

Tests of the Effective Weak Interaction with the World's Smallest Neutrino Detector

by

Ryan Bouabid

Department of Physics  
Duke University

Defense Date: October 15, 2025

Approved:

Phillip Barbeau, Chair

Ashutosh Kotwal

Ronen Plesser

Kate Scholberg

Michael Troxel

Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of  
Philosophy in the Department of Physics in The Graduate School of Duke University  
2025

ABSTRACT

Tests of the Effective Weak Interaction with the World's Smallest Neutrino Detector

by

Ryan Bouabid

Department of Physics  
Duke University

Defense Date: August 29, 2025

Approved:

Phillip Barbeau, Chair

Ashutosh Kotwal

Ronen Plesser

Kate Scholberg

Michael Troxel

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Physics in The Graduate School of Duke University  
2025

Copyright by  
Ryan Bouabid  
2025

# Abstract

The Standard Model of particle physics is a remarkably successful predictive theory. Searches for physics beyond the Standard Model require precise predictions and precise experiments to look for subtle and unexpected deviations. This thesis investigates predictions and experiments of low energy weak interactions in the neutrino sector. Operating within the theoretical frame-work of the effective weak interaction, this thesis shows how careful experiments can serve as broadband searches for new kinds of physics.

In particular this thesis presents results on the first measurement of coherent elastic neutrino nucleus scattering by the COHERENT Collaboration on germanium. It also includes a follow up measurement that is the most precise measurement of CE $\nu$ NS to date.

These measurements, and the theoretical interventions that support them, show how future precision tests of the Standard Model in low energy neutrino physics are a complement to large colliders in their ability to probe the fundamental laws of nature.

# Dedication

*À mes parents,  
pour leur amour, leur soutien,  
et pour m'avoir donné le goût d'apprendre.*

# Contents

Abstract . . . . .	iv
List of Tables . . . . .	xi
List of Figures . . . . .	xii
List of Abbreviations . . . . .	xvi
Acknowledgements . . . . .	xviii
1 Introduction . . . . .	1
2 The Effective Weak Interaction . . . . .	4
2.1 Effective Lagrangian . . . . .	6
2.2 Constructing Nuclear Currents . . . . .	8
2.2.1 Nucleon-Level Currents . . . . .	8
2.2.2 Nuclear Currents and Matrix Elements . . . . .	9
3 $2\nu\beta\beta$ : A Case Study in Effective Theories of the Weak Interaction . . . . .	13
3.1 Double Beta Decay . . . . .	14
3.1.1 Transition Amplitude . . . . .	15
3.1.2 Nuclear Structure and Closure Approximation . . . . .	18
3.1.3 Differential Rates and Uncertainties . . . . .	20
3.2 NLO Corrections in Nuclear Current . . . . .	21
3.3 Double Beta Decay beyond the Standard Model . . . . .	23
3.4 Lessons for Other Weak Tests . . . . .	25
4 Coherent Elastic Neutrino Nucleus Scattering . . . . .	27
4.1 Theoretical Formalism . . . . .	28

4.1.1	Calculation . . . . .	28
4.1.2	CEvNS as a Probe . . . . .	31
4.1.3	Corrections . . . . .	32
4.1.4	Challenge 1: Interaction Rates . . . . .	34
4.1.5	Challenge 2: Recoil Energy Thresholds . . . . .	35
4.1.6	Challenge 3: Backgrounds . . . . .	36
4.2	Summary . . . . .	37
5	The COHERENT Collaboration at the Spallation Neutron Source . . . . .	38
5.1	The Spallation Neutron Source . . . . .	38
5.2	$\pi$ Decay at Rest . . . . .	39
5.3	Neutrino Alley . . . . .	42
5.4	Backgrounds . . . . .	43
5.5	Existing Measurements . . . . .	45
5.5.1	CEvNS . . . . .	46
5.6	Experimental Sources of Uncertainty . . . . .	47
6	The Ge-Mini Detector System . . . . .	50
6.1	Inverse Coaxial Point Contact Detector . . . . .	50
6.1.1	Operating Principle . . . . .	51
6.1.2	Signal Readout . . . . .	52
6.1.3	Ge-Mini Detectors . . . . .	53
6.2	Detector System Design . . . . .	54
6.2.1	Cryogenic Operation . . . . .	54
6.2.2	Shielding . . . . .	55
6.2.3	Muon Veto . . . . .	56
6.3	Commissioning and Campaign 1 . . . . .	57
6.3.1	Slow Controls . . . . .	57
6.3.2	Data Acquisition . . . . .	57

