules of the MAJORANA DEMONSTRATOR, Module 1 and Module 2. More details about these calculations can be found in the 2022 neutrinoless double-beta decay unidoc [60].

Data Set	Access	M1 Exposure (kg-y)	M2 Exposure (kg-y)	$NT\epsilon_{tot}\epsilon_{ROI}$ (10^{24} A-y)
DS0	open	1.12 ± 0.02	0.00 ± 0.00	$5.33^{+0.04}_{-0.04}$
DS1	open	1.81 ± 0.03	0.00 ± 0.00	$8.47^{+0.04}_{-0.04}$
DS1	blind	0.45 ± 0.01	0.00 ± 0.00	$2.08^{+0.04}_{-0.04}$
DS2	open	0.27 ± 0.00	0.00 ± 0.00	$1.25^{+0.05}_{-0.05}$
DS2	blind	0.85 ± 0.01	0.00 ± 0.00	$3.96^{+0.05}_{-0.05}$
DS3	open	0.97 ± 0.01	0.00 ± 0.00	$4.51^{+0.05}_{-0.05}$
DS4	open	0.00 ± 0.00	0.26 ± 0.00	$1.24^{+0.06}_{-0.05}$
DS5a	open	2.15 ± 0.03	0.56 ± 0.01	$12.79^{+0.05}_{-0.06}$
DS5b	open	1.32 ± 0.02	0.45 ± 0.01	$8.32^{+0.05}_{-0.06}$
DS5c	open	0.40 ± 0.01	0.14 ± 0.00	$2.51^{+0.05}_{-0.06}$
DS5c	blind	1.21 ± 0.02	0.42 ± 0.01	$7.62^{+0.05}_{-0.06}$
DS6a	open	3.01 ± 0.04	1.32 ± 0.02	$20.32^{+0.04}_{-0.04}$
DS6a	blind	7.35 ± 0.11	3.07 ± 0.04	$48.93^{+0.04}_{-0.04}$
DS6b	open	1.81 ± 0.03	0.78 ± 0.01	$12.14^{+0.04}_{-0.04}$
DS6b	blind	4.98 ± 0.07	2.21 ± 0.03	$33.70^{+0.04}_{-0.04}$
DS6c	open	3.35 ± 0.05	1.40 ± 0.02	$22.18^{+0.04}_{-0.04}$
DS6c	blind	5.96 ± 0.09	2.55 ± 0.04	$39.75^{+0.04}_{-0.04}$
DS7	open	1.53 ± 0.02	0.00 ± 0.00	$7.15^{+0.05}_{-0.06}$
DS7	blind	2.95 ± 0.04	0.00 ± 0.00	$13.80^{+0.05}_{-0.06}$
DS8P	open	1.47 ± 0.02	0.76 ± 0.01	$10.57^{+0.05}_{-0.05}$
DS8P	blind	3.19 ± 0.05	1.56 ± 0.02	$22.43^{+0.05}_{-0.05}$
DS8I	open	0.00 ± 0.00	0.89 ± 0.01	$4.05^{+0.06}_{-0.08}$
DS8I	blind	0.00 ± 0.00	1.93 ± 0.03	$8.73^{+0.06}_{-0.08}$
DS0-8	all	46.14 ± 0.68	18.32 ± 0.26	$301.81^{+1.13}_{-1.18}$

Table 5.2: The active, enriched mass exposure for each data set [50]. The uncertainty for the totals accounts for correlated uncertainties such as the active mass uncertainty. The total exposure for datasets 0-8 is presented in the last line.

For the ICPC detectors, we use:

```
!isMuVeto && !(isLNFill2&&C==2) && mHL==1 && isGood && isICPC && wfDCBits==0
&& ORNL_AoverE>0 && ORNL_DCR<0 && ORNL_AoverE<10 && ORNL_LQ<0 && Final_Energy>100
```

These cover all the cuts described in chapter 3. A summary of the number of events and the ROI for each dataset can be found in Table 5.1. To optimize the signal ROI, we numerically solve for the upper and lower energy bounds that optimize the 3σ discovery potential. This calculation is made using the energy parameters determined for each dataset in [91] and the measured background index [60]. Post-processing on background model simulations combines the spectra from all detectors and decay chains to produce the background spectra for each component modeled. These are then weighted by component activity and combined to produce the full background model. A background index is calculated by integrating over the background estimation window. A framework was developed to handle the inputs to the model as distributions – as opposed to the previously used single values. Monte Carlo uncertainty propagation was used to build background index distributions from which uncertainties are extracted [18].

5.2 Feldman-Cousins Method

The Feldman-Cousins (FC) method as described in their seminal publication [70], provides a straightforward computation of confidence intervals for $S(C, B)$. This method throws out some of the sensitivity by ignoring energy information, although it is still normally fairly close in value to the final result. We first obtain the overall product of the exposure and the efficiency, $NT\epsilon$. This can be read off from Table 5.2, which summarizes the information on exposure and efficiency found in Chapter 3. We can then quickly compute the FC upper limit on S, or the number of signal events, by plugging in the values of the background counts in the ROI, here, $b = 1.52$ as in Table 5.1 for DS0-8, and the observed number of counts in the ROI, here, $n_0 = 1$ into a Feldman Cousins calculator developed for this purpose by the author, and tested against the Feldman-Cousins publication [70], as well as against the

official MAJORANA DEMONSTRATOR calculator, using RooStats [114, 95]. For these values, the result to two decimal places is $S = 2.89$. Using $NT\epsilon = 3.018 \pm 0.012 \times 10^{26}$, we obtain:

$$T^{0\nu}_{1/2} > \frac{ln(2)NT\epsilon}{S(C,B)} = \frac{ln(2) \times 3.018 \pm 0.012 \times 10^{26}}{2.89} \text{ yr} = 7.24 \pm 0.03 \times 10^{25} \text{ yr}. \quad (5.2)$$

The half-life limit obtained in Eq. 5.2 is the same as that in the official 2023 MAJORANA DEMONSTRATOR final result [26], which yields $T^{0\nu}_{1/2} > 7.2 \times 10^{25}$ yr.

Cross-checks have been done on the newly developed Feldman-Cousins calculator, to verify it yields correct results. The first of these checks is to reproduce Table IV from the Feldman-Cousins publication [70]. In this table appear the 90% C.L. intervals for the Poisson signal mean μ, for total events observed n_0, for known mean background b ranging from 0 to 3. The Feldman-Cousins table can be seen in Fig. 5.1. We note great agreement, with the exception of a couple of entries at low n_0 and high b, $n_0 = 0, b = 2$ and $n_0 = 0, b = 3$, as can be seen in Fig. 5.3. Another check is to use RooStats, courtesy of Wenqin Xu's technical document [123]. This is the official ROOT Feldman-Cousins Calculator, used for the MAJORANA DEMONSTRATOR results previously published. Here are presented 90% C.L. intervals for the Poisson signal mean μ, for total events observed n_0, for known mean background b ranging from 0 to 3, equivalent to Table IV from the Feldman-Cousins publication. We note that this official ROOT RooStats Feldman-Cousins Calculator, Fig. 5.2 agrees with the calculator developed for this parallel analysis, as seen in Fig. 5.3, and not the Feldman-Cousins publication, for $n_0 = 0, b = 2$ and $n_0 = 0, b = 3$. As such, the conclusion is that, in this low-background and low observed counts regime, the parallel calculator implemented in python for this purpose works well and yields the expected results. Consequently, this calculator is safe to use to derive a Feldman-Cousins upper limit on signal for this dissertation. One other thing to note is that the Feldman-Cousins method always over-covers the confidence level [70], as can be seen in Fig. 5.3.

$n_0 \backslash b$	0.0	0.5	1.0	1.5	2.0	2.5	3.0
0	0.00, 2.44	0.00, 1.94	0.00, 1.61	0.00, 1.33	0.00, 1.26	0.00, 1.18	0.00, 1.08
1	0.11, 4.36	0.00, 3.86	0.00, 3.36	0.00, 2.91	0.00, 2.53	0.00, 2.19	0.00, 1.88
2	0.53, 5.91	0.03, 5.41	0.00, 4.91	0.00, 4.41	0.00, 3.91	0.00, 3.45	0.00, 3.04
3	1.10, 7.42	0.60, 6.92	0.10, 6.42	0.00, 5.92	0.00, 5.42	0.00, 4.92	0.00, 4.42
4	1.47, 8.60	1.17, 8.10	0.74, 7.60	0.24, 7.10	0.00, 6.60	0.00, 6.10	0.00, 5.60
5	1.84, 9.99	1.53, 9.49	1.25, 8.99	0.93, 8.49	0.43, 7.99	0.00, 7.49	0.00, 6.99
6	2.21,11.47	1.90,10.97	1.61,10.47	1.33, 9.97	1.08, 9.47	0.65, 8.97	0.15, 8.47
7	3.56,12.53	3.06,12.03	2.56,11.53	2.09,11.03	1.59,10.53	1.18,10.03	0.89, 9.53
8	3.96,13.99	3.46,13.49	2.96,12.99	2.51,12.49	2.14,11.99	1.81,11.49	1.51,10.99
9	4.36,15.30	3.86,14.80	3.36,14.30	2.91,13.80	2.53,13.30	2.19,12.80	1.88,12.30

Figure 5.1: Table IV from the Feldman-Cousins publication [70], to compare to the home-made calculator in Table 5.3.

$n_0 \backslash b$	0.0	0.5	1.0	1.5	2.0	2.5	3.0
0	[0, 2.4825]	[0, 1.9575]	[0, 1.5975]	[0, 1.3275]	[0, 1.0725]	[0, 1.1775]	[0, 0.9525]
1	[0.11, 4.34]	[0.02, 3.83]	[0, 3.3525]	[0, 2.9025]	[0, 2.5125]	[0, 2.1825]	[0, 1.8675]
2	[0.53, 5.90]	[0.05, 5.40]	[0, 4.8975]	[0, 4.4025]	[0, 3.9075]	[0, 3.4425]	[0, 3.0075]

Figure 5.2: A replication of Table IV from the Feldman-Cousins publication [70], using RooStats, courtesy of Wenqin Xu's technical document [123].

$n_0 \backslash b$	0.0	0.5	1.0	1.5	2.0	2.5	3.0
0	0.00, 2.44	0.00, 1.94	0.00, 1.61	0.00, 1.33	0.00, 1.08	0.00, 1.18	0.00, 0.95
1	0.11, 4.36	0.00, 3.86	0.00, 3.36	0.00, 2.91	0.00, 2.53	0.00, 2.19	0.00, 1.88
2	0.53, 5.91	0.03, 5.41	0.00, 4.91	0.00, 4.41	0.00, 3.91	0.00, 3.45	0.00, 3.04
3	1.10, 7.42	0.60, 6.92	0.10, 6.42	0.00, 5.92	0.00, 5.42	0.00, 4.92	0.00, 4.42
4	1.47, 8.60	1.17, 8.10	0.74, 7.60	0.24, 7.10	0.00, 6.60	0.00, 6.10	0.00, 5.60
5	1.84, 9.99	1.53, 9.49	1.25, 8.99	0.93, 8.49	0.43, 7.99	0.00, 7.49	0.00, 6.99
6	2.21, 11.47	1.90, 10.97	1.61, 10.47	1.33, 9.97	1.09, 9.47	0.65, 8.97	0.15, 8.47
7	3.56, 12.53	3.06, 12.03	2.56, 11.53	2.09, 11.03	1.59, 10.53	1.18, 10.03	0.89, 9.53
8	3.96, 13.99	3.46, 13.49	2.96, 12.99	2.51, 12.49	2.14, 11.99	1.81, 11.49	1.51, 10.99
9	4.36, 15.30	3.86, 14.80	3.36, 14.30	2.91, 13.80	2.53, 13.30	2.19, 12.80	1.88, 12.30

Table 5.3: Outputs of the author's home-made Feldman-Cousins calculator, to replicate Table IV from the Feldman-Cousins publication [70], as seen in Figure 5.1. These are the 90% C.L. intervals for the Poisson signal mean μ, for total events observed n_0, for known mean background b ranging from 0 to 3. We see a perfect replication of the official FC results for $n_0 \geq 1$ where b is between 0 and 3.

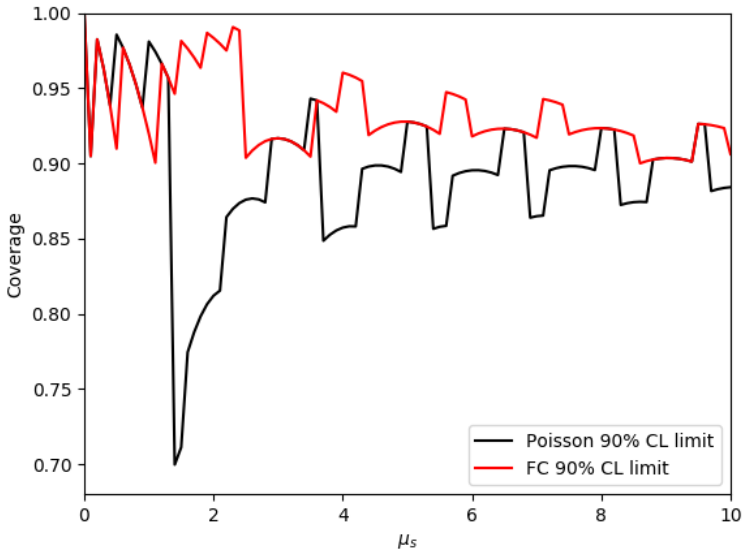

Figure 5.3: The coverage of the Feldman Cousins method and the Poisson statistical method. Note that Feldman-Cousins always over-covers the target confidence level (in this case, 90%).

5.3 Extended Unbinned Negative Log Likelihood (EUNLL) Method

The Frequentist Extended Unbinned Negative Log Likelihood (EUNLL) method can be used for a better estimate of the $0\nu\beta\beta$ decay half-life limit, taking into account the observed energies in the MAJORANA DEMONSTRATOR rather than just the number of counts in the ROI. For this dissertation, we use iminuit, which is a python frontend used to liaise with the Minuit2 C++-based statistical and numerical minimization software, and the MIGRAD minimization algorithm [68]. The energy of each event passing the cuts described above is obtained, using the estimator called Final_Energy.

5.3.1 A Brief First Look

As a first step, a method combining all datasets into one large dataset was used to check that the iminuit outputs are in the same ballpark as the expectation based on the FC

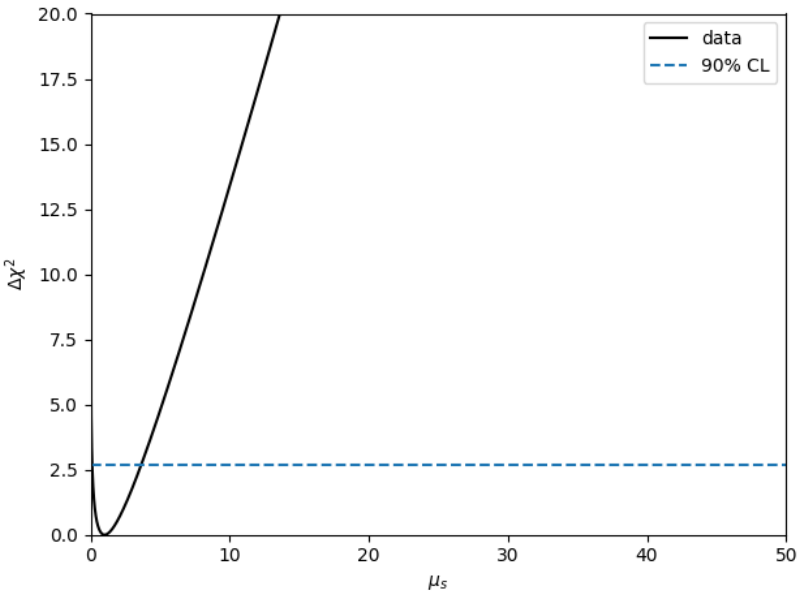

Figure 5.4: $\Delta\chi^2$ vs μ_s. The upper bound on μ_s in the plot corresponds to an upper limit on S.

method presented in the previous section. Naively, we expect a flat background and some peak shape, initially taken as a Gaussian peak plus an exponential decay towards lower energies. In Minuit, we can fix the μ and σ values of the Gaussian, as the peak should be at the $Q_{\beta\beta}$ value of 2039.061 keV with a spread related to our resolution: FWHM of ~2.5 keV, corresponding to a σ of ~1.05 keV. We can also include the parameters describing the exponential decay of the exponentially-modified Gaussian. Using iminuit, the procedure is to step through the values of S (the number of signal events) from 0 to 20 in steps of 0.01, and compute the χ^2 value for each step. Calculating the difference between the minimum chi-squared and the actual chi-squared for a given step, we find where the curve of $\Delta\chi^2$ crosses the equivalent of 90% confidence level, specifically at $\Delta\chi^2 = 1.645^2$. This method is also known as Wilks' Approximation. In our preliminary example, this yields an upper limit on S of 2.57.

This corresponds to a lower limit on the half-life of $0\nu\beta\beta$ in ^{76}Ge. We exclude ICPC events

for a total exposure of $2.909 \pm 0.011 \times 10^{26}$ A-yr. The naive implementation conducted for the purposes of verifying that the code runs consistently with expectation yielded the following result:

$$T_{1/2}^{0\nu} > \frac{ln(2)NT\epsilon}{S} = \frac{ln(2) \times 2.909 \pm 0.011 \times 10^{26}}{2.57}\text{yr} = 7.84 \pm 0.03 \times 10^{25} \text{ yr} \tag{5.3}$$

This result is indeed close to that which was expected, and is lower than the official Majorana Demonstrator result of 8.3×10^{25} yr [26], which makes sense considering it includes less information, i.e. no information from DS8I, the dataset containing information from the new ICPC detectors.

5.3.2 An Improved Implementation

For a better estimate of the $0\nu\beta\beta$ decay half-life limit, we split the data into its individual datasets, as described in Chapter 3. This is the correct thing to do, as the different Majorana Demonstrator datasets have different conditions, such as expected background or noise, and factors such as this can influence the performance during a dataset. For the best possible interpretation of the Majorana Demonstrator data these different conditions must be accounted for in the iminuit computations.

A few issues appeared when trying to calculate the overall half-life limit. To include datasets in which there were 0 events in our region of interest, the code needed to be modified – the initial implementation failed when using these datasets with zero events. For this dissertation the method described by Wenqin Xu in the Majorana Demonstrator Statistical Unidoc [123] was used. This document details how to approach this problem, and the solution is to add $2S \times \ln(2) \times NT\epsilon$ to the χ^2 for each zero-count dataset, for each value of S, the number of signal events stepped through. Extensive studies were carried out to check the RooStats implementation of this method and its performance with Majorana Demonstrator's extremely small sample size [123]. We focused on the challenging cases where zero events are observed in real experiments or Monte-Carlo simulations, i.e. the

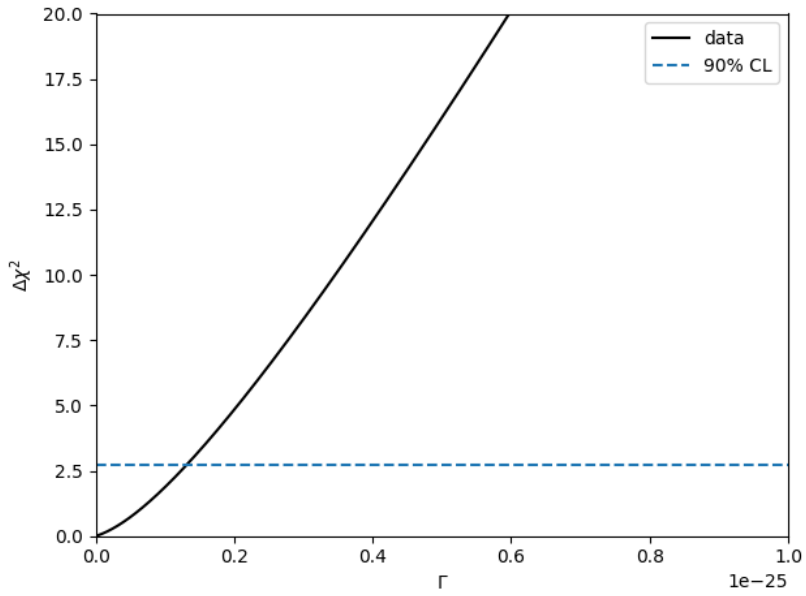

Figure 5.5: $\Delta\chi^2$ vs Γ. This curve shows graphically where the 90% CL maximum decay rate occurs. In this case, with all detectors, $\Gamma < 0.13 * 10^{-25}/yr$, corresponding to a lower half-life limit of $7.8 * 10^{25} yr$. Here Γ is not the rate, quite the rate but the inverse of the half-life. The two differ by a factor of ln(2).

empty datasets. Theoretical values of the test statistic were derived for empty datasets to be $2S$, where S is the number of expected events from all Poisson sources of interest. In the distribution of the test statistic, a peak of half of the test statistic corresponds to empty dataset, i.e. with zero observed events $N_{obs} = 0$. Since half of the test statistic is assigned to be 1 by RooStats [114], $2S$ corresponds to the full value [123]. This will increase the χ^2 and naturally the $\Delta\chi^2$, making the $\Delta\chi^2$ cross the critical threshold earlier, i.e., the best-fit maximum Γ (decay rate) gets smaller, i.e. the half-life limit increases.