6.3.3	Data Processing . . . . .	60
7	Campaign 2: First Light . . . . .	62
7.1	Detector Characterization . . . . .	62
7.2	Data Collection . . . . .	64
7.2.1	Triggering Scheme . . . . .	65
7.2.2	Exposure Summary . . . . .	65
7.2.3	Calibration Runs . . . . .	66
7.2.4	Data Selection . . . . .	67
7.2.5	Blinded Analysis . . . . .	67
7.3	HPGe Waveform Analysis . . . . .	68
7.3.1	Pre-processing . . . . .	70
7.3.2	Energy Estimation . . . . .	70
7.3.3	Timing Extraction . . . . .	72
7.3.4	Reconstruction Uncertainty and Thresholds . . . . .	75
7.4	Data Quality and Cuts . . . . .	76
7.4.1	Timing Based Cuts . . . . .	78
7.4.2	Waveform Based Cuts . . . . .	80
7.4.3	Muon Veto . . . . .	81
7.4.4	Dead-Time . . . . .	83
7.5	Backgrounds . . . . .	84
7.5.1	Steady-State Backgrounds . . . . .	84
7.5.2	Beam-Related Backgrounds . . . . .	84
7.5.3	Background Model . . . . .	86
7.6	Signal Prediction . . . . .	86
7.6.1	Overview . . . . .	86
7.6.2	Flux Averaged Rates . . . . .	87
7.6.3	Quenching Factor . . . . .	88
7.6.4	Detector Smearing and Efficiencies . . . . .	90

7.6.5	Signal PDF . . . . .	92
7.7	Expected Counts . . . . .	92
7.8	Statistical Analysis . . . . .	92
7.9	Discussion . . . . .	94
8	Campaign 3: Towards Precision . . . . .	97
8.1	Data Collected . . . . .	98
8.1.1	Triggering Scheme . . . . .	98
8.1.2	Exposure Summary . . . . .	98
8.1.3	Calibration Studies . . . . .	99
8.1.4	Data Selection . . . . .	99
8.1.5	Blinded Analysis . . . . .	100
8.2	HPGe Waveform Analysis . . . . .	101
8.2.1	Pre-processing . . . . .	101
8.2.2	Energy Estimation . . . . .	101
8.2.3	Timing Extraction . . . . .	101
8.2.4	Pulse Shape Discrimination . . . . .	104
8.3	Data Quality and Preparation . . . . .	112
8.3.1	Timing Based Cuts . . . . .	113
8.3.2	Waveform Based Cuts . . . . .	114
8.3.3	Muon Veto . . . . .	116
8.3.4	Dead-Time . . . . .	116
8.4	Backgrounds . . . . .	118
8.4.1	Off-Beam Data After All Cuts . . . . .	118
8.4.2	High Statistics Internally Triggered Dataset . . . . .	122
8.4.3	Simultaneous Fits to Background Data . . . . .	124
8.4.4	Beam Related Backgrounds . . . . .	125
8.4.5	Background Model . . . . .	125
8.5	Signal Expectation . . . . .	127

8.5.1	Overview . . . . .	127
8.5.2	Improved Smearing . . . . .	128
8.5.3	Signal PDF . . . . .	130
8.5.4	Expected Counts . . . . .	131
8.6	Statistical Analysis . . . . .	131
8.6.1	Results . . . . .	133
8.7	Discussion . . . . .	133
9	Non-Standard Interactions . . . . .	140
9.1	Formalism . . . . .	140
9.2	Statistical Analysis . . . . .	142
9.3	Results . . . . .	144
10	Conclusion . . . . .	145
	Bibliography . . . . .	147

# List of Tables

2.1	Non-relativistic reduction of the nuclear currents. . . . .	10
6.1	Manufacturer specifications for the mass of the Ge-Mini ICPC detectors. . . . .	54
7.1	Summary of HPGe detectors used during Campaign 2. . . . .	63
7.2	Summary of weekly calibration runs during Campaign-2 data-taking. . . . .	67
7.3	Optimal trapezoidal filter parameters during Campaign-2 data-taking. . . . .	73
7.4	Campaign-2 LN2 refill exclusion windows. . . . .	80
7.5	Optimized cut value for trapezoidal filter minimum in Campaign-2. . . . .	81
7.6	Campaign-2 dead-times from data-quality cuts. . . . .	83
7.7	Systematic uncertainties in signal prediction for Campaign-2. . . . .	93
8.1	Optimal trapezoidal filter parameters during Campaign-3 data-taking. . . . .	102
8.2	Additional periods of microphonics-related noise identified in Campaign-3. . . . .	114
8.3	Summary of cuts in Campaign-3 and the resultant dead-time. . . . .	118
8.4	Model parameters and whether they are shared or fixed in the Campaign-3 statistical analysis. . . . .	124
8.5	Total Campaign-3 exposure before and after cuts in the unit of GWhkg. . . . .	132
8.6	Standard model signal expectation for Campaign-3. . . . .	132
8.7	Unblinding results for Campaign-3 dataset. . . . .	139

# List of Figures

2.1	Schematic representation of Standard Model interaction and four-fermion interaction. . . . .	7
2.2	Schematic of the construction of nuclear currents from quark-level operators. . . . .	11
3.1	Induced spectral distortions from BSM tensor current as well as Standard Model weak magnetism. . . . .	24
3.2	Schematic illustration of the effective field theory hierarchy linking new high-energy interactions to measurable low-energy observables. . . . .	25
4.1	Diagram of the tree-level Coherent Elastic Neutrino Nucleus Scattering (CEvNS) interaction in the effective theory in the laboratory frame. . . . .	29
4.2	An example of differential rates of CEvNS on three target nuclei. . . . .	35
5.1	A recent example of the average current measured from protons on target at the SNS. . . . .	40
5.2	Dominant components of the SNS neutrino flux. . . . .	42
5.3	SNS engineering drawings that inform neutrino flux predictions. . . . .	43
5.4	Schematic showing the currently (and soon-to-be) operating neutrino detectors in the COHERENT collaboration. . . . .	44
5.5	The measured flux of neutrons at several locations in the SNS. . . . .	45
5.6	The full on-beam for the CsI dataset is shown as a residual against the CEvNS best fits. . . . .	46
5.7	Graphical representation of the error budget for the first germanium result by COHERENT. . . . .	48
6.1	Schematic of a HPGe WELL <sup>TM</sup> detector made by Mirion Technologies. . . . .	51
6.2	A schematic of the Ge-Mini detector system with key components labeled. . . . .	55