5.3.3 Description of iminuit Procedure

To use iminuit and the Extended Unbinned Negative Log Likelihood method, one has to provide a density function as well as a range of values on which the density function is applied. The software uses these pieces of information in order to compute a cost function which to minimize. In our case, we compute said cost function for each dataset with non-zero events in our region of interest (ROI). As discussed previously, for datasets with zero events, one later adds $2S \times \ln(2) \times NT_i\epsilon_i$ to the χ^2 for each zero-count dataset i.

The likelihood function over all datasets can be described as follows:

$$\mathcal{L} = \prod_{i=DS0}^{DS8I} \mathcal{L}_i \tag{5.4}$$

where

$$\mathcal{L}_i = \mathcal{L}_{i,shape} \times \mathcal{L}_{i,shape} \times \mathcal{L}_{i,nuisance}. \tag{5.5}$$

The rate and shape parts of the likelihood function are:

$$\mathcal{L}_{i,rate} = \frac{(n_{s,i} \times b_i)^{N_i}}{N_i} \tag{5.6}$$

and

$$\mathcal{L}_{i,shape} = \prod_{k=1}^{N_i} \left(\frac{n_{s,i}}{n_{s,i} + b_i} \frac{dP_{sig}(E_i)}{dE} + \frac{b_i}{n_{s,i} + b_i} \frac{dP_{bg}(E_i)}{dE} \right) \tag{5.7}$$

while the nuisance parameter part is

$$\mathcal{L}_{i,nuisance} = \prod_{k=1}^{n} \frac{e^{-\frac{(\vec{\theta}_{k,i} - \hat{\theta}_{k,i})^2}{2\sigma^2_{k,i}}}}{\sqrt{2\pi\sigma^2_{k,i}}} \tag{5.8}$$

where θ_k describes the k-th nuisance parameter.

The density function tracks the total number of events for each dataset i:

$$n_i = \ln(2)\Gamma_{1/2} N_A \epsilon_i (NT)_i f_{76} / M_{\text{mol}} + b_i = n_{s,i} + b_i. \tag{5.9}$$

The dP is described below:

$$dP/dE = dP_{signal}/dE + dP_{bg}/dE \tag{5.10}$$

$$dP_{signal}/dE = ((1 - f_{tail,DS}) \times \frac{1}{\sigma_{DS}\sqrt{2\pi}} \times e^{\frac{-(x-\mu)^2}{2\sigma_{DS}^2}} + f_{tail,DS} \times \frac{1}{2\sigma_{DS}\tau_{DS}} \times exp((\frac{1}{2 * \sigma_{DS}^2 \times \tau_{DS}^2} - \frac{x}{\sigma_{DS}\tau_{DS}}) \times erfc(\frac{-1 - x\sigma_{DS}\tau_{DS}}{\sigma_{DS}\tau_{DS} \times \sqrt{2}})) \tag{5.11}$$

$$dP_{bg}/dE = \frac{b_{DS}}{(xr[1] - xr[0])} \tag{5.12}$$

An explanation of important factors used in the Minuit minimization is presented below:

- N_A is Avogadro's Number $(6.02 * 10^{23})$
- ϵ_{DS} is the overall efficiency of our cuts for dataset DS
- $(NT)_{DS}$ is the overall exposure for dataset DS
- f_{76} is the fraction of ^{76}Ge in enriched detectors in Majorana Demonstrator
- M_{mol} is the average molar mass of the germanium
- Γ is the neutrinoless double beta decay rate for ^{76}Ge
- $f_{tail,DS}$ is the tail fraction, i.e. the fraction of events that are in the exponential decay for dataset DS
- the term associated with 1-$f_{tail,DS}$ is a normal distribution with mean μ and standard deviation σ_{DS} for each DS

- the term associated with $f_{tail,DS}$ is an exponentially-modified normal distribution with mean μ and standard deviation σ_{DS} with an exponential tail towards lower energy, with an exponential decay τ_{DS} for each DS.

- b_{DS} is the number of background events in each dataset DS, with the denominator term indicating a uniform distribution between xr[0] and xr[1], in this case between 1950-2350 keV. The background is treated like a constant, flat in E

The first five factors combine to yield the rightmost column in Table 5.2.

iminuit uses the MIGRAD algorithm to minimize the density function over the interval on which it is defined. The overall cost function c is computed using iminuit's in-built cost.ExtendedUnbinnedNLL method [68].

As this is a log likelihood, we can describe it as

$$-2 * log(\mathcal{L}) = \sum_{DS} -2 * log(\mathcal{L}_{DS}). \tag{5.13}$$

We instantiate a Minuit object with a large number (93) of parameters, as described above, with cost function c. We fix all the known values, such as $\mu = 2039.061$, each with its associated error, e.g. $\mu_{err} = 0.007$, since the $Q_{\beta\beta}$ for ^{76}Ge is 2039.061 ± 0.007 keV [96]. The values and errors of the parameters can be found in exhaustive detail in Chapter 3 for each dataset separately in case the values vary by dataset. Some parameters have fixed values, such as the tail fractions, the σs and the τs for each dataset. On the other hand, the total efficiencies, the exposures, the αs, the μs for each dataset, the isotopic fractions for PPCs and ICPCs and the molar masses for PPCs and ICPCs are allowed to float – these are the nuisance parameters θ_k. The total efficiencies and their uncertainties used in the iminuit analysis can be found in Table 3.2, in the second rightmost column. The exposures and their uncertainties can be found in Table 3.1, in the rightmost column. The isotopic fraction is found in [52] to be $87.4 \pm 0.5\%$ for the PPCs and $88.0 \pm 1.0\%$ for the ICPCs. The molar mass M_{mol} of the enriched material and is 75.668 ± 0.010 g/mol for the PPCs and 75.681 ± 0.020 g/mol for the ICPCs [52], as also described in Chapter 3. The values for the

αs and their uncertainties, the uncertainties on the μ values, as well as the values of the σs and the τs are presented in Fig. 5.6, taken from [60].

We use two different methods to compute the half-life limit. First, we use Wilks' Approximation. Second, we use the full treatment.

5.3.4 Wilks' Approximation

We use a linear space in $T_{1/2}^{-1}$, or the inverse half-life of the hypothetical ^{76}Ge neutrinoless double beta decay process. The linear space is from 0 to 2×10^{-26}, in 201 steps. For each step, we compute the value of the χ^2 at each value of $T_{1/2}^{-1}$, which is the output of the migrad minimization, or migrad.fval. We add $2 \times \Gamma \times \ln(2) \times f_{76} \times NT\epsilon$ to the χ^2 for each zero-count dataset, again for each step in the inverse half-life. $\ln(2)\Gamma$ is the same as $T_{1/2}^{-1}$. Then, we compute the minimum of this function over the linear space in $T_{1/2}^{-1}$. To obtain a 90% C.L. value on the decay rate, we find where the χ^2 curve crosses the value of 1.645^2. Due to computing time limitations, linear interpolation between the values of χ^2 on either side of where χ^2 passes the above threshold value is used to determine exactly where the crossing occurs. We obtain the rate $\Gamma_{critical}$, which for Wilks' Approximation is where $\chi^2 = 1.645^2$. This informs us of the 90% confidence interval for the value of the inverse half-life $T_{1/2}^{-1}$, which in this case is one-sided, i.e. we obtain an upper limit on the inverse half-life at 90% C.L, at $\Gamma_{critical}$. To convert into a lower limit on the half-life $T_{0\nu}^{1/2}$, one must simply invert the upper limit on the inverse half-life.

The output is:

$$T_{1/2}^{0\nu} > 8.71 \times 10^{25} \text{ yr} \tag{5.14}$$

at 90% C.L.

This is 4.9% higher than the official value of $8.3 \times 10^{25} yr$ [26].

Data Set	σ (keV)	f_{tail}	τ (keV)	δ_μ (keV)	δ_α (keV)
DS0	1.08	0.220	1.25	0.137	0.029
DS1	1.10	0.191	1.51	0.136	0.030
DS2	1.08	0.203	1.55	0.190	0.030
DS3	1.08	0.207	1.48	0.172	0.030
DS4	1.06	0.094	1.17	0.147	0.049
DS5a	1.21	0.141	1.66	0.089	0.027
DS5b	1.11	0.148	1.58	0.087	0.032
DS5c	1.08	0.173	1.53	0.176	0.032
DS6a	1.06	0.167	1.55	0.164	0.031
DS6b	1.05	0.183	1.45	0.137	0.031
DS6c	1.05	0.186	1.45	0.158	0.031
DS7	1.16	0.230	1.52	0.162	0.028
DS8P	1.02	0.134	1.64	0.166	0.032
DS8I	1.08	0.068	2.77	0.114	0.036

Figure 5.6: The values of parameters used in iminuit. Taken from [60]. The tail fractions, the σs and the τs are fixed. The δ_αs and the δ_μs represent the uncertainties of the αs and the μs.

5.3.5 Full Statistical Treatment

We use a linear space in $T_{1/2}^{-1}$, or the inverse half-life of the hypothetical ^{76}Ge neutrinoless double beta decay process, from 0 to 2×10^{-26}, in 201 steps. For each step, we compute the value of the χ^2 at each value of the inverse half-life, which is the output of the migrad minimization, or migrad.fval. We add $2 \times \Gamma \times \ln(2) \times f_{76} \times NT\epsilon$ to the χ^2 for each zero-count dataset, again for each step in $T_{1/2}^{-1}$. Then, we compute the minimum of this function over a linear space in the inverse half-life. For each value, we find the χ^2_{crit} such that 90% of toy MC simulations are below that χ^2 value. A preliminary graph of $T_{1/2}^{-1}$ vs χ^2_{crit}, without full statistics, is presented in Fig. 5.7. Due to computing time limitations, linear interpolation between the values of χ^2 on either side of where χ^2 passes the above threshold value is used to determine exactly where the crossing occurs. This informs us of the 90% confidence interval for the value of the inverse half-life $T_{1/2}^{-1}$, which in this case is one-sided, i.e. we obtain an upper limit on $T_{1/2}^{-1}$ at 90% C.L, at $\Gamma_{critical}$. To convert into a lower limit on the half-life $T_{0\nu}^{1/2}$, one must simply invert the upper limit on $T_{1/2}^{-1}$.

The output is:

$$T_{0\nu}^{1/2} > 8.41 \times 10^{25} \text{ yr} \tag{5.15}$$

at 90% C.L.

This is about 1.4% higher than the official value of $8.3 \times 10^{25} yr$ [26] – very close!

5.3.6 Correlation Matrix Investigations

The correlation matrix has been produced in order to verify correlations between variables. No huge correlations between the different parameters are observed (the maximum absolute value is 0.62, with a negative correlation). The two largest positive and negative correlations are in entries 0 and 5, 10 and 13, 10 and 17, 0 and 54, 17 and 49, 17 and 50 in Fig. 5.8. The definition of the parameters can be found in Table 5.4. Most of the relatively highly correlated and anti-correlated parameters occur in DS6a which is one of the larger datasets and has a large number of events associated with it. Perhaps more importantly, it features

Param. Number	Param. Name	Param. Number	Param. Name	Param. Number	Param. Name
0	isotopic fraction (PPC)	20	exposure (5a)	40	μ (2)
1	atomic mass (PPC)	21	exposure (5c)	41	α (2)
2	isotopic fraction (ICPC)	22	ϵ_{tot} (5a)	42	exposure (2)
3	atomic mass (ICPC)	23	ϵ_{tot} (5c)	43	ϵ_{tot} (2)
4	μ (6a)	24	μ (0)	44	μ (1)
5	α (6a)	25	α (0)	45	α (1)
6	μ (6b)	26	exposure (0)	46	exposure (1)
7	α (6b)	27	ϵ_{tot} (0)	47	ϵ_{tot} (1)
8	μ (6c)	28	μ (7)	48	rate (Γ)
9	α (6c)	29	μ (8)	49	bkg. events (0)
10	exposure (6a)	30	α (7)	50	bkg. events (1)
11	exposure (6b)	31	α (8)	51	bkg. events (2)
12	exposure (6c)	32	exposure (7)	52	bkg. events (7)
13	ϵ_{tot} (6a)	33	exposure (8)	53	bkg. events (8)
14	ϵ_{tot} (6b)	34	ϵ_{tot} (7)	54	bkg. events (8i)
15	ϵ_{tot} (6c)	35	ϵ_{tot} (8)	55	bkg. events (6a)
16	μ (5a)	36	μ (8i)	56	bkg. events (6b)
17	α (5a)	37	α (8i)	57	bkg. events (6c)
18	μ (5c)	38	exposure (8i)	58	bkg. events (5a)
19	α (5c)	39	ϵ_{tot} (8i)	59	bkg. events (5c)

Table 5.4: Parameter numbers in the correlation matrix vs. their name and meaning to the right of the parameter number. Dataset of the parameter in parantheses.

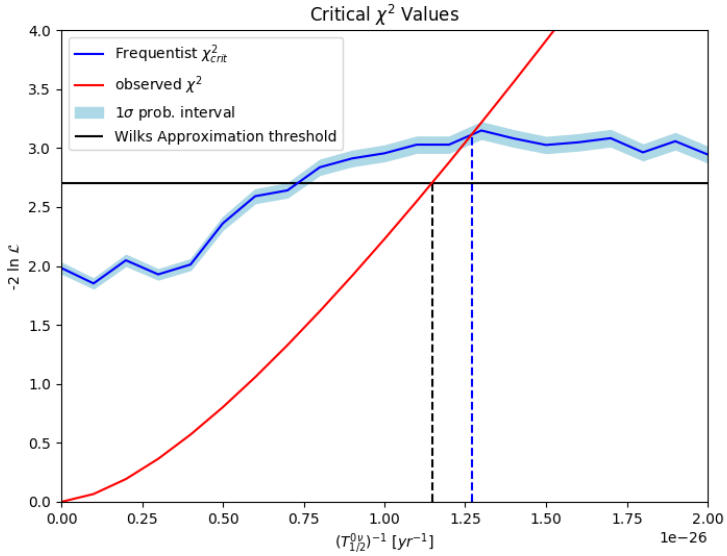

Figure 5.7: For each rate, or inverse half-life, $\frac{1}{T^{0\nu\beta\beta}_{1/2}}$, we find the χ^2_{crit} such that 90% of toy MC simulations are below that χ^2 value.

the event which is closest to the $Q_{\beta\beta}$ value of 2039 keV. This event affects the iminuit results more than others, leading to increased correlations between its important parameters. A high negative correlation between the alpha (peak shape width) and the active exposure of DS6a seems to make sense. DS6a is the only DS with any constraint on the width – the presence of the one event in the ROI pulls the rate and shape uncertainties at the same time as we test different values of a possible signal. Interestingly, the highest negative correlations are between the isotopic fraction and the active exposure of DS6a and between the isotopic fraction and the number of background events in DS6a. This is consistent with the discussion below in subsection 5.5.5, regarding the bias checks. If there are fewer than expected events, there should be an anti-correlation, because iminuit would have to pull the other variables lower in order to satisfy its constraints using the cost function c. All other correlations are below an absolute value of 0.5, and we ignore these here as they are deemed small enough.

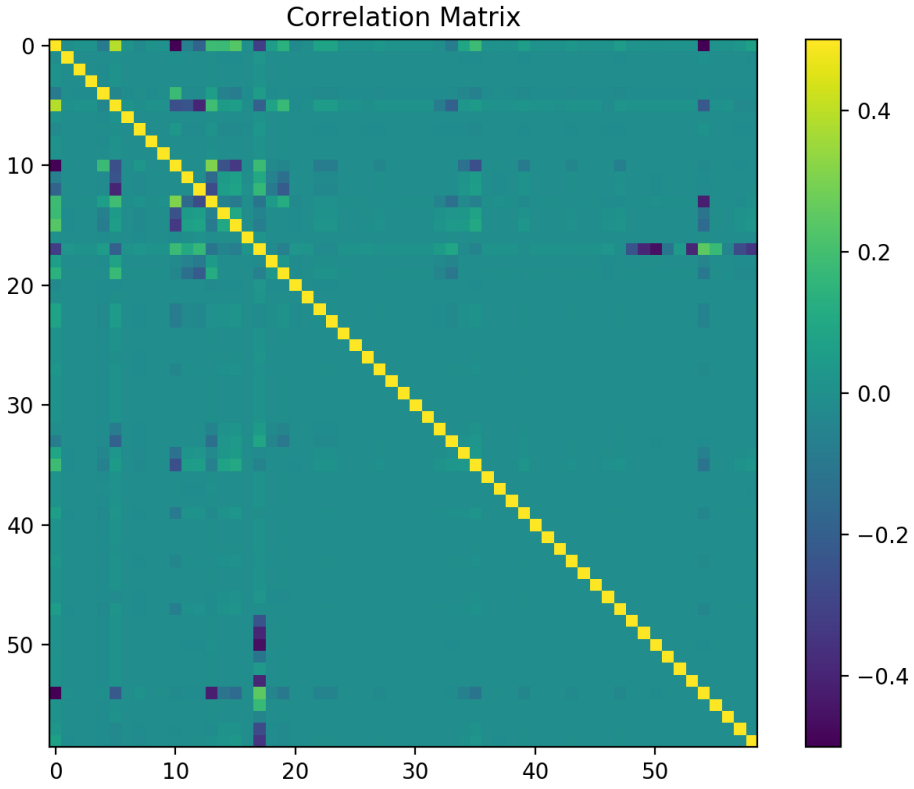

Figure 5.8: Correlation Matrix. As in any correlation matrix, the diagonal entries are equal to 1. Otherwise, the matrix colors displayed are normalized to -0.5 and 0.5.

5.4 Half-life Sensitivity Study

We calculate the sensitivity of the MAJORANA DEMONSTRATOR experiment by performing a Monte Carlo (MC) based parallel analysis.

For each dataset in particular, the approximate background rate and the exposure are known. To create the MC simulation, random floating point numbers are generated with a flat distribution, taking into account the background rate, which can be read off from Table 5.1, for a given energy range (1950-2350 keV) and the given exposure of a dataset, following Table 5.2. Due to low statistics, DS2, 3, 4 and 5b are handled together, using the DS1-4, 5b-6 background rate listed in the table. If the generated random number is smaller than the background level for the given dataset, a new event is added to the dataset's event list, with an energy randomly selected from a uniform distribution within 1950 to 2350 keV, excluding the simulated events falling in the excluded energy ranges described above. In the end, for each simulation and each dataset, a simulated set of energies are obtained. We repeat this process N times, where N is maximized based on available computing resources; N is taken to be 10000 for a 1% statistical uncertainty.

As a first step, Wilks' Approximation is used in order to calculate a sensitivity for the experiment. We use similar code to that described above to calculate the χ^2 curve for each individual simulation. For each value of $T_{1/2}^{-1}$ from 0 to 2×10^{-26} in increments of 10^{-27}, the χ^2 value is calculated, and converted to a 2-sided p-value based on the Gaussian distribution. Then, to calculate the 90% C.L. sensitivity, linear interpolation is used to see where to p-value crosses the threshold value of $p = 0.1$. In our case, this happens at $\frac{1}{T_{1/2}^{0\nu\beta\beta}} = 1.24 \times 10^{-26}$, corresponding to a half-life sensitivity of $T_{1/2}^{0\nu\beta\beta} = 8.1 \times 10^{25} yr$, the same as in the official final MAJORANA DEMONSTRATOR result [26]. A visual representation of the half-life sensitivity study can be seen in Fig. 5.9. The $0\nu\beta\beta$ half life where the p-value of the median of the simulation ensemble (dashed line) is 0.1 corresponds to the sensitivity, here 8.06×10^{25} years. The $0\nu\beta\beta$ half life where the p-value of the observed Majorana data (solid line) equals 0.1 corresponds to the observed 90% lower limit on the half life, here 8.71×10^{25} years.