6.3	Photograph of the Ge-Mini detector system during assembly. . . . .	56
6.4	A photograph of the Ge-Mini detector system, on November 28, 2022 . . . . .	58
6.5	Example digitized waveform for a HPGe detector. . . . .	59
7.1	Drift-time calculations for COHERENT ICPC HPGe. . . . .	64
7.2	Timing diagram of the 120 Hz triggering scheme. . . . .	65
7.3	Cumulative beam exposure recorded over the duration of Campaign 2 as a function of date. . . . .	66
7.4	Average measured beam power versus time during Campaign 2. . . . .	68
7.5	Example waveform showing the pre-trigger baseline and the signal-arrival region. . . . .	69
7.6	An example kernel of the energy-estimating trapezoidal filter shown in the time domain. . . . .	71
7.7	An example of a rise time optimization performed on Ge-21 after the data-taking period. . . . .	72
7.8	HPGe waveform overlaid with waveform analysis filters used for Campaign-2. . . . .	75
7.9	Example timing reconstruction residuals and resolution ( $t_{50}$ ) as a function of true energy for Campaign-2. . . . .	77
7.10	Example pre-amplifier reset event. . . . .	79
7.11	An example of an event recorded during a liquid nitrogen Dewar fill. . . . .	79
7.12	Example pre-amplifier reset induced cross-talk. . . . .	82
7.13	Campaign 2 externally triggered background spectra integrated over the full data-taking period. . . . .	85
7.14	Campaign-2 internally triggered and externally triggered backgrounds compared. . . . .	87
7.15	Germanium quenching factor measurements and models. . . . .	89
7.16	Campaign-2 unblinded data and model best-fits. . . . .	95
7.17	The profiled log-likelihood of the amplitude of the signal distribution for Campaign-2. . . . .	95

8.1	An example Campaign-3 waveform from the externally triggered off-beam dataset and the analysis windows used. . . . .	99
8.2	Average recorded beam power delivered during Campaign 3 as a function of time.	100
8.3	A HPGe waveform and matched filter output . . . . .	103
8.4	Distributions of the error in reconstruction of pulse onset from Campaign-3 matched filter. . . . .	104
8.5	Energy vs onset reconstruction error for Campaign-3. . . . .	105
8.6	Example fast pulse and slow pulse. . . . .	106
8.7	Example slow pulse and fast pulse shown with the analysis filters used to determine PSD. . . . .	108
8.8	Example of T/E distributions using internally triggered calibration data for Ge-28.	109
8.9	Example fit of a fast pulse extracted from the $^{71}\text{Ge}$ x-ray to the analytic model described above. . . . .	110
8.10	Bulk event rise-time distributions as measured using uniformly distributed intrinsic radiation events. . . . .	110
8.11	Distribution of T/E for samples of toy waveforms generated with signal-like rise times. . . . .	111
8.12	Application of 99.7% signal-accepting PSD parameter on internally triggered calibration data for Ge21. . . . .	112
8.13	Example of a digitized waveform in an HPGe detector containing a so-called “negative pulse.” . . . .	116
8.14	An example distribution of energy deposited one of the scintillator panels that comprises the muon veto. . . . .	117
8.15	Full Campaign 3, externally triggered background energy spectra for HPGe detectors. . . . .	120
8.16	Full Campaign 3, externally triggered background pulse onsets for HPGe detectors.	121
8.17	Combined, internally triggered calibration data spectra for Ge detectors. . . . .	123
8.18	Simultaneous fit of calibration data and off-beam data for Ge21. . . . .	125

8.19	Simultaneous fit of calibration data and off-beam data for Ge25. . . . .	126
8.20	Simultaneous fit of calibration data and off-beam data for Ge26. . . . .	126
8.21	Simultaneous fit of calibration data and off-beam data for Ge28. . . . .	127
8.22	Average beam trace over Campaign-3 after application of the beam stability cut. A FWHM of $\sim 350$ ns is estimated. . . . .	128
8.23	Differential recoil spectra, integrated over recoil energies with and without convo- lution over POT trace. . . . .	129
8.24	An example of smearing functions Campaign-3 signal prediction. . . . .	130
8.25	Campaign-3 signal PDFs used in statistical analysis. . . . .	131
8.26	Ge21 unblinded data and best fits to signal and background are shown. . . . .	134
8.27	Ge25 unblinded data and best fits to signal and background are shown. . . . .	135
8.28	Ge26 unblinded data and best fits to signal and background are shown. . . . .	136
8.29	Ge28 unblinded data and best fits to signal and background are shown. . . . .	137
8.30	Campaign-3 profiled log-likelihoods for the number of counts in each detector. . . . .	138
8.31	Campaign-3 PLL of each detector is expressed as a ratio of the observed to predicted number of counts. . . . .	138
8.32	Combined unblinded datasets across all four detectors in Campaign-3 . . . . .	139
9.1	Two-dimensional, profiled allowed regions for $(\varepsilon_d^{\mu\mu}, \varepsilon_u^{\mu\mu})$ plane, with all other parameters fixed to zero. . . . .	141
9.2	Two-dimensional, profiled allowed regions for $(\varepsilon_u^{ee}, \varepsilon_u^{\mu\mu})$ plane, with all other parameters fixed to zero. . . . .	143

# List of Abbreviations

$2\nu\beta\beta$	Two neutrino double beta decay . . . . .	13
<b>CEvNS</b>	Coherent Elastic Neutrino Nucleus Scattering . . . . .	xii
<b>CC</b>	charged-current . . . . .	4
<b>EFT</b>	Effective Field Theory . . . . .	4
<b>ENC</b>	equivalent noise charge . . . . .	53
<b>evt-39</b>	event-39 . . . . .	59
<b>evt-61</b>	event-61 . . . . .	59
<b>FWHM</b>	full-width half-maximum . . . . .	53
<b>HPDE</b>	high density poly ethylene . . . . .	55
<b>HPGe</b>	High Purity Ge . . . . .	50
<b>ICPC</b>	Inverse Coaxial Point-Contact . . . . .	50
<b>LEGEND</b>	Large Enriched Germanium Experiment for Neutrinoless $\beta\beta$ Decay . . . . .	2
<b>LN2</b>	liquid nitrogen . . . . .	54
<b>NC</b>	neutral-current . . . . .	6
<b>NLO</b>	next-to-leading order . . . . .	14
<b>NME</b>	nuclear matrix element . . . . .	17
<b>NSIs</b>	non-standard interactions . . . . .	140
<b>MW</b>	megawatt . . . . .	39
<b>m.w.e.</b>	meters water-equivalent . . . . .	44
<b>ORNL</b>	Oak Ridge National Lab . . . . .	38
<b>PDF</b>	probability density function . . . . .	86
<b>PMTs</b>	photomultiplier tubes . . . . .	56
<b>POT</b>	Protons On Target . . . . .	64
<b>PSD</b>	pulse shape discrimination . . . . .	101
<b>SMEFT</b>	Standard Model Effective Field Theory . . . . .	23

<b>SNS</b>	Spallation Neutron Source . . . . .	38
<b>TRP</b>	transistor reset pre-amplifiers . . . . .	53
<b>TTL</b>	transistor-transistor logic . . . . .	78
<b>WM</b>	weak magnetism . . . . .	21

# Acknowledgements

PhDs, unfortunately, cannot bubble out of the vacuum. They are produced in a dense field of support and contributions from countless people. I certainly do not even know the names of each person who has created the context in which I could contribute a small piece of our shared understanding of the world. But, for those that I do know I have the privilege of being grateful to them.

My mother Majida Bargach and my father Ali Bouabid are fundamentally courageous, curious, and loving parents and without them there would be no PhD. I owe my love of learning to them, a debt I could never repay.

My older brother Neil Bouabid is without peer in his unconditional love and support.

Throughout my PhD I have been supported by the Department of Energy through the Monitoring Technology and Verification university consortium as well as the Office of Science Graduate Student Research Program. Both of these fellowships have given me enriching experiences as a young researcher as well as generous material support.

What can I say about my advisor Phil Barbeau? There is a dye that stains the hands of people that build experiments that bump up against the universe. Phil's hands are covered in that dye<sup>1</sup>. I will always be grateful that he opened his doors to me to get my hands dirty with him.

The members of Phil Barbeau's research group, past and present, defined my experience as a student. There are contributions to this work that cannot be enumerated — without a doubt this has been the most collaborative environment I have ever been in. Grayson, Long, Sam, Connor, and Peibo all helped create what it means to be part of this group.