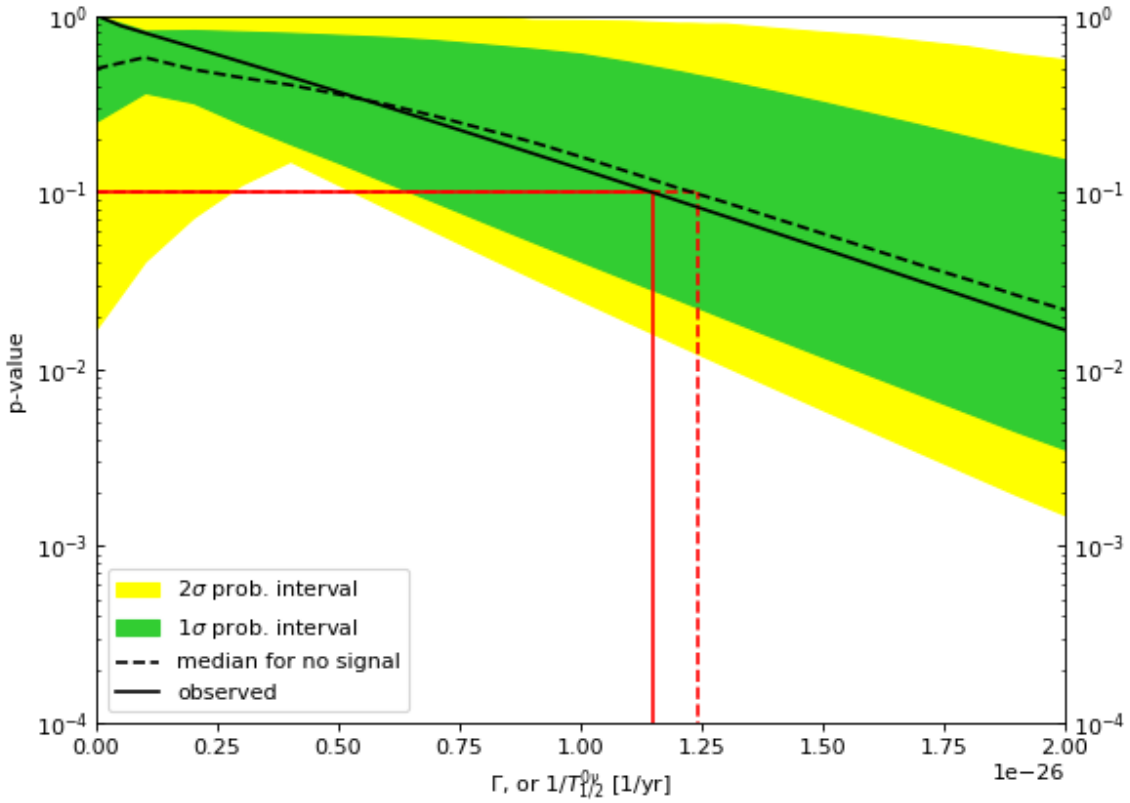

Figure 5.9: The p-value as a function of $0\nu\beta\beta$ decay half life obtained from the Frequentist profile likelihood method for DS0-8 using Wilks' Approximation. Here Γ is not the rate, quite the rate but the inverse of the half-life. The two differ by a factor of ln(2).

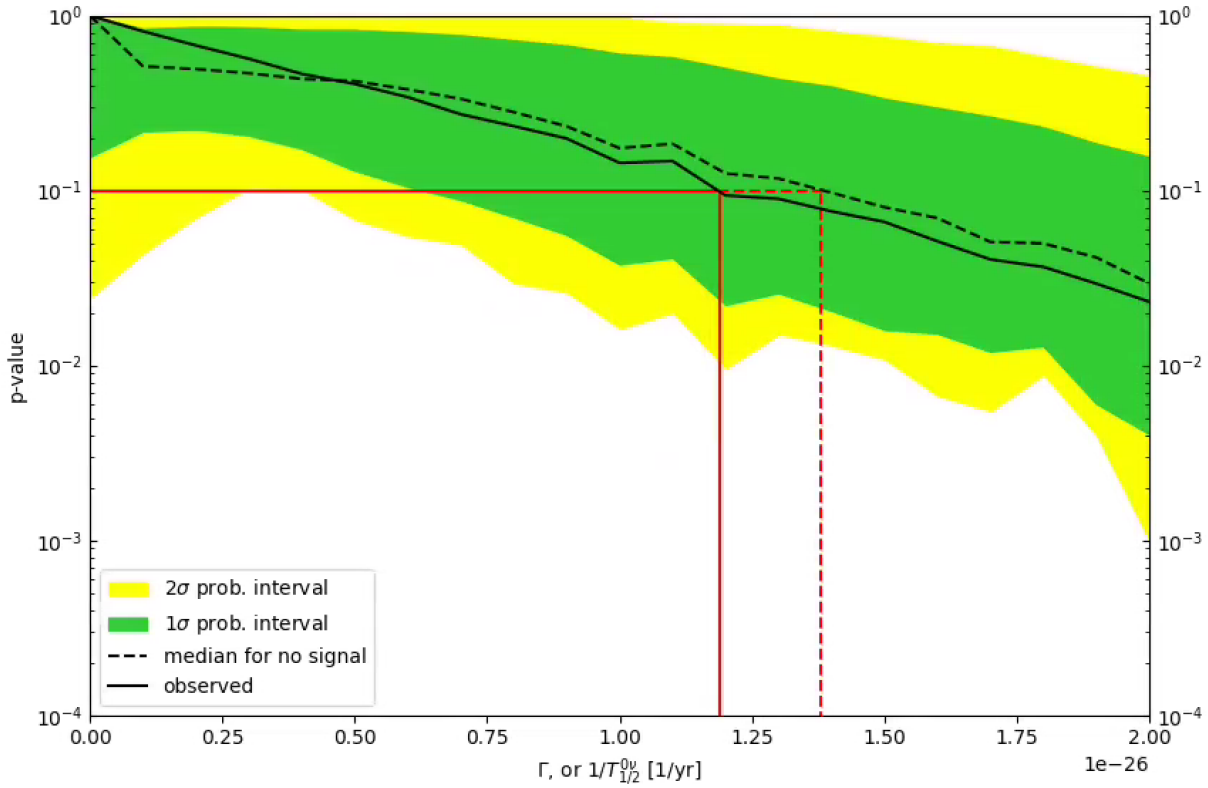

Figure 5.10: Full stats plot for the $0\nu\beta\beta$ half-life and sensitivity study. The p-value as a function of $0\nu\beta\beta$ decay half life obtained from the Frequentist profile likelihood method for DS0-8. Here Γ is not the rate, quite the rate but the inverse of the half-life. The two differ by a factor of ln(2).

The Wilks' Approximation uses a Gaussian distribution to compute the p-value, which is expected to be an approximation, albeit a decent one. For the full treatment, MC simulations are done, to compute an array of CDF distributions for a given rate $T_{1/2}^{-1}$, from 0 to 2×10^{-26} in increments of 10^{-27}. We simulate N datasets using each $T_{1/2}^{-1}$, and compute how far away from the mean of the distribution each event is, in units of σ. Then, these values are converted to reference p-values, for each $T_{1/2}^{-1}$, using the actual CDF distributions, instead of using the Gaussian approximation. Afterwards, we use iminuit to recompute the χ^2 values for another set of simulations, in the same manner as for the Wilks' Approximation method. To compute the observed MAJORANA DEMONSTRATOR data's half-life limit, we compare the obtained χ^2 values for each $T_{1/2}^{-1}$ to our set of reference p-values. This allows us to determine the final p-values for a given rate $T_{1/2}^{-1}$, from 0 to 2×10^{-26} in increments of 10^{-27}. Specifically, we find the percentile of the newly obtained χ^2 values relative to the reference p-values. These values are plotted in Fig. 5.10. Then, to calculate the 90% C.L. sensitivity, linear interpolation is used to see where to p-value crosses the threshold value of $p = 0.1$. We also find the median, 1 σ and 2 σ bands of the simulated p-values, and plot them in Fig. 5.10.

We see many issues with low statistics, as the Gaussian approximation is well known and implemented in numpy whereas the CDF distributions and reference p-values are computed using self-developed code, and can create a more jagged appearance of Fig. 5.10. compared to Fig. 5.9. To significantly improve the output of this code, e.g. its smoothness, many more simulations would be required. In theory this should be straightforward, but it is rather time intensive and also quite computationally expensive. With low statistics, we note that the output is similar to before for the observed 90% lower limit on the half life, i.e. where the p-value of the observed Majorana data (solid line) equals 0.1, $\frac{1}{T_{1/2}^{0\nu\beta\beta}} = 1.19 \times 10^{-26}$, equivalent to 8.41×10^{25} years. However, the sensitivity of this process, i.e. where the p-value of the median of the simulation ensemble (dashed line) is 0.1, now corresponds to $\frac{1}{T_{1/2}^{0\nu\beta\beta}} = 1.38 \times 10^{-26}$, or a sensitivity of 7.25×10^{25} years, significantly lower than for Wilks' Approximation.

5.5 Cross-Checking the Code

Here, we provide an overview of the cross checks, and describe the strategy of using Wilks' approximation to study them. The first cross-check is to verify what happens when the parameter uncertainties provided to iminuit are made larger by several factors, or made 0, which is equivalent to fixing said parameters. This is done for both rate (exposure, efficiency) and shape (μ, α) parameters, separately and in combination. We also double check the rate variation by using the Rolke-Lopez method [89] implemented in TRolke in ROOT. Another cross-check is done in which we assume uncertainties for the different parameters are either fully correlated, partly correlated and uncorrelated across datasets – in the end, this makes a negligible difference on our results, as can be seen below. A cross-check was done in order to check that if we combine two datasets or have them separated but with parameters tied together, we obtain the same half-life limit, to check whether the fitting code is correctly handling the likelihood computations for combinations of datasets. Another cross-check was performed to verify that the outputs from iminuit make sense, and are consistent with the parameters input to iminuit, given their uncertainties. For each parameter, we compared the output of the iminuit best fit to the expected value of the parameter, i.e. the value that was used to create the simulated dataset. For simplicity and ease of computation, we use Wilks' Approximation to undergo the cross-check studies. These studies are meant to double check iminuit itself, rather than the statistical methods behind it, and so it makes little difference whether we use the approximation or the full statistical version – the latter only takes longer and is more computationally intensive.

5.5.1 Effect of the Uncertainty on the Half-Life Limit

In Tables 5.5 and 5.6, we present the results of multiplying the errors by several factors. The former includes an event in our data near the $Q_{\beta\beta}$ at 2039.64 keV, while the latter ignores its effects. The specific process was to multiply the uncertainties computed by the Majorana Demonstrator collaboration, as given in their respective technical documents,

by a given factor and use the NormalConstraint method in iminuit. As expected, the higher the errors, the lower the computed half-life limit. The zero error multiplier corresponds to fixed parameters. We fix parameters that are constrained for the final result reported above, to investigate the effect on the half-life limit. Fixing parameters is equivalent to reducing the uncertainty associated with them to 0, i.e. we should expect the half-life limit to increase when some previously-constrained parameters are fixed. Fixing these parameters makes a very minimal difference on the half-life limit, so we conclude a $\sim 1\%$ uncertainty and fixed parameters are almost equivalent in iminuit for our purposes.

Another investigation was to multiply the calculated uncertainties by given factors to investigate the effects. If the uncertainty is multiplied by a small factor, the half-life limit changes very little. E.g., doubling all the uncertainties changes the limit by 0.25%, a small if not negligible change. A large part of this is due to an event in our data near the $Q_{\beta\beta}$ at 2039.64 keV. In Table 5.6, we see that even multiplying the errors by a factor of 10 changes the limit by only 4%. Note that the presence of that one ROI event makes a difference of a factor of 2 in our half life limit! Comparing to Table 5.5, we see that a similar change with the near-ROI event included would be a change in the limit of circa 14%. As such, the studies underwent to fix the systematic uncertainties for all parameters by members of the Majorana Demonstrator Collaboration are still important.

5.5.2 Rate Systematic Study

It is interesting to note that the rate and shape uncertainties affect the results about the same, and a separate study has been done to double-check if this is correct. To cross-check iminuit's output it is quite important to also use a different, verified, method. For this study it was decided to use ROOT's incorporated TRolke class which uses a different implementation based on the method described in the Rolke-Lopez paper [89]. A simplified model based on the Majorana Demonstrator numbers was used. Using the known Majorana Demonstrator total number of background events, 152, but, for simplicity, assuming 0 events in the 4 keV ROI, and assuming an overall efficiency of 0.78, we can

compare the relative change as we vary the uncertainty in the efficiency to the rate numbers in Table 5.6, which removes the event at ROI event at 2039.64 keV from computation. It should make no difference whether this cross-check is completed assuming zero or one events in the ROI, as long as we compare with the appropriate iminuit result table. This can be thought of as a direct cross-check of the validity of the iminuit method. We use the efficiency as the error on the active exposure is much smaller, and would thus have a minor effect. We expect that if the relative change in the upper limit using the TRolke computation is on the same order as the relative change in the iminuit output for the half-life limit, then we can conclude that the iminuit output is verified. Using an initial uncertainty in the efficiency of 0.04, we scale it, as for Table 5.6, by factors of 0, 2, 5 and 10. The results, in units of 1 divided by the upper limit, to compare to the half-life limit (since, by Eq. 5.2, they are inversely proportional), are presented in Table 5.7. The slightly higher percent errors for iminuit can be explained by the fact that the active exposure uncertainty, while smaller, does have some contribution. It is perhaps a bit surprising that a 4% uncertainty on the efficiency only results in a 0.02% change in the upper limit. This shows that the upper limit is dominated almost completely by the Poisson statistics in the ROI. It also means that it is important to choose the central values for the rate uncertainties as accurately as possible. Therefore, we conclude the cross-check yields the expected results and the iminuit outputs are valid and good to use. Due to the small changes even for a huge 10x change in the rate uncertainty, this cross-check is enough to validate rate uncertainties.

5.5.3 Correlation Cross-Checks

Following Wenqin Xu's statistical unidoc [123], we use three different methods to calculate the half-life limit, described below. In our default model, the uncertainties from different datasets are treated as if they are uncorrelated. However, some of the uncertainties could be common across the datasets. It has been speculated that efficiencies and exposures could have a component of 5% uncertainty, relative to the value of the efficiency or exposure being considered, that is the same across all datasets (which is what we call "fully corre-

Error Multiplier (parameters)	Half-life limit (years)
1 (Main Result)	8.7101×10^{25}
0 (shape only, fixed)	8.7139×10^{25}
0 (rate only, fixed)	8.7137×10^{25}
0 (rate and shape, fixed)	8.7174×10^{25}
2 (shape only)	8.6992×10^{25}
2 (rate only)	8.6994×10^{25}
2 (rate and shape)	8.6885×10^{25}
5 (shape only)	8.6358×10^{25}
5 (rate only)	8.6243×10^{25}
5 (rate and shape)	8.5505×10^{25}
10 (shape only)	8.0818×10^{25}
10 (rate only)	8.3512×10^{25}
10 (rate and shape)	7.5095×10^{25}

Table 5.5: This is how the half-life limits vary for different constraints. The half-life limit is reported. As expected, the higher the errors, the lower the computed half-life limit. 0 error multiplier corresponds to fixed parameters. Fixing the parameter makes a very minimal difference on the half-life limit.

Error Multiplier (parameters)	Half-life limit (years)
1 (Main Result)	1.6604×10^{26}
0 (shape only, fixed)	1.6606×10^{26}
0 (rate only, fixed)	1.6607×10^{26}
0 (rate and shape, fixed)	1.6609×10^{26}
2 (shape only)	1.6597×10^{26}
2 (rate only)	1.6595×10^{26}
2 (rate and shape)	1.6588×10^{26}
5 (shape only)	1.6547×10^{26}
5 (rate only)	1.6532×10^{26}
5 (rate and shape)	1.6475×10^{26}
10 (shape only)	1.6214×10^{26}
10 (rate only)	1.6302×10^{26}
10 (rate and shape)	1.5926×10^{26}

Table 5.6: This is how the half-life limits vary for different constraints when we exclude the 2039.64 keV event.

Error Multiplier (E.M.)	Rolke Upper Limit	Inverse Rolke Upper Limit	% Difference from E.M. 1	% Difference between Table 5.6 E.M. 1 and E.M. for rate
0	1.1304	0.8846	0.02%	0.02%
1	1.1306	0.8845	0	0
2	1.1311	0.8841	0.04%	0.07%
5	1.1347	0.8813	0.36%	0.45%
10	1.1480	0.8710	1.54%	1.87%

Table 5.7: TRolke vs iminuit percent difference with an error multiplier on rate uncertainties. This table shows that iminuit and TRolke give roughly consistent results.

lated"), and the rest of the uncertainty could be dataset-independent, which is what we call "uncorrelated". For the semi-correlated, or partly correlated case, we assume that the maximum uncertainty, relative to the value of the parameter, considered across all datasets represents the correlated part of the uncertainty, and anything else is uncorrelated across datasets. While the unidoc applies this to only the efficiency uncertainty, we apply it to all uncertainties, and split it into rate and shape. For many of these uncertainties, 5% relative uncertainty is vastly greater than calculated by the MAJORANA DEMONSTRATOR team, making the uncertainty greater than expected. The limits of the three approaches are listed in Table 5.8. There is no change to the obtained limits to two significant figures, and the observed changes are only about O(0.1%). This comparison demonstrates that the treatment of the correlations among the uncertainties only has a negligible quantitative effect on the final limits. The tiny effect is not unexpected since the uncertainties themselves have little impact on the limit to begin with, as demonstrated in the previous section. Indeed we expect the fully correlated numbers to be somewhere between those of the error multiplier 1 and 2 in Table 5.5, and probably closer to 1. Similarly, the semi-correlated numbers should end up in between the non-correlated and fully correlated numbers – and they do. As such, we conclude that Table 5.8 is virtually as expected and whether the uncertainties are fully, partly or not correlated only makes an extremely minor difference on the half-life limit of only O(0.1%). In retrospect, it would have been good to know this in advance, as we could have reduced the number of free parameters to just a few, rather than having a large number of parameters for each dataset.

5.5.4 Combining and Splitting Datasets

A cross-check was done in order to check that if we combine two datasets, here taken to be DS6b and DS6c, or have them separated but with parameters tied together, we obtain the same half-life limit. This test checks whether the fitting code is correctly handling the likelihood computations for combinations of datasets. The value of the half-life limit in these cases is 8.69×10^{25} years, slightly lower than the value with these datasets completely separate,

Description	Half-life limit (years)
DSs with non-correlated uncertainties	8.7101×10^{25}
Fully correlated (5% uncertainty) – rate	8.7043×10^{25}
Fully correlated (5% uncertainty) – shape	8.7101×10^{25}
Fully correlated (5% uncertainty) – rate and shape	8.7043×10^{25}
Semi-Correlated uncertainties (rate)	8.7070×10^{25}
Semi-Correlated uncertainties (shape)	8.7101×10^{25}
Semi-Correlated uncertainties (rate and shape)	8.7070×10^{25}

Table 5.8: This is how the half-life limits vary for different ways of defining uncertainties. We reproduce the method and followed by Wenqin Xu in the Statistical Unidoc [123]. I report the half-life limit. As in the unidoc, there is little variation with correlated, uncorrelated and fixed relative uncertainties.

which makes sense, as some information is being essentially lost when we tie parameters together.

Alternatively, we split DS6c into 2 sub-datasets, with parameters tied together, and split its events properly into 2 bins depending on the run number. We assume the parameters such as the efficiency, the mean of $Q_{\beta\beta}$ and the alpha and their uncertainties do not vary, that is, that they remain the same for each sub-dataset. The active exposure is split such that the exposure of the sub-datasets adds to the full active exposure of DS6c. We obtain the same result as for the full-DS6c treatment, 8.71×10^{25}, indicating splitting datasets and tying their parameters together or having one larger dataset yields the same answer in iminuit.

5.5.5 Bias Check

Another cross-check was performed to verify that the outputs from iminuit make sense, and are consistent/close to what's expected given the parameters calculated by the MAJORANA DEMONSTRATOR Collaboration, and the uncertainties of these parameters. We performed

repeated fits to simulated datasets, which were obtained by using the known MAJORANA DEMONSTRATOR parameters. For each parameter, we compared the output of the iminuit best fit to the expected value of the parameter, i.e. the value that was used to create the simulated dataset. The success of this cross-check is visible in Fig. 5.11. Overall, we see that the individual outputs from iminuit (each individual blue point) is very close to the expected mean of that parameter. We also see that the means of the distribution made up of individual iminuit outputs are very close to zero, and within the MAJORANA DEMONSTRATOR calculated uncertainties (i.e. the green stars have an absolute value of less than 1). The 1-sigma spread (interval, depicted in red, between the red pluses) captures the expected mean of the distribution in virtually all cases. The error on the mean is depicted as a green line around the green dot, which represents the mean value of all the simulations. We note that in some cases the error of the mean does not extend to 0. This is the case for the isotopic fraction, active exposures of DS6a-c, the epsilon total, or total efficiencies, of DS6a-c, and for some other total efficiencies in other DSs. The biased-looking deviations were expected when a simulation with an event in the ROI occurs, which is more likely to happen in the longer datasets. This outcome also speaks to the fact that there could be relatively fewer events than expected in a number of datasets when compared to the case of computed values of the total efficiencies, active exposures and isotopic fraction, leading the iminuit minimization algorithm to find a lower value for these numbers as the best fit. This is consistent with the fact that the solid black curve in Fig. 5.10 is lower than the dashed black line at where p is 0.1, i.e. in an average simulation of MAJORANA DEMONSTRATOR we see more events. Overall, we conclude that the iminuit outputs are consistent with expectations.