Jay Runge is a flywheel of creativity and I look forward to one day measuring the orientation of temples in Malta with him.

Tyler Johnson is my personal high-water mark for taking on impossible challenges and

---

<sup>1</sup> He might get a kick about the fact that this metaphor is in part inspired by a passage from *The Republic*, or he might think I am once again being a “that kid” — hopefully it's a bit of both.

he inspires me to do the same. He also helped keep me relatively sane throughout my time at Duke. I hope we have many years ahead of us to talk physics and not-physics.

I owe much of what I understand about the statistics of particle physics to Charlie Prior. I'm grateful for his enthusiasm for both learning and teaching.

There are other members of the group who played important roles in the every day experience of physics graduate school. Aitor Bracho, among other things, was an important companion in the era of Zoom learning during the pandemic. Emma van Nieuwenhuizen played a central role in deploying the experiment written about in this thesis. I owe Ana Zaalishvili in no small part for cheerfully helping me to stay on track towards the end of my degree. Natalie Jones deserves mention for, among other things, tolerating after-hours theft of gummy candies. And Isabel Colón Rivera was a consistently enthusiastic appreciator of physics which served as a steady reminder that this should all be fun.

The COHERENT collaboration, and its home at Oak Ridge National Lab, are full of people that made this work possible. I could not name them all though their contributions are essential. Kate Scholberg in particular has succeeded in creating a uniquely enriching and supportive environment for students through this collaboration. In that same spirit, I owe Jason Newby for always being willing to roll up his sleeves and confront the tactile reality of pulling off experiments at Oak Ridge.

In 2023 I spent six months at Los Alamos National Lab as a graduate student fellow. I am deeply appreciative to the Weak Interactions Team there for being gracious hosts and mentors.

During this time I also benefited from the immeasurable support of my new family which was one of the many blessings I received from my wife. My new sister Saanya Maini has been a constant source of laughter. My new matashri Ashima Maini has been a well of support and wisdom.

Finally, my wife Shreya Maini is the unnamed co-author of any contribution I made in this work. She heard every idea, every challenge, every rant, every breakthrough and she gave me the motivation and joy to see all of it to its end.

# Chapter 1

## Introduction

What are particle physicists supposed to study? The question is not at all rhetorical and, indeed, there are many answers that float around the community. As a small contribution, it is worth considering the limited perspective of a researcher who has joined the field at a time where this question did not have an obvious answer.

For a portion of the modern history of the field of particle physics “discoveries” were a thing that could be anticipated. Remarkable as it is, particles have repeatedly been theorized before their observation. Clear road-maps have existed for planned progress in the understanding of the laws of nature. The same cannot be said for this current moment.

For a generation of particle physicists who have taken up the profession in a post-Higgs world, that certainty seems alien. The super-symmetry did not show up in the colliders, the dark matter is still a mystery, and the Standard Model of particle physics has stubbornly held on as humanity’s greatest predictive theory. Of course, a WIMP discovery could be announced tomorrow or the Majorana nature of neutrinos could be revealed within the year and everything would fall into place just as everyone predicted it would. And yet, just as easily those discoveries might be generations away.

So it is worth really considering the question: what are particle physicists supposed to study? The work of this thesis is a small slice of one type of answer: *look everywhere and have fun doing it.*

Decades of institutional effort have carved out literal and metaphorical mines to search for new physics. These Big Projects that are well motivated should be contributed to by cohorts of creative and eager young scientists. And at the same time, it is important to look for smaller rocks that have not yet been overturned because there is simply no way to predict where the next hint will come for the new physics.

History tells us, so far, that the neutrino might hold some clues. Some of the first hints

at the new physics have come from the unexpected phenomena of neutrino oscillations — the first cracks in the frustrating accidental symmetries of the Standard Model. Fortuitously there are still many rocks to turn over in neutrino physics.