5.6 Limits on $m_{\beta\beta}$

With all cross-checks having passed, the obtained half-life limit of $0\nu\beta\beta$ in ^{76}Ge is taken to be, using the full Frequentist implementation,

$$T_{1/2}^{0\nu\beta\beta} > 8.41 \times 10^{25}\text{yr}. \tag{5.16}$$

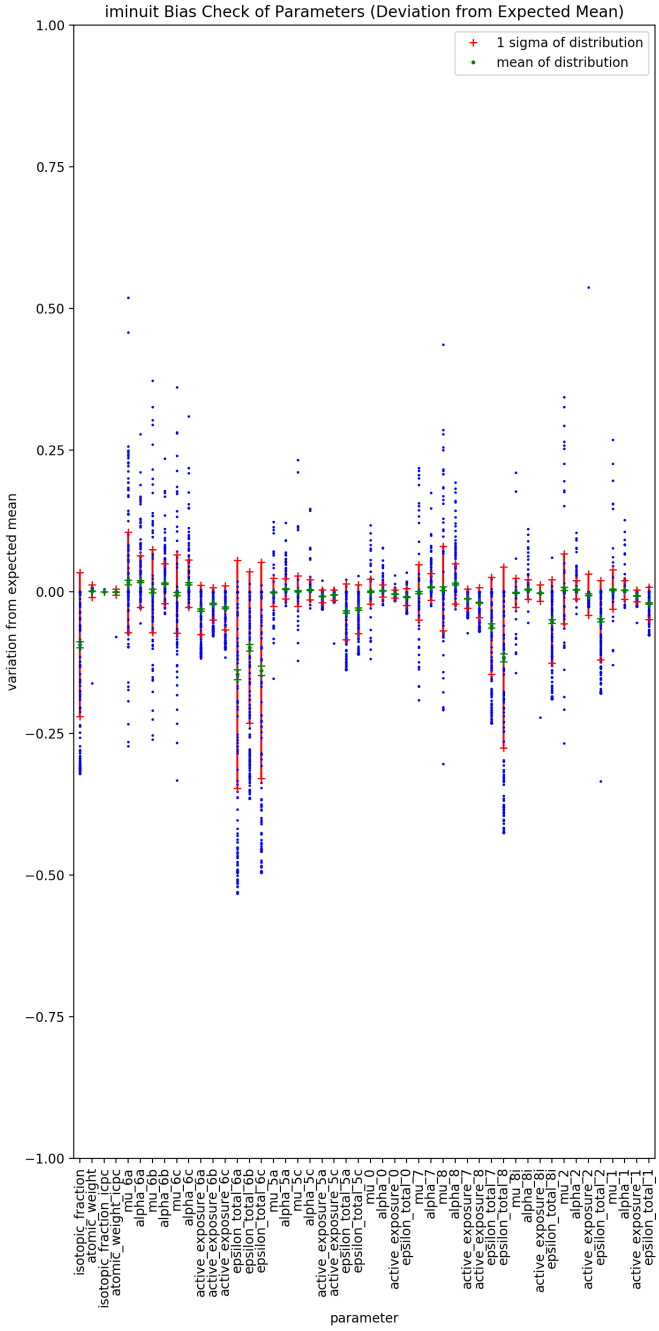

Figure 5.11: Deviations of expected means of parameters across bias check simulations. The means are drawn with error bars signifying the errors on the means.

Similarly, we report the half-life sensitivity as

$$T_{1/2}^{0\nu\beta\beta} > 7.25 \times 10^{25}\text{yr}. \tag{5.17}$$

.

The half-life of the $0\nu\beta\beta$ decay process is the experimental observable which has been probed by the MAJORANA DEMONSTRATOR. In this section we connect the half life limit to the physics underlying neutrinoless double-beta decay. Although the are many possible mechanisms, we focus on light left-handed neutrino exchange and report a limit on $m_{\beta\beta}$, the coherent sum of the neutrino masses. To reach a limit on the $m_{\beta\beta}$, we must also use nuclear matrix elements (NMEs), calculated through many body nuclear theory [56]. To calculate an effective Majorana mass $m_{\beta\beta}$, one needs to include the expected half-life of neutrinoless double beta decay (or, specifically, a lower limit on this half-life), a phase space factor $G_{0\nu}$, an effective Majorana mass $m_{\beta\beta}$ and an NME $M_{0\nu}$. Recall that the equation relating these factors to the half-life is:

$$[T_{1/2}^{0\nu\beta\beta}]^{-1} = G_{0\nu}\ |M_{0\nu}|^2\ |m_{\beta\beta}|^2 \tag{5.18}$$

or, to calculate a limit on the effective neutrino mass $\langle m_{\beta\beta}\rangle$,

$$\langle m_{\beta\beta}\rangle = \frac{511\text{ keV}}{g_A^{eff^2}|M_{Ge}^{0\nu}|\sqrt{T_{1/2}\times G_{\text{Ge}}^{0\nu}}} \tag{5.19}$$

Above,

- $g_A^{eff} = 1.27$ is the effective axial weak coupling factor for a free neutron.
- $|M_{Ge}^{0\nu}|$ is the nuclear matrix element for $0\nu\beta\beta$ in ^{76}Ge – here, we are using a range of values from $2.66-6.34$ due to different theoretical values obtained by different research groups [93, 78, 48, 98, 80, 81, 69, 118, 115, 30, 51].
- $G_{\text{Ge}}^{0\nu}$ is the phase space factor, for which the latest derived value is 2.37×10^{-15}/yr [94].

We report a range of $\langle m_{\beta\beta} \rangle$ corresponding to the range of nuclear matrix elements used, 2.66-6.34, and $T^{0\nu}_{1/2} > 8.41 \times 10^{25}$ yr. We find $\langle m_{\beta\beta} \rangle < 112 - 267$ meV. This limit depends on theoretical uncertainties in $|M^{0\nu}_{Ge}|$ and g^{eff}_A.

Chapter 6

CONCLUSIONS FROM THE MAJORANA DEMONSTRATOR AND FUTURE PATHS

The discovery of the still hypothetical $0\nu\beta\beta$ process would fundamentally change our knowledge of the Universe, and of physics in particular. While the MAJORANA DEMONSTRATOR did not find evidence for $0\nu\beta\beta$, we set new limits on the value of its half-life, and also searched for other new physics beyond the standard model. That said, the search is not nearly over, and it will likely continue for a number of decades (unless luck smiles on us and the discovery is made sooner!). Lessons learned from the MAJORANA DEMONSTRATOR, from both software and hardware sides will be used in the next phase of germanium-based $0\nu\beta\beta$ searches. For instance, the enriched germanium detectors will be re-used. Software will be based on the MAJORANA DEMONSTRATOR as well as GERDA implementations, and the lessons learned during the work done for this dissertation are particularly important in making efficient progress towards the collaboration's future goals.

6.1 Hardware

The next generation ^{76}Ge experiment is the LEGEND experiment, planned to be operated in two phases, LEGEND-200 and LEGEND-1000. In the first phase, $\sim$200 kg of ^{76}Ge detectors will be operated at LNGS in the GERDA infrastructure, after a number of upgrades. A reduction of background is expected due to the use of larger detectors, which require fewer cables and components per unit mass, and materials with improved radiopurity, such as MAJORANA DEMONSTRATOR UGEFCu. An energy resolution at least as good as that in MAJORANA DEMONSTRATOR is envisioned. The estimate for the LEGEND-200 background is about a factor of 3 lower than in GERDA, specifically 2×10^{-4} events per keV-kg-yr. LEGEND-200

is expected to achieve a discovery sensitivity on the level of 10^{27} yr in its 5 year planned live time. LEGEND-1000 will feature roughly 1 ton of Ge detectors. Further upgrades in low-radioactivity cables and electronics, as well as the use of UGLAr (UnderGround Liquid Argon), are expected to yield a 20-fold background reduction. LEGEND-1000 is expected to achieve a discovery sensitivity on the level of 10^{28} yr in its 10 year planned live time [16].

6.2 Software

For LEGEND, software development should move towards a more automated, more modern paradigm. While ROOT has been a mainstay of particle (including neutrino) physics for a significant amount of time, it can end up being limiting and black-box-ish. Reinventing the wheel is not necessary, but if one of the shared goals of the collaboration and collaboration members is to improve the latter's software skills, which is mutually beneficial, tools with wider applicability than ROOT should be used. This is indeed already in progress, as DCR has moved away from ROOT to a more python-heavy approach. For LEGEND, the development of pygama [88], using, among others, data-management tools like hdf5 [116] and pandas [103] dataframes, sheds a more modern light on the project. Julia [33], an up-and-coming programming language, will also be heavily used, especially for a parallel analysis by the ex-GERDA part of the collaboration. Overall, a rejection of clunky tools for sleeker and more modern ones should be preferred, when there are no performance issues with the latter.

6.2.1 Paths forward: What could be improved?

As a general rule, improvements are very usually possible. DCR is, of course, no exception. Perhaps the most important thoughts regarding the future of germanium $0\nu\beta\beta$ software, in the opinion of the author, can be summarized below:

- It is best to avoid methods which require parameters to be tuned manually. While barely acceptable for roughly 40 detectors, as in MAJORANA DEMONSTRATOR, such a

method will quickly become unacceptable as a large number of detectors are commissioned for LEGEND.

- We should not use parameters whose values depend on other discriminators (e.g. DCR calculated after the AvsE cut), at least where manual checks are required. This process is not only cumbersome and repetitive, but it leads to huge delays if situations are encountered where each parameter tuning is waiting on another. Furthermore, the background modelling investigation shows that it is perhaps overall better to have a 'purer' form of DCR, tuned without the preliminary application of AvsE, available for use in cases where the use of AvsE is unwarranted or leaves a lot to be desired, due to background model simulation complexity.

- It is also advantageous to use parameters which do not require (re-)tuning if one needs to use a different cut, i.e. in the previous version of dcr, a full retune was needed in order to use the parameter with a different acceptance, so there were lots of different dcrs (dcr90, dcr95, dcr99, etc). When possible, parameters should be constructed such that in order to obtain a different acceptance, one needs to just simply change the cut value.

6.2.2 DCR-specific lessons learned

A recent development has been a short investigation into DCR by Chris Haufe, of the Majorana Demonstrator background modelling group. Using one detector as an example, a parabolic shape in DCR efficiency versus energy was observed (see Fig. 6.1). However, none of the results are unexpected, or problematic, for DCR itself. The overall shape of the DCR efficiency with the AvsE cut applied yields a flat efficiency with energy for most of the energy range, with the exception of a region around 300-500 keV, and around 2500 keV (bin which suffers from low statistics). The efficiency by energy for the same detector after applying the AvsE cut can be seen in Fig. 6.2. The efficiency by energy for the all detectors

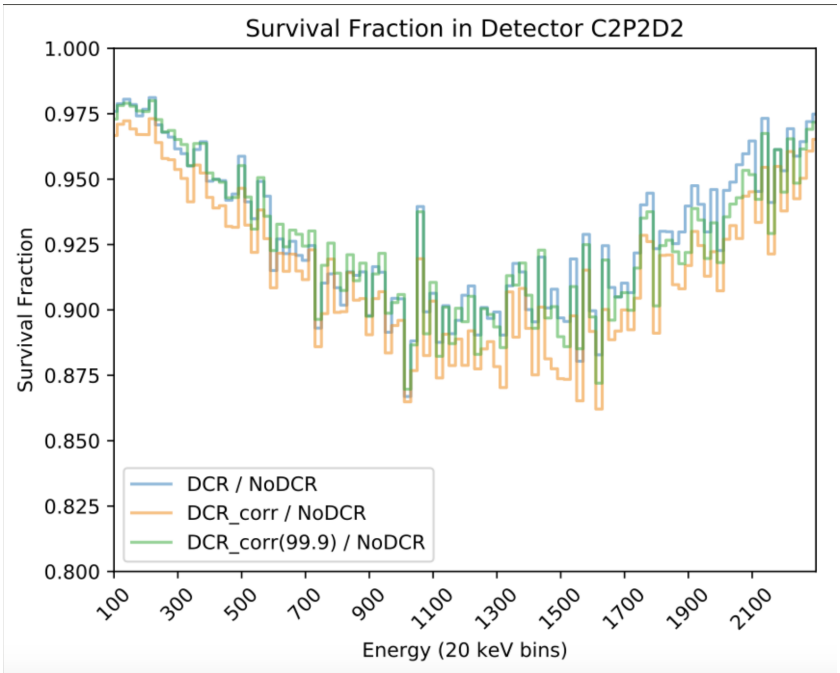

Figure 6.1: DCR acceptance (survival fraction) vs Energy in detector C2P2D2 without AvsE cut applied. Figure courtesy of Chris Haufe.

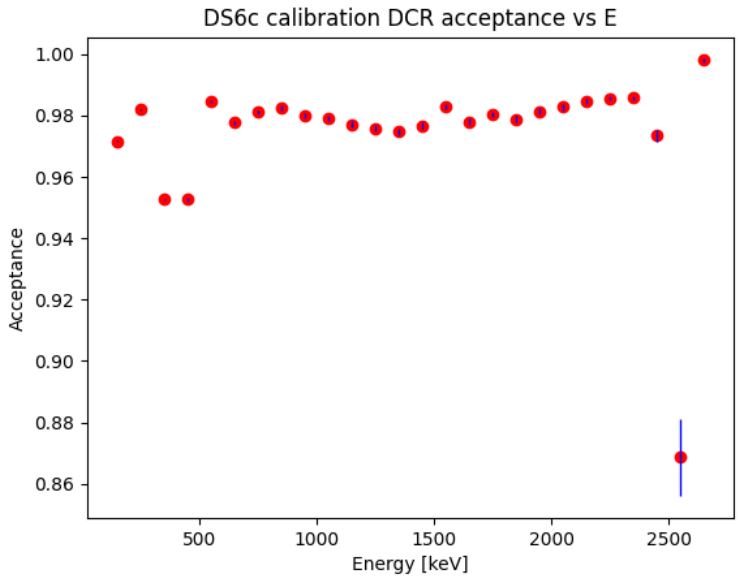

Figure 6.2: DCR acceptance (survival fraction) vs Energy in detector C2P2D2 with the AvsE cut applied.

after applying the AvsE cut can be seen in Fig. 6.3. As a reminder, the main application of DCR was removing alphas around our $Q_{\beta\beta}$ of 2039 keV. This is unaffected by any shape in efficiency – all that's needed to be known is the efficiency around $Q_{\beta\beta}$.

The variation in efficiency by energy is potentially somewhat important for future background modelling studies, as the application of AvsE is not simple – there is only a heuristic determining MC values of AvsE for background simulations, which is known to not be quite accurate. Future investigations could improve the DCR application to background modelling by calculating DCR without applying the AvsE cut, and perhaps adding another anchor around 500 keV to potentially remove the lower efficiency in the 300-500 keV region. However, this is outside the scope of this thesis and is only presented as a possible path forward when developing a similar parameter for LEGEND, or for the MAJORANA Background Modelling group if deemed appropriate.

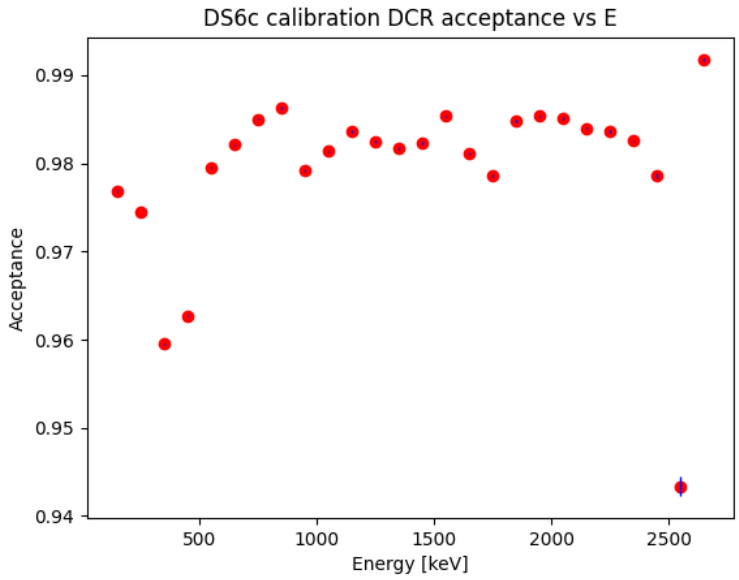

Figure 6.3: DCR acceptance (survival fraction) vs Energy for all detectors with the AvsE cut applied.

It is also worth mentioning that the investigation into the parabolic DCR efficiency is an example as to why a statistics-based derivation of DCR is useful. During this time, an issue that came up was the desire to use a higher DCR cut to do a 99.9% efficiency cut (see Fig. 6.1). While not perfect due to DCR tailing effects (multi-site events with high DCR), it was straight-forward to derive the value of an 'ideal' DCR for a 99.9% cut, $DCR > 3.1$. In the previous version of DCR, each specific acceptance level would require a manual, labor-intensive re-tuning of DCR so that the cut of X% would correspond to a $DCR > 0$ cut.

6.3 Final Words

Overall, the next-generation LEGEND 1000 experiment has a good chance of discovering whether the neutrino is a Majorana particle, i.e. its own anti-particle. If the properties match certain theoretical predictions, the neutrino's Majorana nature might tell us why

there is more matter than anti-matter, and as a consequence why we exist, and not anti-people or, perhaps, nothing at all. A discovery of the Majorana neutrino would also be the first time matter would be created in the laboratory.

Incremental progress, such as presented in this thesis, is always needed in order to reach lofty goals. Methodical improvements, for example in offline analysis cut performance, can lead to big discoveries. The implementation of current, bleeding-edge technology, can be mutually beneficial for both scientific collaborations and the individuals who acquire these skills while working as Ph.D. students. Progress is progress, and hopefully, within a few decades, the effort will have been worth it.

BIBLIOGRAPHY

[1] Multi-site event discrimination for the MAJORANA DEMONSTRATOR data. Technical Report M-TECHDOCUNIDOC-2020-049, Univ. of Washington, 2022. Internal Report, May, 20, 2020.

[2] C. E. Aalseth et al. Search for neutrinoless double-β decay in ^{76}Ge with the majorana demonstrator. *Phys. Rev. Lett.*, 120:132502, Mar 2018.

[3] S. Abe et al. First Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen. 3 2022.

[4] N. Abgrall et al. The MAJORANA DEMONSTRATOR neutrinoless double-beta decay experiment. *Adv. High Energy Phys.*, 2014:365432, Aug 2014.

[5] N. Abgrall et al. The majorana parts tracking database. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 779:52–62, 2015.

[6] N. Abgrall et al. High voltage testing for the majorana demonstrator. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 823:83–90, 2016.

[7] N. Abgrall et al. The majorana demonstrator radioassay program. *Nucl. Instrum. Meth. A*, 828:22–36, 2016.

[8] N. Abgrall et al. The majorana demonstrator calibration system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 872:16–22, 2017.

[9] N. Abgrall et al. Charge trapping correction and energy performance in the MAJORANA DEMONSTRATOR. *Phys Rev C*, 107(4), 2023.

[10] I. Abt, L. Garbini, C. Gooch, S. Irlbeck, X. Liu, M. Palermo, and O. Schulz. Alpha-event and surface characterisation in segmented true-coaxial {HPGe} detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 858:80 – 89, 2017.

[11] D. Q. Adams et al. Improved limit on neutrinoless double-beta decay in ^{130}Te with cuore. *Phys. Rev. Lett.*, 124:122501, Mar 2020.

[12] Ian Affleck and Michael Dine. A New Mechanism for Baryogenesis. *Nucl. Phys. B*, 249:361–380, 1985.

[13] M. Agostini, M. Barnabe-Heider, T. Bode, D. Budjáš, A. Lazzaro, and S. Schönert. Bege detector response to alpha-induced energy depositions on the p+ electrode and groove surfaces. bege detector response to alpha-induced energy depositions on the p+ electrode and groove surfaces. GERDA Scientific/Technical Report GSTR-13-006, Physik Department E15, Technische Universität München, July 2013.

[14] M. Agostini et al. Final results of gerda on the search for neutrinoless double-β decay. *Phys. Rev. Lett.*, 125:252502, Dec 2020.

[15] Matteo Agostini, Giovanni Benato, and Jason A. Detwiler. Discovery probability of next-generation neutrinoless double-β decay experiments. *Phys. Rev. D*, 96:053001, Sep 2017.

[16] Matteo Agostini, Giovanni Benato, Jason A. Detwiler, Javier Menéndez, and Francesco Vissani. Toward the discovery of matter creation with neutrinoless double-beta decay. 2 2022.

[17] D. Hervas Aguilar. Data cleaning: Pile-up rejection in ds6. Technical Report M-TECHDOCUNIDOC-2019-041, 2019. Internal Report, June 13, 2019.

[18] D. Hervas Aguilar. As-built background model. Technical Report M-TECHDOCUNIDOC-2019-044, 2021. Internal Report, February 10, 2021.

[19] M. Aker et al. Direct neutrino-mass measurement with sub-electronvolt sensitivity. *Nature*, 18:160–166, 2022.

[20] V. Alenkov et al. First results from the amore-pilot neutrinoless double beta decay experiment. *Eur. Phys. J. C*, 79:791, Sep 2019.

[21] S. I. Alvis et al. Search for neutrinoless double-β decay in ^{76}Ge with 26 kg yr of exposure from the majorana demonstrator. *Phys. Rev. C*, 100:025501, Aug 2019.

[22] Sebastian Alvis et al. Search for neutrinoless double-β decay in ^{76}ge with 26 kg yr of exposure from the Majorana Demonstrator. *Phys. Rev. C*, 100:025501, 2019.

[23] J. Anderson, R. Brito, D. Doering, T. Hayden, B. Holmes, J. Joseph, H. Yaver, and S. Zimmermann. Data acquisition and trigger system of the gamma ray energy tracking in-beam nuclear array (gretina). *IEEE Transactions on Nuclear Science*, 56(1):258–265, 2009.

[24] G. Anton et al. Search for neutrinoless double-β decay with the complete exo-200 dataset. *Phys. Rev. Lett.*, 123:161802, Oct 2019.

[25] R. Arnold et al. Results of the search for neutrinoless double-β decay in ^{100}Mo with the nemo-3 experiment. *Phys. Rev. D*, 92:072011, Oct 2015.

[26] I. J. Arnquist et al. Final result of the Majorana Demonstrator's search for neutrinoless double-β decay in ^{76}Ge. *Phys. Rev. Letters*, 130:062501, Feb 2023.