This thesis presents new results in the theoretical and experimental study of the weak interaction as probed by low-energy neutrinos. It begins in Chapter 2 with a review of the theoretical framework that underpins these investigations. Chapter 3 then details novel contributions to the precise calculation of double beta decay rates, including new strategies for using double beta decay to search for physics beyond the Standard Model. This work appears in the *Journal of High Energy Physics* as “ $2\nu\beta\beta$  spectrum in chiral effective field theory” (Volume 2025, Article 82, 2025). Chapter 4 shifts focus to coherent elastic neutrino–nucleus scattering, summarizing the theoretical predictions for this process. The remainder of the thesis presents the primary experimental efforts and new results in coherent neutrino scattering, including work published in *Physical Review Letters* as “Evidence of Coherent Elastic Neutrino–Nucleus Scattering with COHERENT’s Germanium Array” (Phys. Rev. Lett. 134, 231801, 9 June 2025), as well as new, yet-unpublished results that build upon this measurement.

This thesis will belabor why exactly these are interesting places to look for the new physics but what will be harder to capture is, narratively, the type of personal development that occurs throughout the doctorate.

The theoretical interventions in this thesis are some of the results of a six month residency at Los Alamos National Laboratory. The main focus of this stay was to contribute to one of these big projects: the Large Enriched Germanium Experiment for Neutrinoless  $\beta\beta$  Decay (LEGEND). Contributions to the experimental efforts to LEGEND, including looking for ways to poke at the Standard Model with the experiment beyond the pressing search for Lepton Number Violation, do not appear in this thesis. This work however was an instrumental learning experience and eventually gave rise to fruitful collaborations with the Theory Division that hopefully continue for many years.

The experimental interventions are, broadly speaking, the result of a unique and nurturing

environment at both Duke University and Oak Ridge National Laboratory, manifested in the COHERENT collaboration. Within this collaboration, the doctoral research matured into independent contributions, culminating in new results on coherent elastic neutrino–nucleus scattering with germanium detectors. Beyond the analyses presented in this thesis, additional contributions were made to a range of collaborative efforts, including detector deployments and the organization of working groups devoted to sterile neutrino searches. Although these projects do not appear in detail here, they formed an important part of a broader engagement with the collaboration and helped shape the perspective and skills that inform the work that follows.

Taken together, the theoretical and experimental avenues of this work reflect an interplay between the frameworks that guide curiosities and the experiments that test them; and the joys of doing both.

# Chapter 2

## The Effective Weak Interaction

The focus of this work is on low-energy neutrinos, with energies comparable to the scale of nuclear excitations, where the neutrino wavelengths are large enough that they probe nuclei in their entirety. At these energies, nucleons and nuclei comprise the natural laboratory for studying the weak interaction, and what neutrinos may reveal about fundamental physics. Here, the phenomena of the weak interaction observed experimentally emerge as the aggregate of microscopic interactions hidden inside of composite particles. Handling this microscopic physics requires inputs from non-perturbative interactions in the strong force, and so is often prohibitively difficult. Despite this, the framework of Effective Field Theory (EFT) provides a prescription for making precise predictions, often treating nucleons and nuclei as fundamental particles.

In this picture, effective interactions, and even the matter-fields, are the low-energy limit of a theory that describes the physics above a certain energy scale. The development of the theory of the weak interaction is useful example of this concept [1]. The early development of the charged-current (CC) weak interaction was described with an interaction Lagrangian that exhausted all possibilities:

$$\mathcal{L}_{\text{CC}} = \sum_i G_i (\bar{n}(O_i)^\mu p) (\bar{e}(O_i)_\mu \nu) + h.c. \quad (2.1)$$