[27] J. Arnquist et al. α-event characterization and rejection in point-contact HPGe detectors. *Eur. Phys. J. C*, 82:226, Mar 2022.

[28] O. Azzolini et al. Final result of cupid-0 phase-i in the search for the ^{82}Se neutrinoless double-β decay. *Phys. Rev. Lett.*, 123:032501, Jul 2019.

[29] John N. Bahcall and Raymond Davis Jr. Solar neutrinos: A scientific puzzle. *Science*, 191(4224):264–267, 1976.

[30] J. Barea, J. Kotila, and F. Iachello. $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements in the interacting boson model with isospin restoration. *Phys. Rev. C*, 91:034304, Mar 2015.

[31] P. Barton, P. Luke, M. Amman, Y. Chan, J. Detwiler, J. Loach, R. Martin, A. Poon, C. Tindall, and K. Vetter. Low-noise low-mass front end electronics for low-background physics experiments using germanium detectors. In *Proceedings of the 2011 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC2011)*, pages 1976–1979. IEEE, Oct 2011.

[32] Antoon Theuwis Ben Depuydt and Igor Romandic. Germanium: From the first application of Czochralski crystal growth to large diameter dislocation-free wafers. *Materials Science in Semiconductor Processing*, 9(4-5):437–443, 2006.

[33] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.

[34] Laurie M. Brown. The idea of the neutrino. *Physics Today*, 31(9):23, 1978.

[35] Wilfried Buchmüller. Leptogenesis. *Scholarpedia*, 9:11471, 2014.

[36] W. Bugg, Yu. Efremenko, and S. Vasilyev. Large plastic scintillator panels with wls fiber readout: Optimization of components. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 758:91 – 96, 2014.

[37] Micah Buuck. Calculating the active mass of the MAJORANA module 1 detectors, 2016. Internal Report, May 16, 2016.

[38] Thomas Caldwell et al. Estimated limit on $T^{0\nu}_{1/2}$ from the MAJORANA DEMONSTRATOR using ds0-6. Technical Report M-TECHDOCUNIDOC-2018-031, Univ. of North Carolina, 2018. Internal Report, June 19, 2018.

[39] J. Chadwick. The existence of a neutron. *Proceedings of the Royal Society A*, 1932.

[40] Pinghan Chu. Unidoc of pulser tag of the mjd. Technical Report M-TECHDOCUNIDOC-2017-025, Los Alamos National Laboratory, 2017.

[41] F. Reines et al. C.L Cowan Jr. Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956.

[42] James M. Cline. Baryogenesis. In *Les Houches Summer School - Session 86: Particle Physics and Cosmology: The Fabric of Spacetime*, 9 2006.

[43] Alain Coc and Elisabeth Vangioni. Primordial nucleosynthesis. *Int. J. Mod. Phys. E*, 26(08):1741002, 2017.

[44] GALLEX Collaboration. Solar neutrinos observed by gallex at gran sasso. *Physics Letters B*, 285(4):376–389, 1992.

[45] SAGE Collaboration. Measurement of the solar neutrino capture rate with gallium metal. *Physical Review C*, 60, 1999.

[46] SNO Collaboration. Measurement of the rate of $\nu_e + d \rightarrow \ p + p + e^-$ interactions produced by 8b solar neutrinos at the sudbury neutrino observatory. *Physical Review Letters*, 87(7), 2001.

[47] Super-Kamiokande Collaboration. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81:1562–1567, 1998.

[48] L. Coraggio, A. Gargano, N. Itaco, R. Mancino, and F. Nowacki. Calculation of the neutrinoless double-β decay matrix element within the realistic shell model. *Phys. Rev. C*, 101:044315, Apr 2020.

[49] Clara Cuesta. Multi-site event discrimination for the majorana demonstrator. Technical Report M-TECHDOCUNIDOC-2017-014, CIEMAT, 2017. Internal Report, Oct. 18, 2017.

[50] et al. David Edwins. Livetime and exposure for majorana demonstrator data sets 0-8. Technical Report M-TECHDOCUNIDOC-2018-032, Univ. of South Carolina, 2022. Internal Report, Feb., 2022.

[51] Frank F. Deppisch, Lukas Graf, Francesco Iachello, and Jenni Kotila. Analysis of light neutrino exchange and short-range mechanisms in $0\nu\beta\beta$ decay. *Phys. Rev. D*, 102:095016, Nov 2020.

[52] J. Detwiler. Updated isotopic abundances and molecular weights of the majorana demonstrator enriched detectors. Technical Report M-TECHDOCUNIDOC-2022-068, Univ. of Washington, 2022. Internal Report, May 2022.

[53] Michelle J. Dolinski, Alan W.P. Poon, and Werner Rodejohann. Neutrinoless double-beta decay: Status and prospects. *Ann. Rev. Nucl. Part. Sci.*, 69(1):219–251, 2019.

[54] Steve Elliott. Blindness in the MAJORANA DEMONSTRATOR data analysis - data prescaling. Technical Report M-TECHUNIDOC-2012-064, Los Alamos National Laboratory, 2012.

[55] C.D. Ellis and N.F. Mott. Energy relations in the β-ray type of radioactive disintegration. *Proceedings of the Royal Society A*, 1933.

[56] Jonathan Engel and Javier Menéndez. Status and future of nuclear matrix elements for neutrinoless double-beta decay: a review. *Rep. Prog. Phys.*, 80(4):046301, mar 2017.

[57] C. S. Wu et al. Experimental Test of Parity Conservation in Beta Decay. *Phys. Rev*, 105:1413–1414, 1957.

[58] E. Fiorini et al. A search for lepton non-conservation in double beta decay with a germanium detector. *Physics Letters B*, 25:602–603, 1967.

[59] E.W. Hoppe et al. Reduction of radioactive backgrounds in electroformed copper for ultra-sensitive radiation detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 764:116–121, 2014.

[60] Ian Guinn et al. Updated limit on $t^{0\nu}_{1/2}$ from ds0-8 of the MAJORANA DEMONSTRATOR. Technical Report M-TECHDOCUNIDOC-2022-067, MAJORANA DEMONSTRATOR, 2022. Internal Report, June 14, 2022.

[61] Jürgen Evers et al. Czochralski's Creative Mistake: A Milestone on the Way to the Gigabit Era. *Angew. Chem. Int. Ed.*, 42:5684–5698, 2003.

[62] K. Vetter et al. Performance of the greta prototype detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 452(1):105 – 114, 2000.

[63] N.Abgrall et al. The Processing of Enriched Germanium for the Majorana Demonstrator and R&D for a next Generation Double-Beta Decay Experiment. *NIM A*, 877:314–322, 2018.

[64] S. Fukuda et al. Solar b-8 and hep neutrino measurements from 1258 days of superkamiokande data. *Physical Review Letters*, 86:5651–5655, 2001.

[65] Tommaso Comellato et al. Charge-Carrier Collective Motion in Germanium Detectors for $\beta\beta$-Decay Searches. *The European Physical Journal C*, 81(1):76, 2021.

[66] Steve Elliott *et al.* First estimate of the limit on $T^{0\nu}_{1/2}$ from the MAJORANA DEMONSTRATOR. Technical Report M-TECHDOCUNIDOC-2017-022, Los Alamos National Laboratory, 2017.

[67] C.E. Aaleseth *et al.* (Majorana Collaboration). Search for neutrinoless double-β decay in ^{76}Ge with the MAJORANA DEMONSTRATOR. *Phys. Rev. Lett.*, 120:132502, 2018.

[68] James F. Minuit: Function minimization and error analysis reference manual. Technical Report CERNLIB-D506, CERN, 2017.

[69] Dong-Liang Fang, Amand Faessler, and Fedor Šimkovic. $0\nu\beta\beta$-decay nuclear matrix element for light and heavy neutrino mass mechanisms from deformed quasiparticle random-phase approximation calculations for ^{76}Ge,^{82} Se,^{130} Te,^{136} Xe, and ^{150}Nd with isospin restoration. *Phys. Rev. C*, 97:045503, Apr 2018.

[70] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, Apr 1998.

[71] Enrico Fermi. Tentativo di una teoria dell'emissione dei raggi beta. *Ric. Sci.*, 4:491–495, 1933.

[72] Enrico Fermi. Tentativo di una teoria dei raggi β. *Il Nuovo Cimento*, 11(1), 1934.

[73] M. Fukugita and T. Yanagida. Baryogenesis Without Grand Unification. *Phys. Lett. B*, 174:45–47, 1986.

[74] Sunyar A. W. Goldhaber M., Grodzins L. Helicity of Neutrinos. *Phys. Rev*, 109:1015–1017, 1958.

[75] Julieta Gruszko. Surface Alpha Interactions in P-Type Point-Contact HPGe Detectors: Maximizing Sensitivity of 76Ge Neutrinoless Double-Beta Decay Searches, 2017.

[76] V.E. Guiseppe, J.M. Lopez-Castano, D.C. Radford, R.L. Varner, and C.-H. Yu. Ornl analysis. Technical Report M-TECHDOCUNIDOC-2022-065, 2022. Internal Report, Feb. 16, 2022.

[77] W.C. Haxton. Detection of the free neutrino: a confirmation. *Ann. Rev. Astron. Astrophys.*, 33:459–503, 1995.

[78] Mihai Horoi and Andrei Neacsu. Shell model predictions for ^{124}Sn double-β decay. *Phys. Rev. C*, 93:024308, Feb 2016.

[79] A. Hostiuc and J.Detwiler. Delayed charge recovery tagging. Technical Report M-TECHDOCUNIDOC-2020-048, Univ. of Washington, 2022. internal Report, Mar. 22, 2022.

[80] Juhani Hyvärinen and Jouni Suhonen. Nuclear matrix elements for $0\nu\beta\beta$ decays with light or heavy majorana-neutrino exchange. *Phys. Rev. C*, 91:024613, Feb 2015.

[81] Fedor Šimkovic, Adam Smetana, and Petr Vogel. $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements evaluated in closure approximation, neutrino potentials and su(4) symmetry. *Phys. Rev. C*, 98:064325, Dec 2018.

[82] Jordan Myslik Inwook Kim and Tom Caldwell. Data cleaning in ds0-ds8 for 2020-2022 $0\nu\beta\beta$-decay analysis. Technical Report M-TECHUNIDOC-2020-050, 2022. Internal Report, Feb. 28, 2022.

[83] U.D. Jentschura and B.J. Wundt. Neutrino helicity reversal and fundamental symmetries. *J. Phys. G*, 41:075201, 2014.

[84] Glenn F. Knoll. Radiation detection and measurement, 3rd ed. *John Wiley & Sons*, 2000.

[85] T. Lannen. Exposure recovery. Technical Report M-TECHUNIDOC-2022-061, 2022. Internal Report, Jan. 27, 2022.

[86] Tom Lannen, Inwook Kim, et al. Run selection and data cleaning of ds0-8. Technical Report M-TECHDOCUNIDOC-2022-062, Univ. of North Carolina, 2022.

[87] T. D. Lee and C. N. Yang. Question of Parity Conservation in Weak Interactions. *Phys. Rev*, 104:254–258, 1956.

[88] LEGEND Collaboration. pygama: Python package for data processing and analysis , 2023.

[89] W.A. Rolke A. M. López and J. Conrad. Limits and confidence intervals in the presence of nuisance parameters. *NIM A*, 551(493), 2005.

[90] J.M. Lopez-Castano and Ping-Han Chu. Unidoc of energy performance of majorana demonstrator. Technical Report M-TECHDOCUNIDOC-2017-012, Los Alamos National Laboratory, 2021. Internal Report, Sep. 30, 2021.

[91] J.M. Lopez-Castano and Ping-Han Chu. Energy systematic of majorana demonstrator. Technical Report M-TECHDOCUNIDOC-2017-021, Los Alamos National Laboratory, 2022. Internal Report, Feb. 9, 2022.

[92] Ettore Majorana. Teoria simmetrica dell'elettrone e del positrone. *Il Nuovo Cimento (1924-1942)*, 14(4):171, Sep 2008.

[93] J Menéndez. Neutrinoless$\beta\beta$decay mediated by the exchange of light and heavy neutrinos: the role of nuclear structure correlations. *J. Phys. G: Nucl. Part. Phys.*, 45(1):014003, dec 2017.

[94] M. Mirea, T. Pahomi, and S. Stoica. Values of the phase space factors involved in double beta decay. *Rom. Rep. Phys.*, 67:872–889, Apr 2015.

[95] L. Moneta et al. The roostats project. *PoS ACAT*, 57:2010, 2010.

[96] Brianna Mount et al. Double beta decay q values of ^{74}se and ^{76}ge. *Phys. Rev. C.*, 81:032501, 2010.

[97] P. Mullowney, M.-C. Lin, K. Paul, C. Ahrens, M. Loh, R. J. Cooper, D. C. Radford, and E. Hull. Computational models of germanium point contact detectors. *Nuclear Instruments and Methods in Physics Research A*, 662:33–44, January 2012.

[98] M. T. Mustonen and J. Engel. Large-scale calculations of the double-β decay of ^{76}**Ge**, ^{130}**Te**, ^{136}**Xe**, and ^{150}**Nd** in the deformed self-consistent skyrme quasiparticle random-phase approximation. *Phys. Rev. C*, 87:064302, Jun 2013.

[99] Jordan Myslik and Tom Caldwell. Data cleaning in ds0-ds5 for $0\nu\beta\beta$-decay analysis. Technical Report M-TECHDOCUNIDOC-2017-018, Lawrence Berkeley National Laboratory, 2017.

[100] Emmy Noether. Invariant Variation Problems. *Gott. Nachr.*, 1918:235–257, 1918.

[101] Milton Ohring. 7 - mechanical behavior of solids. In Milton Ohring, editor, *Engineering Materials Science*, pages 299–V. Academic Press, San Diego, 1995.

[102] Gulden Othman. The CAGE Scanner: Investigating Surface Backgrounds in High-Purity Germanium Detectors, 2021.

[103] The pandas development team. pandas-dev/pandas: Pandas, February 2020.

[104] Wolfgang Pauli. Aufsätze und vorträge über physik und erkenntnisthistorie (translation of). *Collected Scientific Papers*, 2:1313, 1964.

[105] B. Pontecorvo. Neutrino experiments and the problem of conservation of leptonic charge. *Zh. Eksp. Teor. Fiz.*, 53:1717–1725, 1968.

[106] D. Radford, I. Y. Lee, and K. Lagergren. Siggen software package, 2017.

[107] Simon Ramo. Currents Induced by Electron Motion. *Proceedings of the IRE*, 27:584–5, 1939.

[108] Reines and Cowan. The Reines-Cowan Experiments. Detecting the Poltergeist. *Los Alamos Science*, 25:1–6, 1997.

[109] Nicholas Ruof. Natural Radioactive Background Suppression in 76Ge Neutrinoless Double-Beta Decay Experiments, 2022.

[110] Ernest Rutherford. The scattering of α and β particles by matter and the structure of the atom. *Philosophical Magazine*, 6(21), 1911.

[111] Z. Maki; M. Nakagawa; S. Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 28(5):870, 1962.

[112] A.D. Sakharov. Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe. *Sov. Phys. Usp.*, 34(5):392–393, 1991.

[113] J. Schechter and J.W.F. Valle. Neutrinoless Double beta Decay in SU(2) x U(1) Theories. *Phys. Rev. D*, 25:2951, 1982.

[114] Gregory Schott. RooStats for Searches. In *PHYSTAT 2011*, pages 199–208, Geneva, 2011. CERN.

[115] L. S. Song, J. M. Yao, P. Ring, and J. Meng. Nuclear matrix element of neutrinoless double-β decay: Relativity and short-range correlations. *Phys. Rev. C*, 95:024305, Feb 2017.

[116] The HDF Group. Hierarchical data format version 5, 2000-2023.

[117] J. J. Thomson. On bodies smaller than atoms. *The Popular Science Monthly*, 1901.

[118] Nuria López Vaquero, Tomás R. Rodríguez, and J. Luis Egido. Shape and pairing fluctuation effects on neutrinoless double beta decay nuclear matrix elements. *Phys. Rev. Lett.*, 111:142501, Sep 2013.

[119] Clint Wiseman. Ln fill tagging for the MAJORANA DEMONSTRATOR. Technical report, Univ. of South Carolina, 2016. internal report June 22, 2016.

[120] Clint Wiseman. Muon veto analysis for the MAJORANA DEMONSTRATOR. Technical report, Univ. of South Carolina, 2016.

[121] Clint Wiseman. Automatic muon veto analysis for the majorana demonstrator. Technical Report M-TECHDOCUNIDOC-2021-056, University of Washington, 2021.

[122] Wenqin Xu. A short note on the estimation of random coincidence rate in ds1 and expectation in ds0. Technical report, Univ. of North Carolina, 2016. internal Report, June 25, 2016.

[123] Wenqin Xu et al. Statistical analysis for setting limit on the $0\nu\beta\beta$ decay for the majorana demonstrator. Technical report, Univ. of South Dakota, 2017. Internal Report, Oct. 20, 2017.

Appendix A
DCR TECHNICAL DETAILS

An open question remained whether to use the ROI or DEP efficiency as the central value. A quick study was done. Fig. A.1 was produced in order to help decide which of the efficiencies to use. A downward trend in DEP efficiency is clear. It seems unnatural to report anything other than the ROI efficiency.

A.1 DCR Channel Selection Details

This section is presented for reference; work done in collaboration with Jason Detwiler.

Next, plots were generated to track the values of the parameters as a function of run number. An example set of plots for one channel is shown in Fig. A.18. In these plots a slightly modified notation is used: α is called μ_m, β is σ_m, γ is σ_b, t_{drift} is dt, m is $DCR:dt$ slope, and μ_c and σ_c are μ_{corr} and σ_{corr}, respectively. Also drawn is $\sigma_{tot} = \beta\, Q_{\beta\beta} + \gamma$. Several parameters show level changes at run boundaries when configurations were changed and detectors were re-biased. One can also see instances where large shifts or spikes are present. These indicate instabilities which may require channel selection.

To hone in on true parameter shifts, the change in each parameter was computed from period to period. Level shifts that followed a dataset boundary were ignored. Parameters following a period that was already selected out for having invalid parameters as described above were compared with the last valid period, if available. An example plot showing evaluated shifts is shown in Fig. A.19. Periods immediately preceding spikes in these plots are candidates for channel selection.

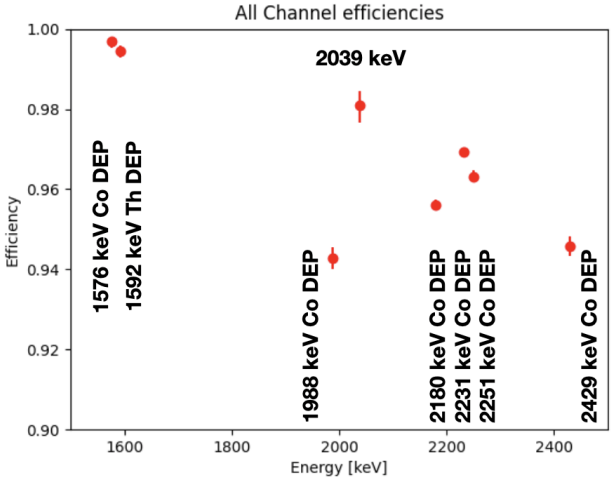

Figure A.1: DS6c and Cobalt data used in this plot. DEP efficiencies trend down with energy. This is evidence of an energy dependence in DCR acceptance. It is also evidence that the correct efficiency to use is that at 2039 keV. As a note, these are all statistical only uncertainties. Systematics are on the order of 0.8%, as can be seen in Table 4.2. This was intended as a quick study and, unlike data in Table 4.2, does not include exposure weighting, hence the differences.

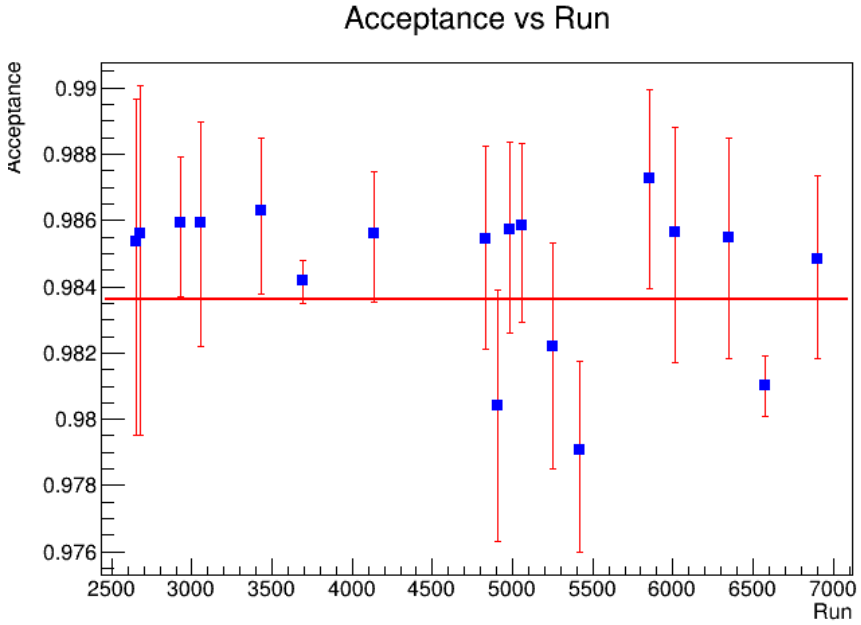

Figure A.2: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS0 for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS0.

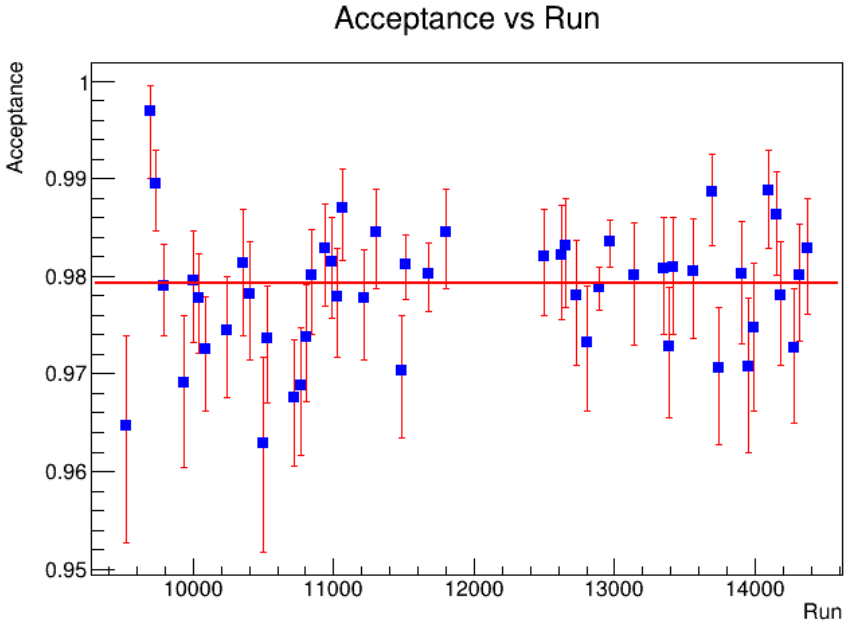

Figure A.3: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS1 for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS1.

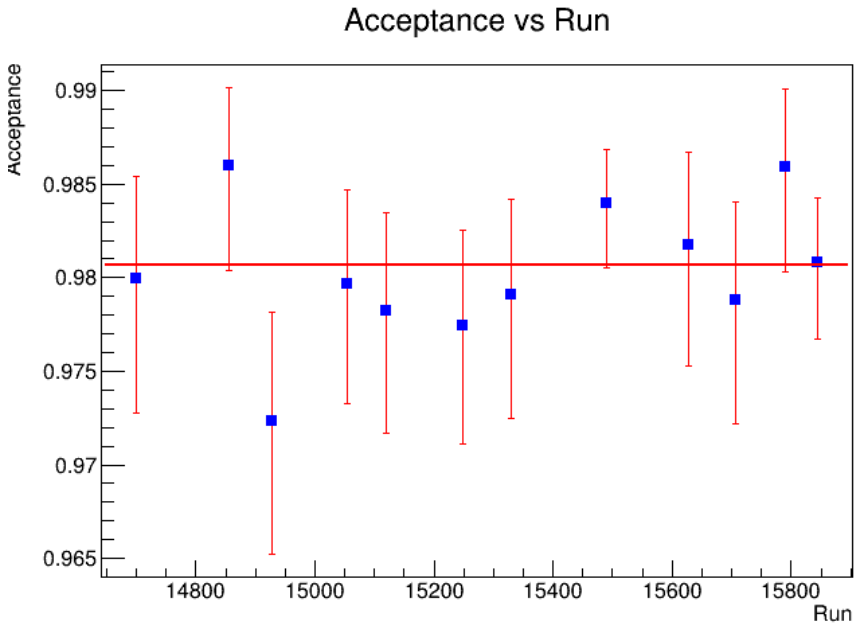

Figure A.4: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS2 for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS2.

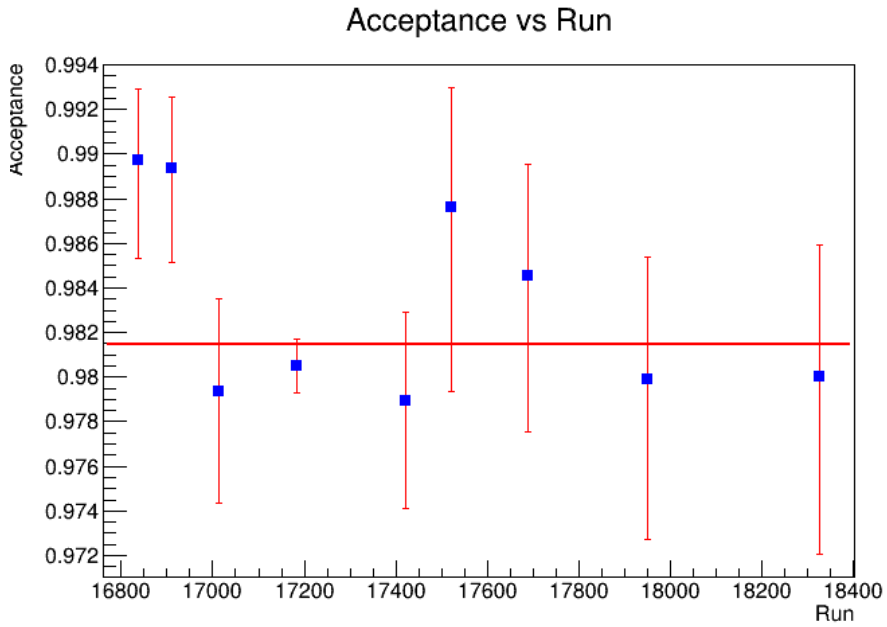

Figure A.5: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS3 for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS3.

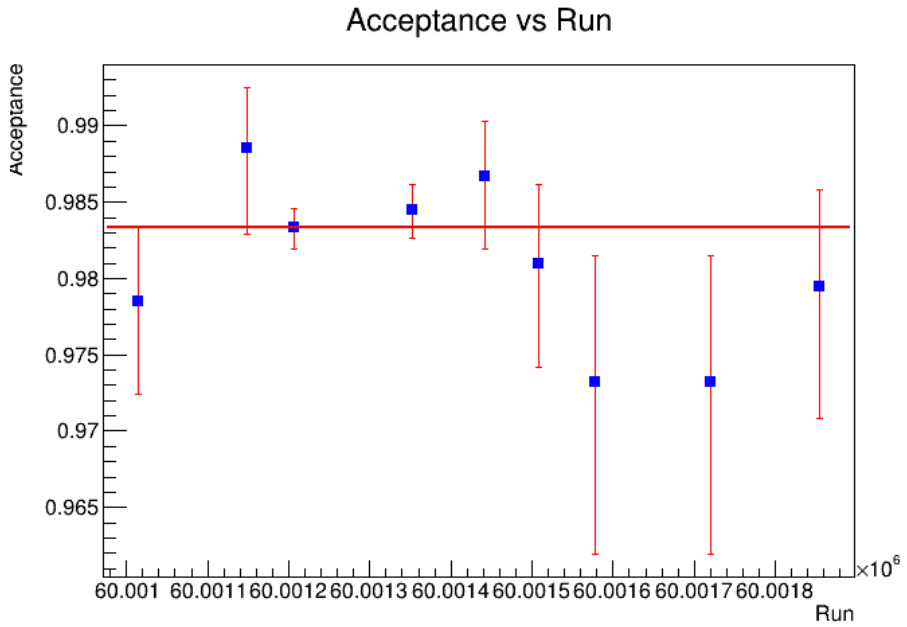

Figure A.6: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS4 for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS4.

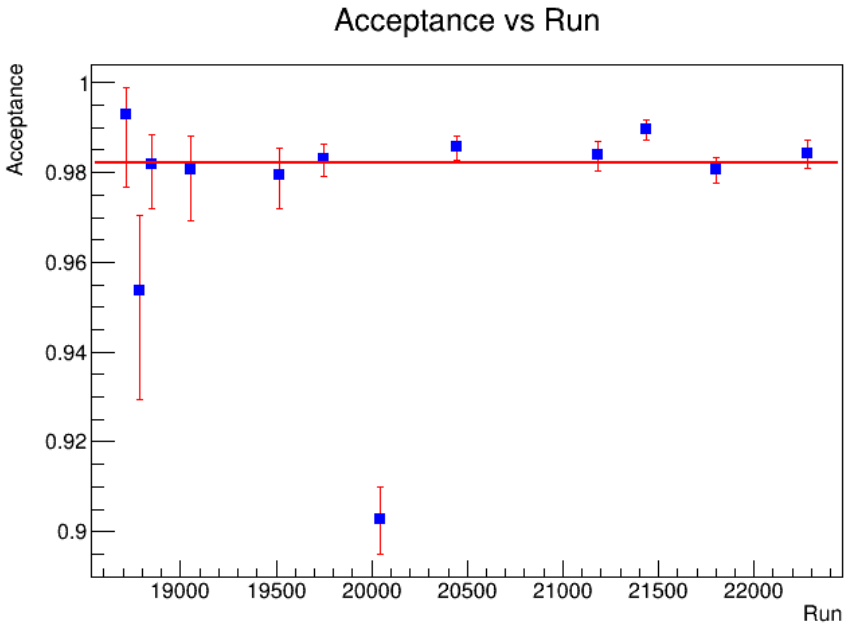

Figure A.7: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS5a for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS5a. This plot will be replaced with the version with natural channels removed for the the calibrations with very low acceptance once the new CS is implemented and the new skim files are produced.

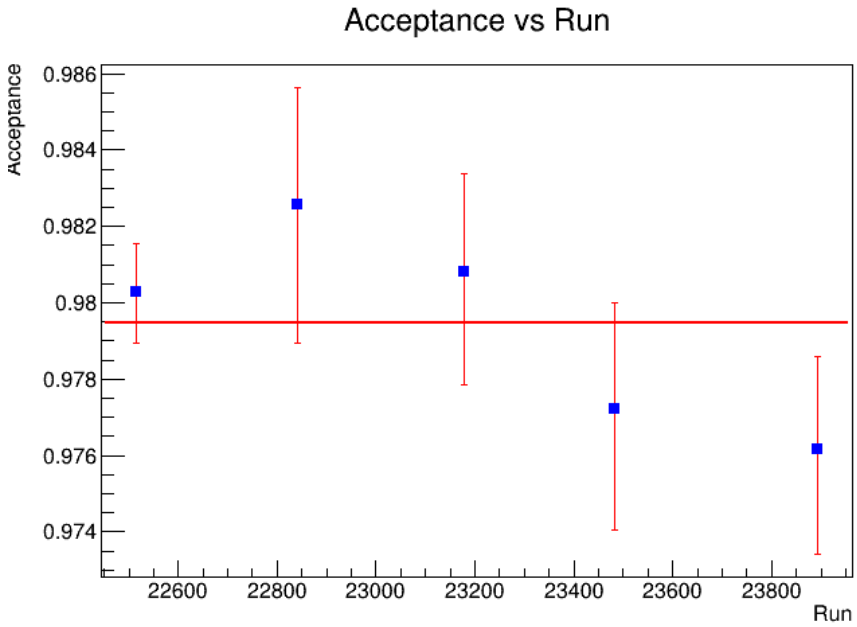

Figure A.8: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS5b for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS5b.

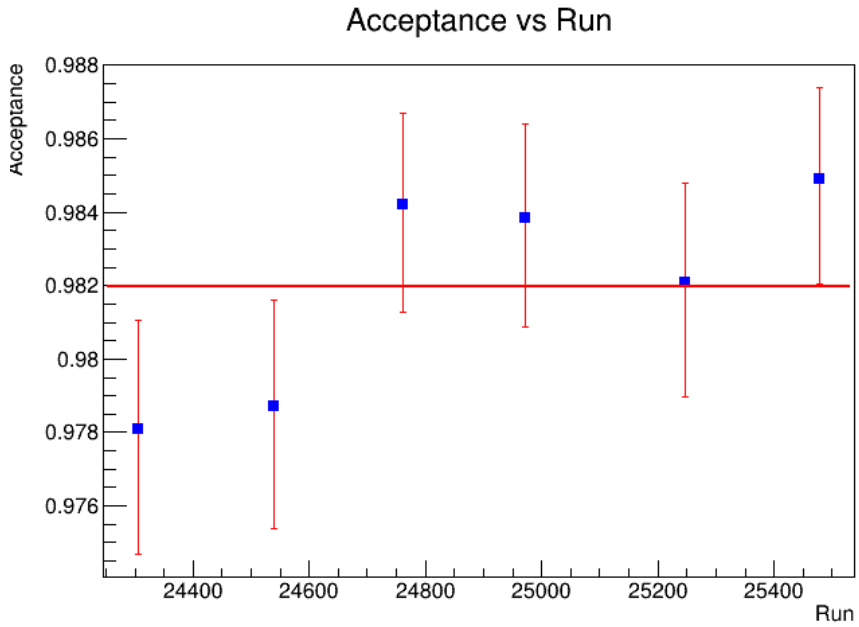

Figure A.9: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS5c for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS5c.

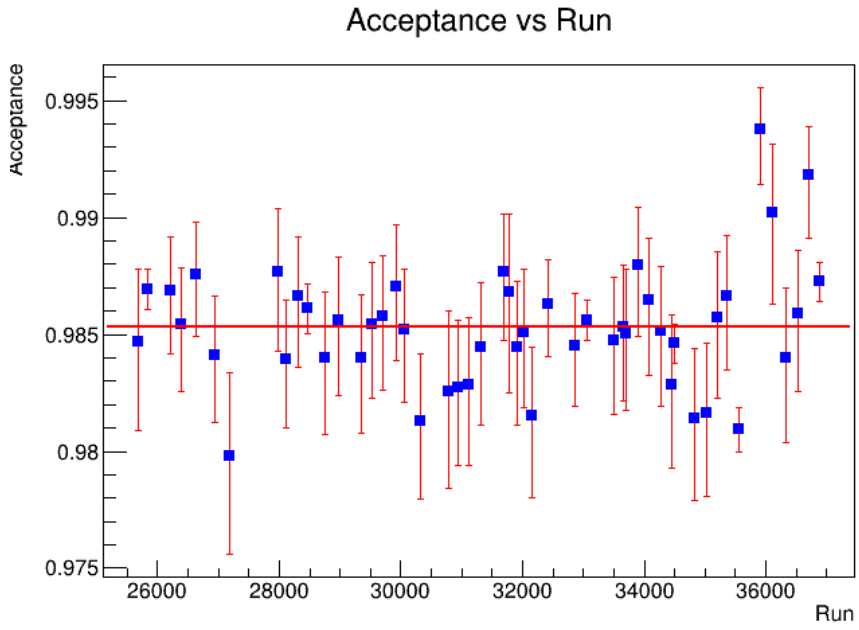

Figure A.10: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS6a for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS6a. The high-acceptance calibration starting at 35915 and the period following it will be channel-selected out. The plot will be updated after this is fully implemented.

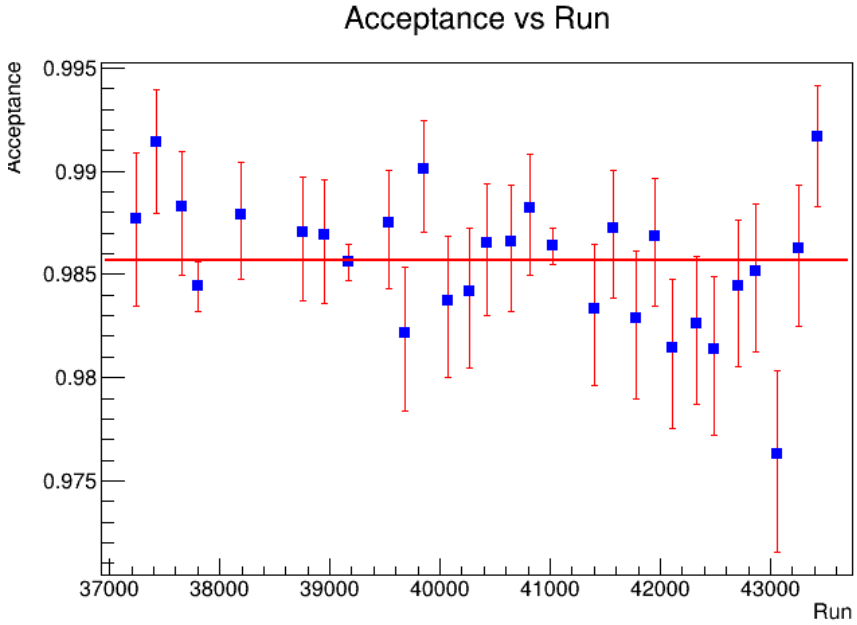

Figure A.11: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS6b for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS6b.

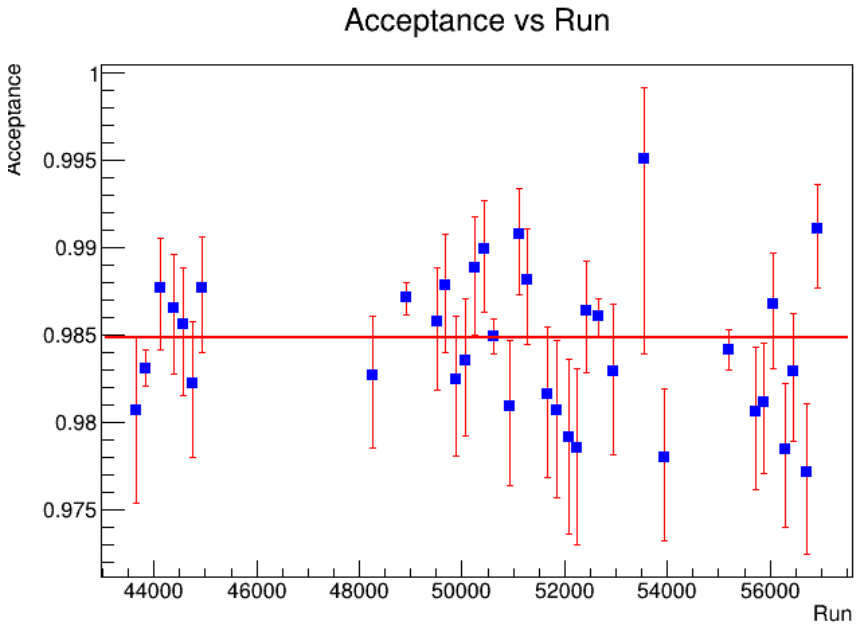

Figure A.12: The weighted-by-calibration acceptance for a dcr_{corr} cut of 2.326 in dataset DS6c for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the average across DS6c.

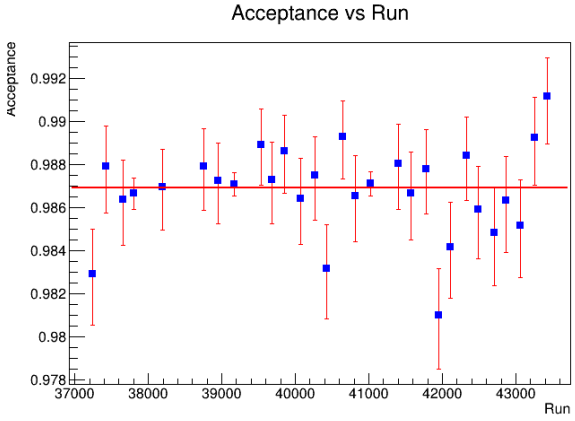

Figure A.13: The weighted acceptance for a dcr_{corr} cut of 2.326 in dataset DS6b for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the weighted average across DS6b.

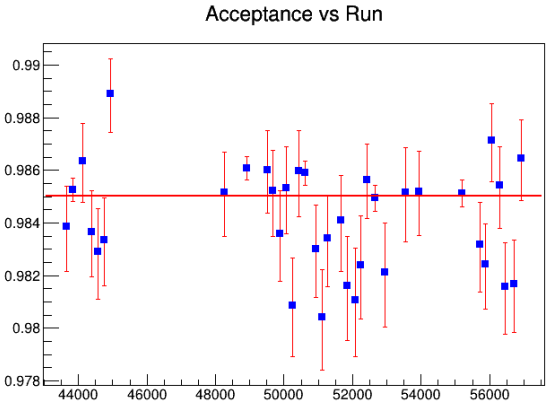

Figure A.14: The weighted acceptance for a dcr_{corr} cut of 2.326 in dataset DS6c for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the weighted average across DS6c.

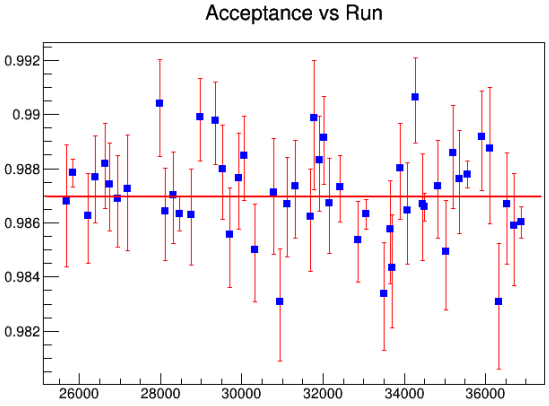

Figure A.15: The weighted acceptance for a dcr_{corr} cut of 2.326 in dataset DS6a for all high-gain channels, for each calibration. Details of the efficiency and its uncertainties can be found in Table 4.2. The flat red line represents the weighted average across DS6a.

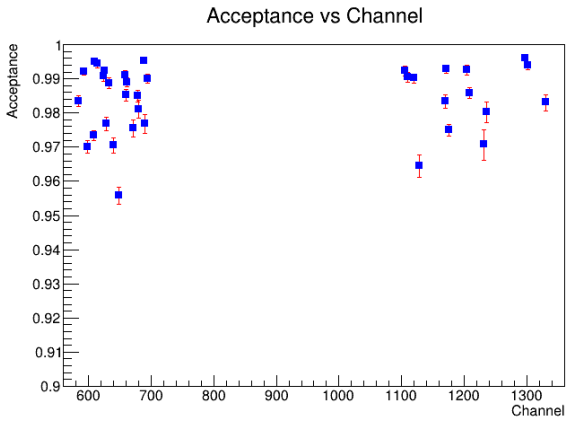

Figure A.16: The acceptance for a dcr cut of 2.326 in dataset DS6b, for each high-gain channel.

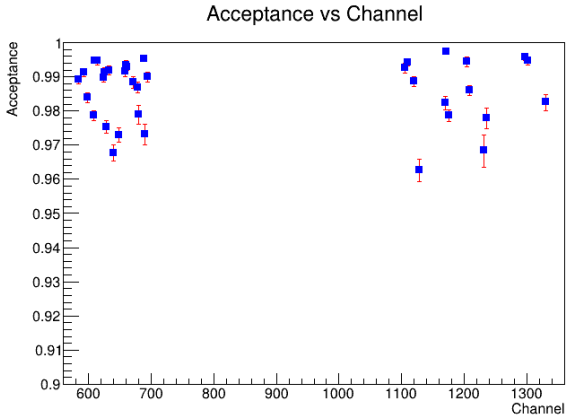

Figure A.17: The acceptance for a dcr_{corr} cut of 2.326 in dataset DS6b, for each high-gain channel.

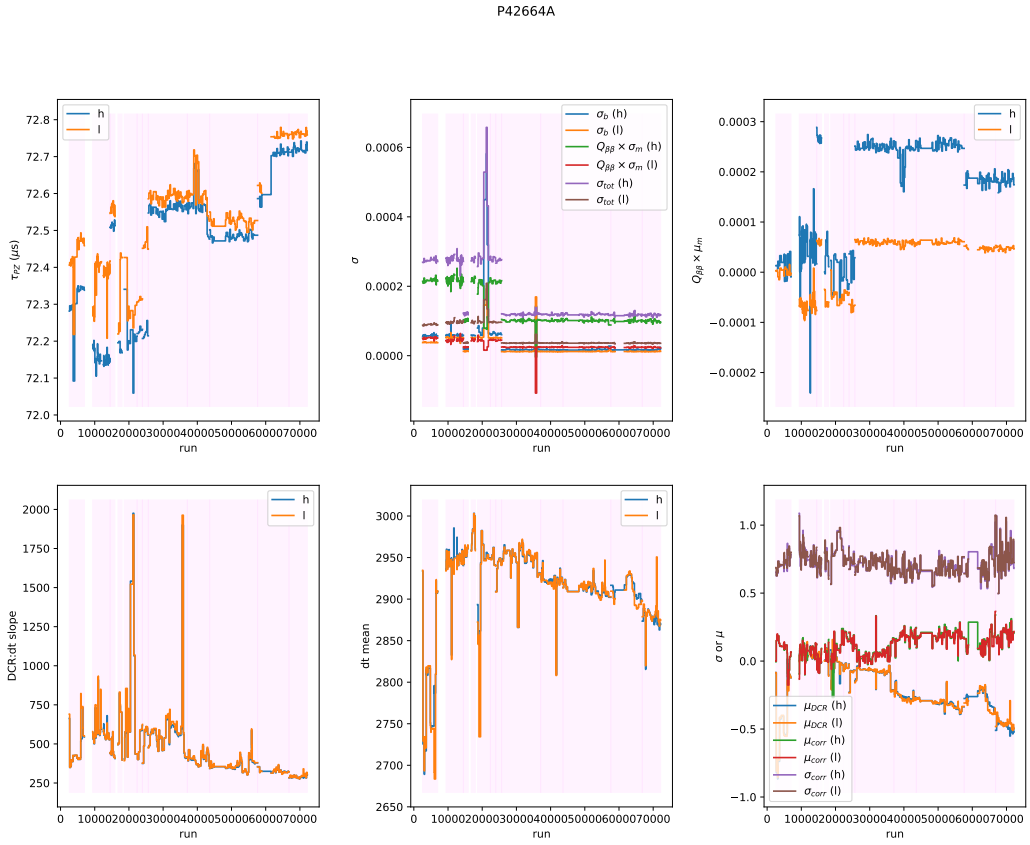

Figure A.18: DCR parameters vs. run for high and low gain channels for detector P42664A. Pink shaded regions denote dataset periods. See the text for parameter descriptions. Plot courtesy of Jason Detwiler [79].

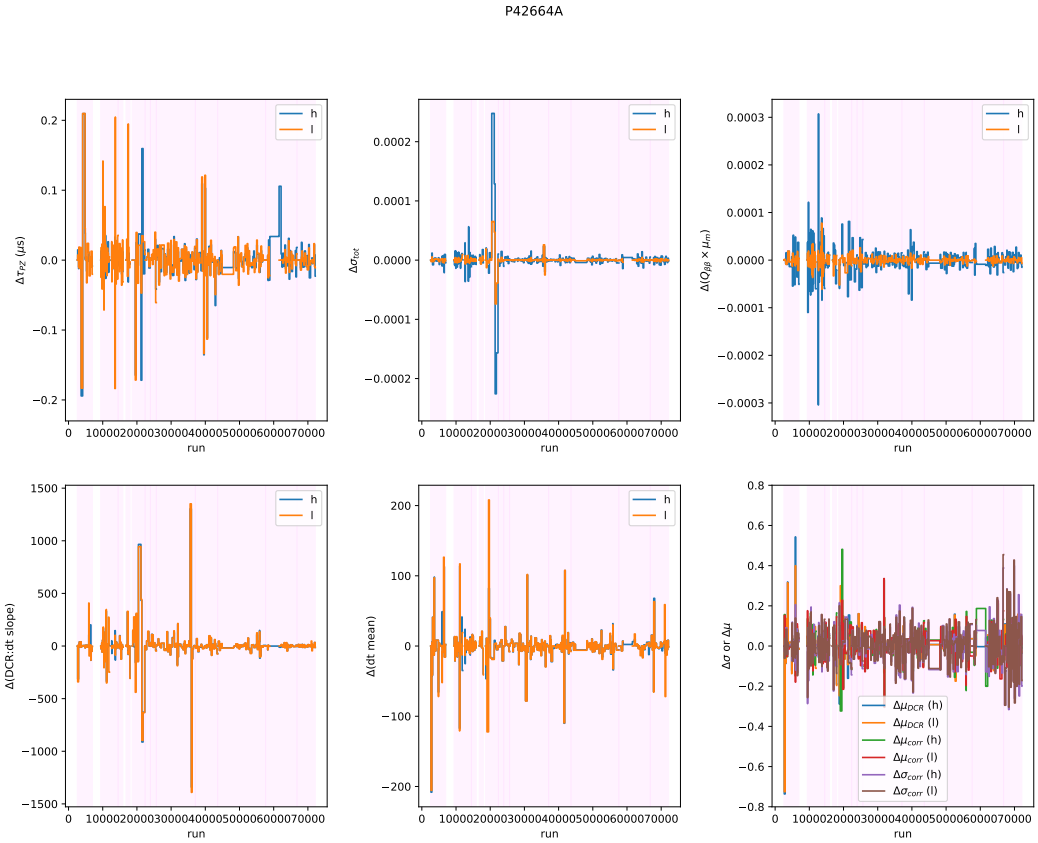

Figure A.19: DCR parameters shifts vs. run for high and low gain channels for detector P42664A (evaluated from Fig. A.18). Plot courtesy of Jason Detwiler [79].

DCR Shifts

When any of the DCR parameters shift, what matters is how the value of $\mathrm{DCR_{corr}}$ corresponding to the cut position $\mathrm{DCR_{cut}}$ changes over time near $E \approx Q_{\beta\beta}$. Call that value D_c: if D_c increases, i.e. either the distribution shifts upwards, widens, or both, applying the cut $\mathrm{DCR_{corr}} < \mathrm{DCR_{cut}}$ results in a loss of efficiency. Conversely, if that D_c decreases, the cut lets in more background events.

The full expression for D_c at $Q_{\beta\beta}$ is:

$$D_c = \left(\frac{\delta_c - \alpha\, Q_{\beta\beta}}{\beta\, Q_{\beta\beta} + \gamma} - \mu_{\mathrm{DCR}} - \frac{\langle t_{\mathrm{drift}} - \mu_{t_{\mathrm{drift}}} \rangle}{m} - \mu_c \right) / \sigma_c \tag{A.1}$$

The value of D_c corresponds to a particular "raw" DCR value δ_c. Since on average $t_{\mathrm{drift}} = \mu_{t_{\mathrm{drift}}}$, δ_c is given by

$$\delta_c = (2.326\sigma_c + \mu_{\mathrm{DCR}} + \mu_c)(\beta\, Q_{\beta\beta} + \gamma) + \alpha\, Q_{\beta\beta} \tag{A.2}$$

where we have set $D_c = 2.326$ corresponding to a 99% cut efficiency in retaining signal for a Gaussian distribution. For a particular parameter p that shifts by a value Δp, the corresponding shift $\Delta D_c(\Delta p)$ can be computed from Eq. A.1. We do not make the approximation $\Delta p \ll p$. For shifts in τ that lead to shifts in δ, we use Eq. 4.4 with $E_{\mathrm{uncal}} = Q_{\beta\beta}/0.4$ for high-gain channels, and $E_{\mathrm{uncal}} = Q_{\beta\beta}/1.2$ for low-gain channels (these approximations should be close enough to compute $\Delta\delta$ within 10% accuracy, which is good enough for this study). We thus compute the following shifts:

$$\Delta D_c(\Delta\tau) = -\frac{\Delta\tau}{\tau^2} Q_{\beta\beta,\,\mathrm{uncal}} / (\beta\, Q_{\beta\beta} + \gamma)\, \sigma_c \tag{A.3}$$

$$\Delta D_c(\Delta\alpha) = \Delta\alpha\, Q_{\beta\beta,\,\mathrm{uncal}} / (\beta\, Q_{\beta\beta} + \gamma)\, \sigma_c \tag{A.4}$$

$$\Delta D_c(\Delta\beta, \Delta\gamma) = -\frac{\delta_c - \alpha\, Q_{\beta\beta}}{\sigma_c} \left(\frac{1}{(\beta + \Delta\beta)\, Q_{\beta\beta} + (\gamma + \Delta\gamma)} - \frac{1}{\beta\, Q_{\beta\beta} + \gamma} \right) \tag{A.5}$$

$$\Delta D_c(\Delta\mu_{\mathrm{DCR}}, \Delta\mu_{t_{\mathrm{drift}}}, \Delta m) = \left[\Delta\mu_{\mathrm{DCR}} - \Delta\mu_{t_{\mathrm{drift}}} / (m + \Delta m) \right] / \sigma_c \tag{A.6}$$

$$\Delta D_c(\Delta\mu_c) = \Delta\mu_c / \sigma_c \tag{A.7}$$

$$\Delta D_c(\Delta\sigma_c) = [(\delta_c - \alpha\, Q_{\beta\beta}) / (\beta\, Q_{\beta\beta} + \gamma) - \mu_{\mathrm{DCR}} - \mu_c] / (\sigma_c + \Delta\sigma_c) - 2.326 \tag{A.8}$$

These shifts are plotted for the high gain channel of detector P42664A in Fig. A.20. In this plot, D_c is labeled DCR_c, $\Delta\alpha$ is referred to as $\Delta\mu(Q_{\beta\beta})$, the combined shift for $\Delta\beta$ and $\Delta\gamma$ is referred to as $\Delta\sigma(Q_{\beta\beta})$, the shifts in the drift time correction parameters are referred to as Δdtc, and μ_c and σ_c are referred to as μ_{corr} and σ_{corr}, respectively.

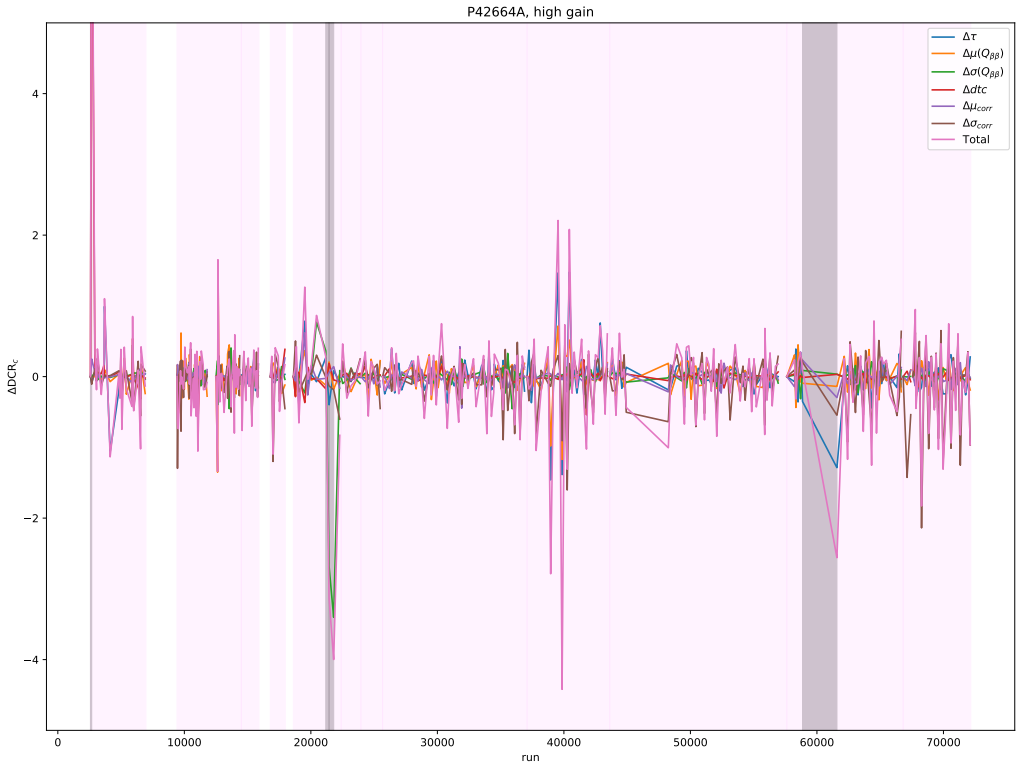

Figure A.20: D_c shifts (ΔDCR_c, same as ΔD_c in the text) vs. run for the high gain channel of detector P42664A (evaluated from Fig. A.19). Regions selected out for having large DCR shifts are demarcated by the grey bands. The spikes in the middle range that exceed the instability threshold are due to the use of disjoint parameter sets and are thus not flagged for instability. The large spike at the start of the window follows the very first data subset in DS0, and is the reason that that first data subset is selected out for this channel. Plot courtesy of Jason Detwiler [79].

Appendix B

DCR ANALYSIS PLOTS

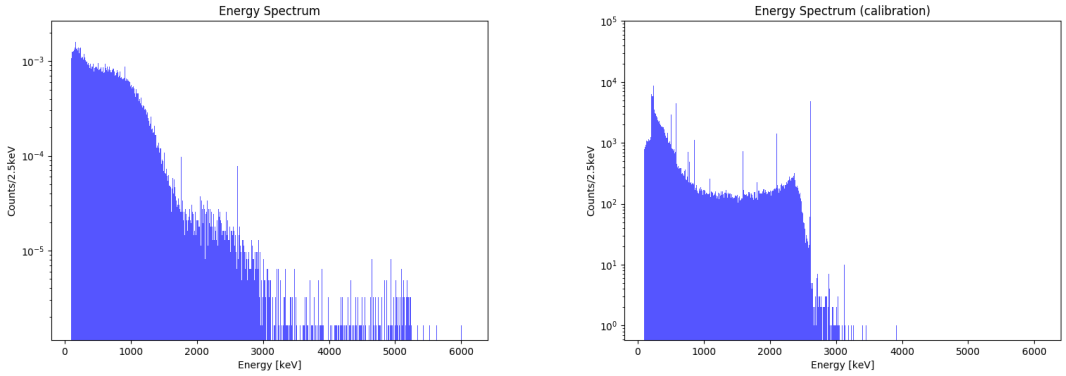

Figure B.1: Energy spectrum, for the full dataset and for one long calibration.

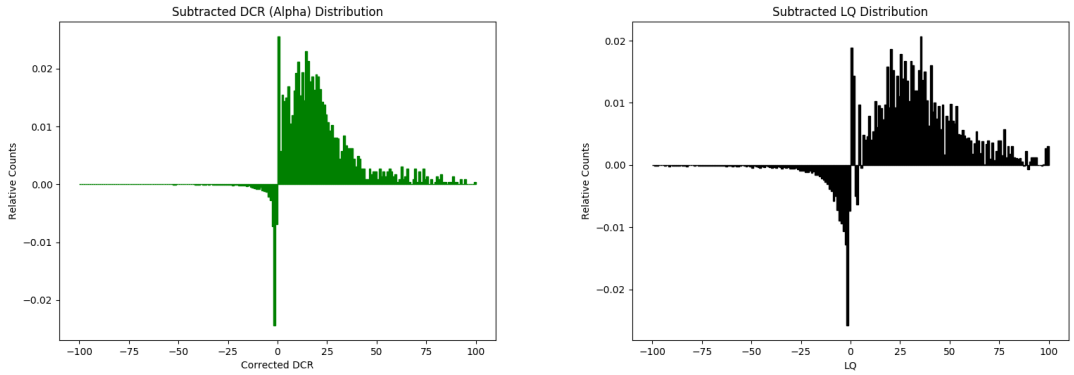

Figure B.2: The DCR and LQ distributions, subtracting the long calibration distributions from the full dataset distributions, in the region between 1500-1800 keV.

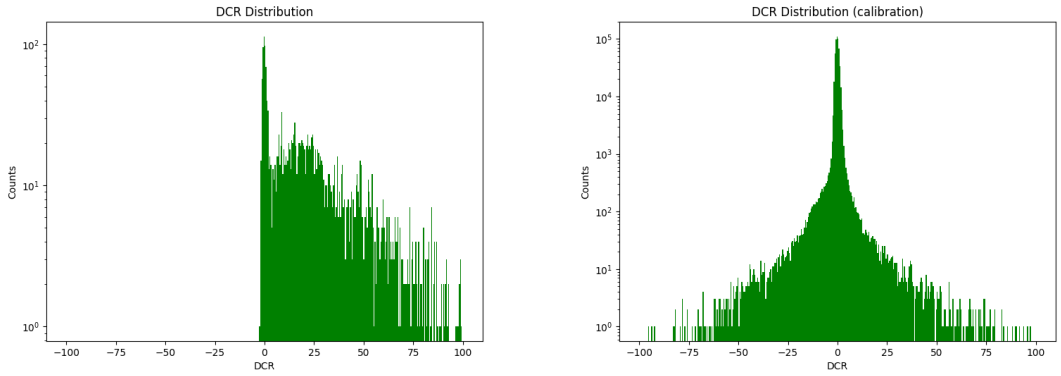

Figure B.3: The corrected DCR distribution, for the full dataset and for one long calibration, in the region between 1950-2350 keV. The distribution shows that above a few units of DCR, the distribution in the physics data is strongly dominated by events not present, or rarely present, in the calibration, which are expected to be alpha particles on the surface of our detector.

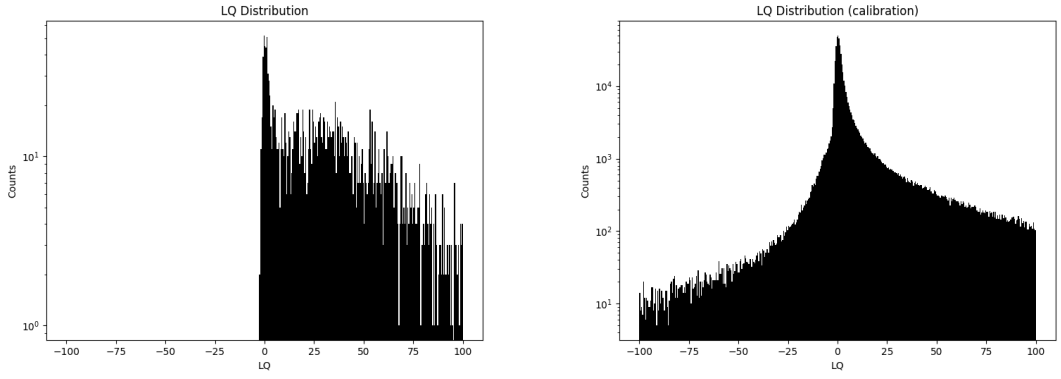

Figure B.4: The LQ distribution, for the full dataset and for one long calibration, in the region between 1950-2350 keV. The distribution shows that...

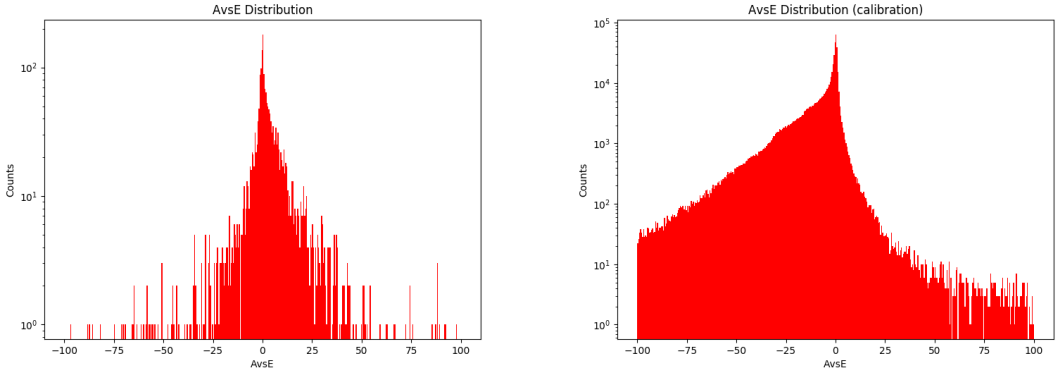

Figure B.5: The corrected AvsE distribution, for the full dataset and for one long calibration, in the region between 1950-2350 keV.

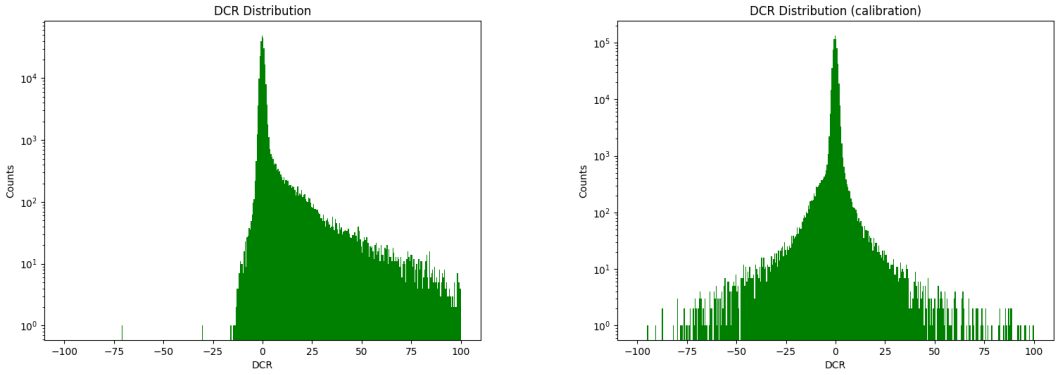

Figure B.6: The corrected DCR distribution, for the full dataset and for one long calibration, with energy greater than 100 keV.

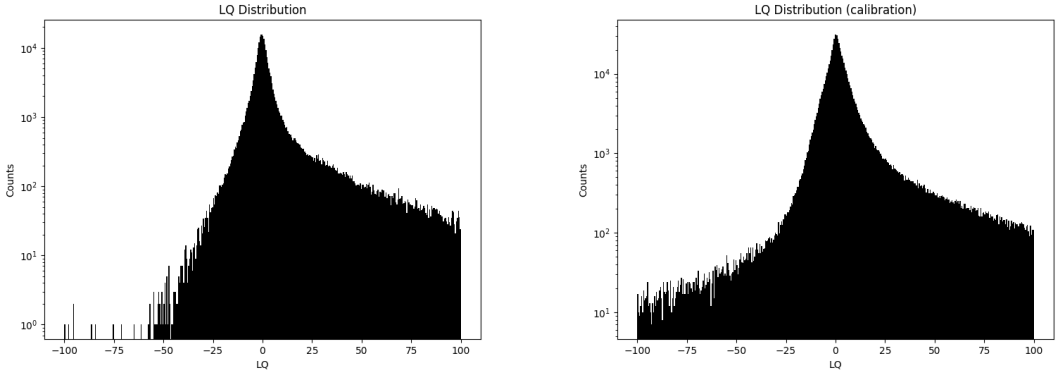

Figure B.7: The LQ distribution, for the full dataset and for one long calibration, with energy greater than 100 keV.

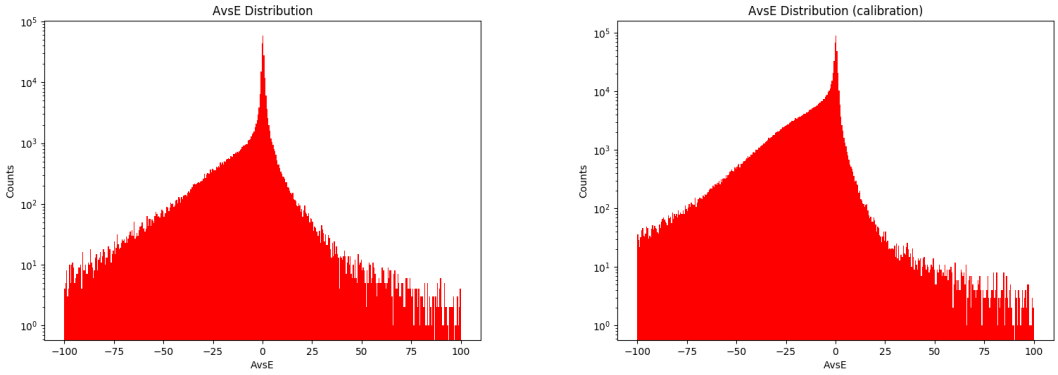

Figure B.8: The corrected AvsE distribution, for the full dataset and for one long calibration, with energy greater than 100 keV.

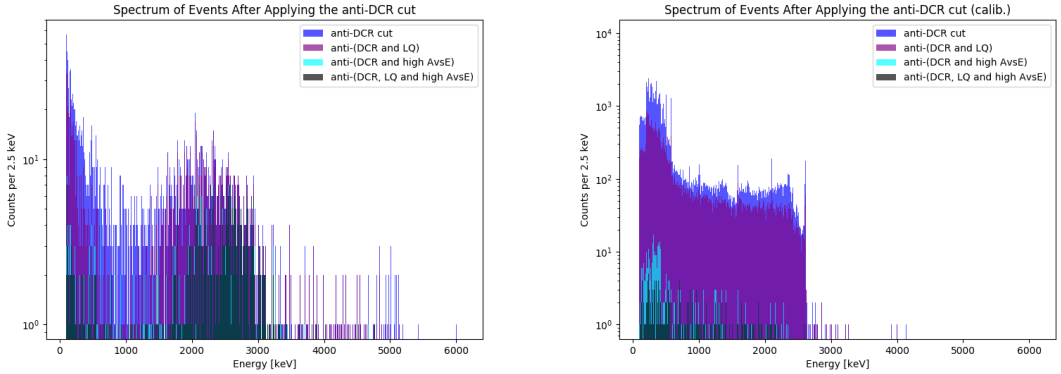

Figure B.9: Events failing the DCR cut, or passing the anti-DCR cut, i.e. with $DCR > 2.326$. Also shown, events also failing AvsE, LQ, or both. Shown for both data and one calibration.

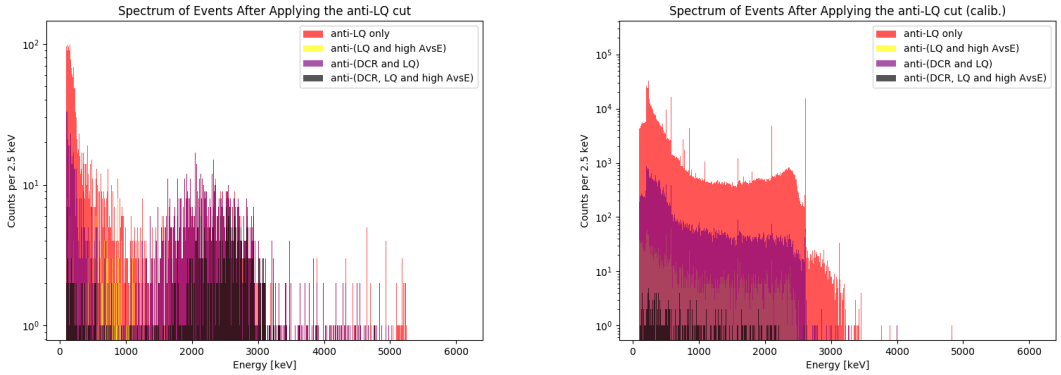

Figure B.10: Events failing the LQ cut, or passing the anti-LQ cut, i.e. with $LQ > 5$. Also shown, events also failing AvsE, DCR, or both. Shown for both data and one calibration.

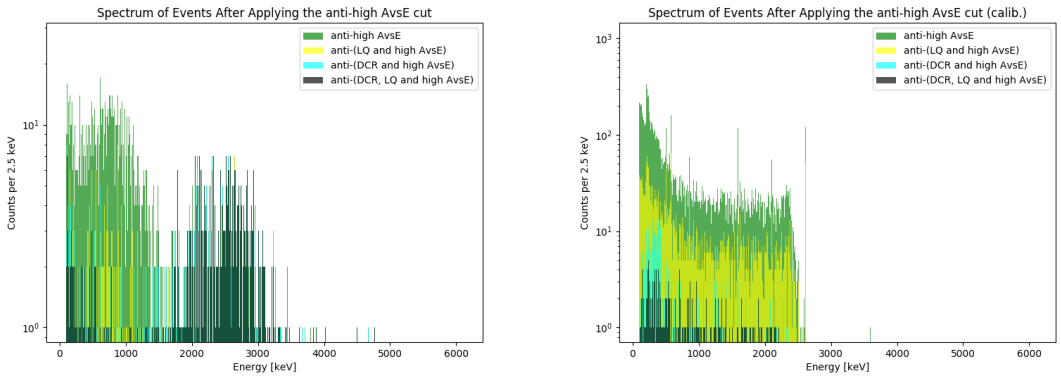

Figure B.11: Events failing the high AvsE cut, or passing the anti-high AvsE cut, i.e. with $AvsE > 9$. Also shown, events also failing LQ, DCR, or both. Shown for both data and one calibration.

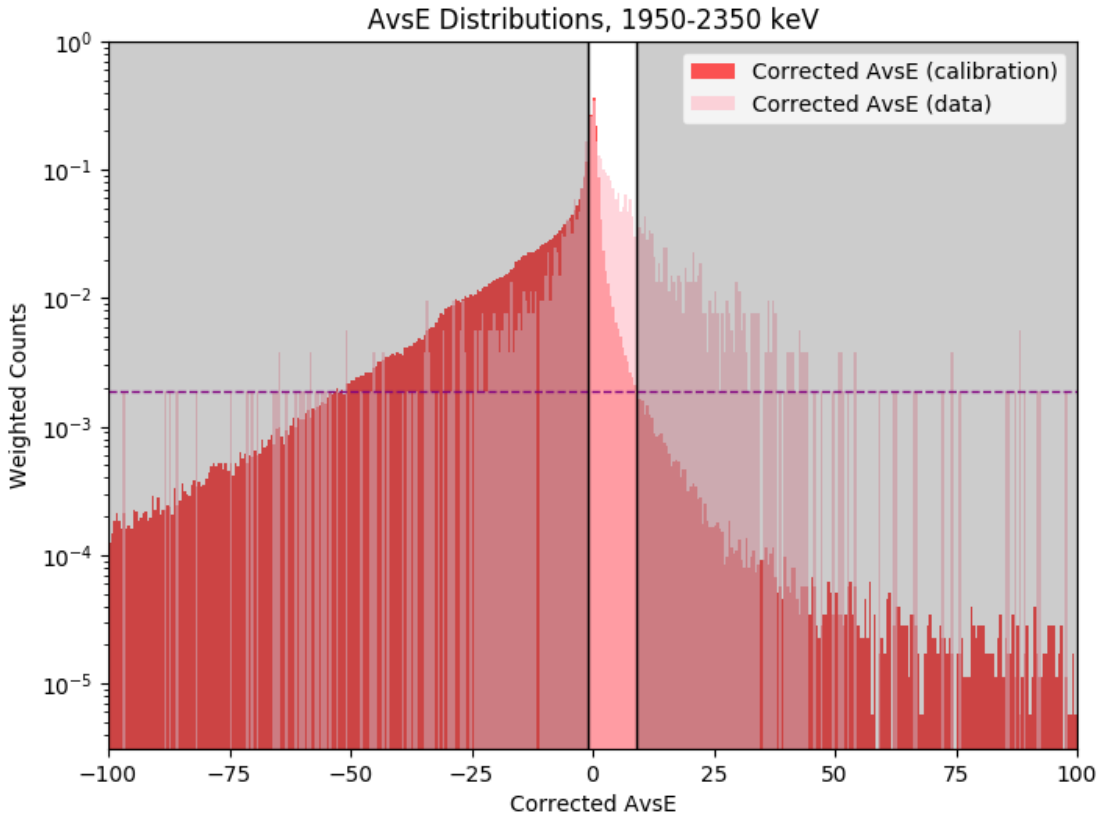

Figure B.12: The corrected AvsE distribution, for the full dataset and for one long calibration, overlaid, in the region between 1500-1800 keV.

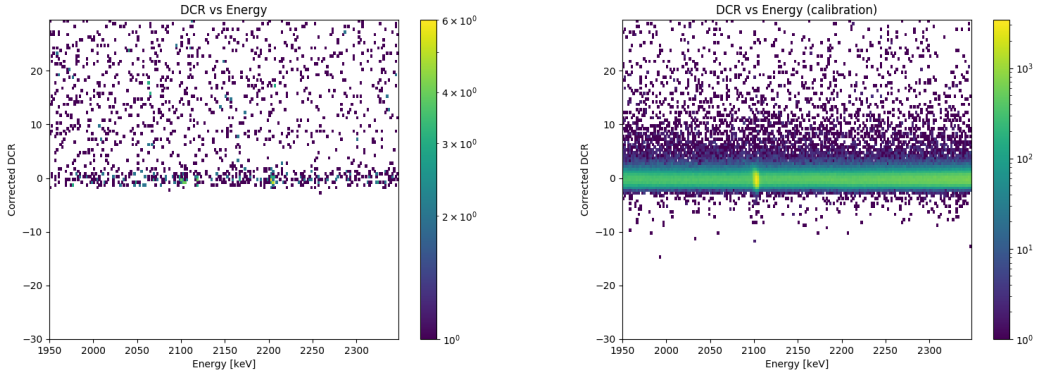

Figure B.13: Corrected DCR vs E, for the full dataset and for one long calibration, in the region between 1950-2350 keV.

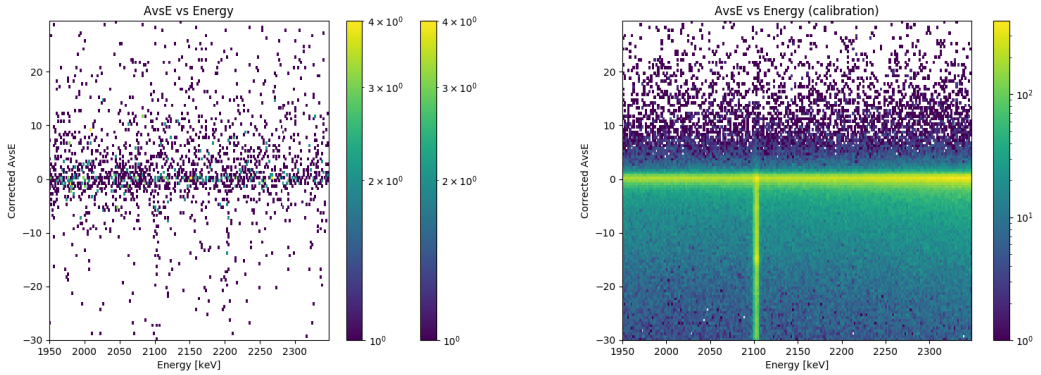

Figure B.14: Corrected AvsE vs E, for the full dataset and for one long calibration, in the region between 1950-2350 keV.

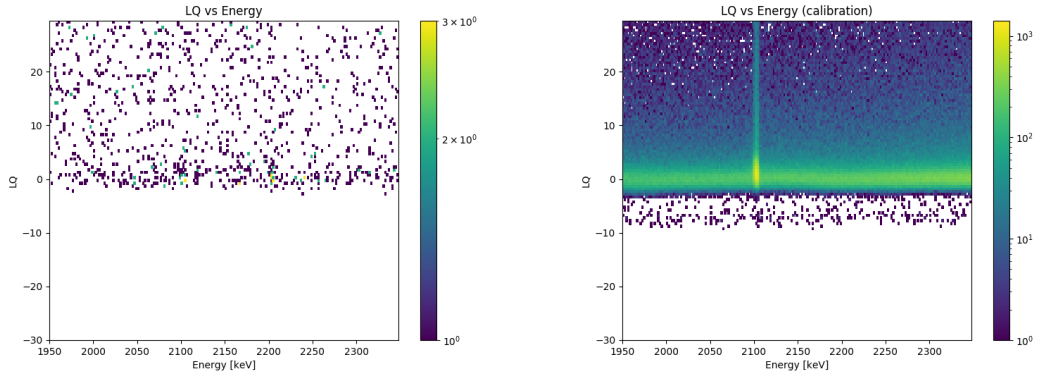

Figure B.15: LQ vs Energy, for the full dataset and for one long calibration, in the region between 1950-2350 keV.

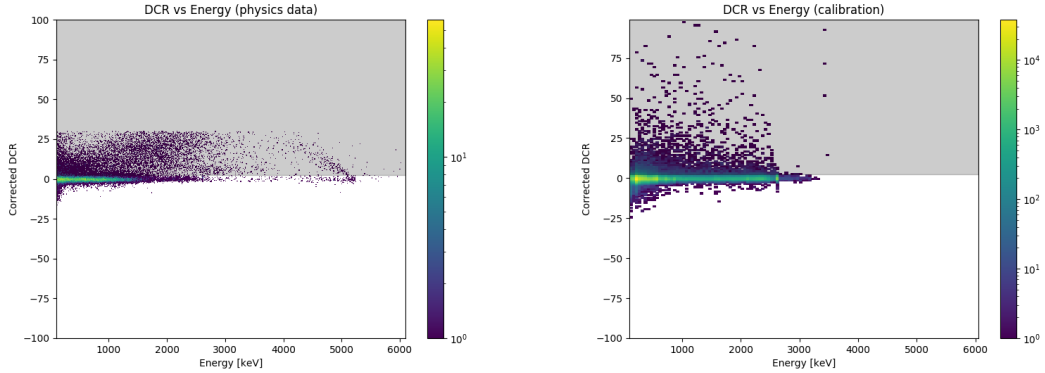

Figure B.16: Corrected DCR vs E, for the full dataset and for one long calibration, with energy greater than 100 keV.

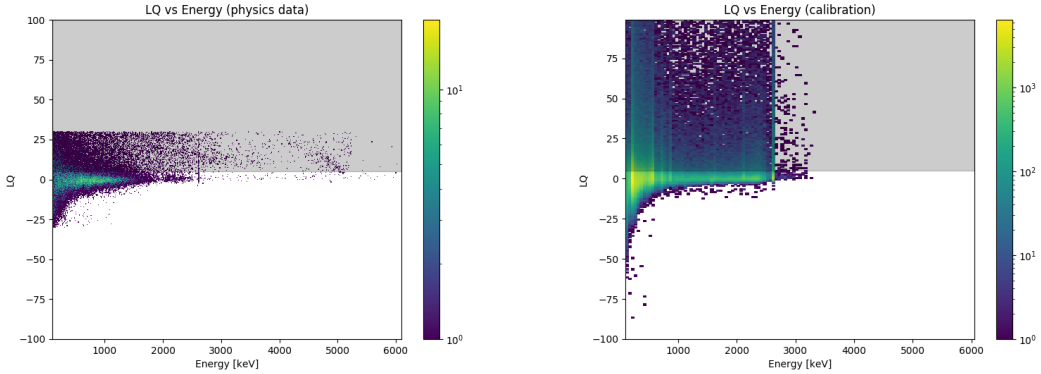

Figure B.17: LQ vs Energy, for the full dataset and for one long calibration, with energy greater than 100 keV.

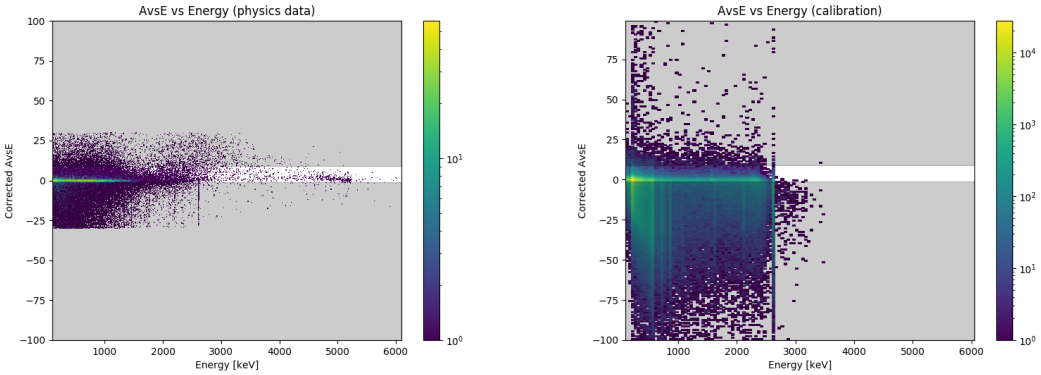

Figure B.18: Corrected AvsE vs E, for the full dataset and for one long calibration, with energy greater than 100 keV.