Here,  $n$ ,  $p$ ,  $e$ , and  $\nu$  are Dirac matter fields that describe neutrons, protons, electrons, and neutrinos respectively. The interaction responsible for beta decay is decomposed as a sum over all possible fermion bilinears of the form  $\bar{\psi}O_i\psi$ , where

$$O_i \in \{1, \gamma^\mu, \gamma_5, \gamma^\mu \gamma_5, \sigma^{\mu\nu}\}. \quad (2.2)$$

are linearly independent Dirac matrices. They correspond to the scalar, vector, pseudoscalar,

axial-vector, and tensor currents. Together, the bilinear currents represent the Lorentz structures that can appear in four-fermion interactions and form a complete basis under Lorentz transformations. In full generality, each could, in principle, contribute to weak interaction currents observed at low energies. Which of these actually occur in nature is determined by observation.

Indeed, each operator introduces a different Lorentz structure, which manifests itself experimentally through different observable phenomena. In experiments, this impacts the differential rates that are the physical observables of the theory.

The development of weak interaction studies revealed that parity is maximally violated in processes involving neutrinos, and that beta decays could only be consistently described if the interaction had a chiral structure [2]. These results led to the formulation of the vector–axial  $V - A$  theory, in which the charged–current weak interaction couples left–handed leptons to hadronic currents through a mixture of vector and axial–vector operators [3]. This structure can be expressed as

$$\mathcal{L}_{CC} = G_F (\bar{n}\gamma^\mu(1 - g_A\gamma_5)p) (\bar{e}\gamma_\mu(1 - \gamma_5)\nu) + \text{h.c.} \quad (2.3)$$

The resulting framework encoded the Lorentz structure required to match experimental observations. Refinements for precision predictions were introduced through the development of form factors, which parametrize nucleon and nuclear structure effects not captured by the effective theory. Higher-order effects, such as electroweak radiative corrections, were also incorporated [4]. This effective description remains a remarkably successful framework for understanding low-energy weak interactions.

Despite its success, the low-energy effective theory predicts neutrino–matter interaction probabilities that grow with energy, eventually exceeding unity. In this context, the “probability” refers to the scattering probability for a neutrino to interact with a single nucleon or nucleus. Because the Fermi interaction treats the weak force as a point-like contact

interaction, the predicted cross section grows as

$$\sigma \sim G_F^2 E_\nu^2. \quad (2.4)$$


At sufficiently high neutrino energies (a couple of hundred GeV), the interaction probability becomes unphysical. The full electroweak theory, with explicit  $W$  and  $Z$  propagators, restores sensible predictions for these energies. Historically, this limitation of the Fermi theory was a key motivation for the development of a unified electroweak model. By contrast, the full Standard Model remains well-behaved across all experimentally accessible energies.

In this thesis, a similar perspective is adopted: the Standard Model itself may be a low-energy limit of a more fundamental theory. By expressing both the Standard Model and possible new physics within an EFT framework, it becomes possible to search for subtle signatures of high-energy phenomena in precision studies of low-energy neutrino interactions. More specifically, this thesis focuses on a new measurement in neutrino-nucleus interactions and how, when cast in this language of EFT, it can be used as a broad net to search for new physics.

## 2.1 Effective Lagrangian

The work of this thesis is primarily focused on interactions between neutrinos and nuclei. Since there are no nuclei in the Standard Model, EFT predictions need a prescription to build up a picture of neutrino-nucleus interactions. That is, there needs to be a way to go from operators that mediate neutrino-quark interactions all the way to operators that mediate neutrino-nucleus interactions. Figure 2.1 illustrates the goal of this prescription schematically.

While the main experimental intervention of this work is in a particular neutral-current (NC) interaction, the theoretical interventions are primarily focused on a particular CC interaction. In general, the application of the framework of EFT to compare precise measurements to precise predictions is the same, and so, in this section, both CC and



**Figure 2.1:** Transition from the Standard Model description of neutrino-quark scattering via  $Z$  exchange (left) to the low-energy effective theory for coherent neutrino-nucleus scattering (right).

NC interactions will be considered. Importantly, the effective weak theory inherits the symmetries of the Standard Model. As a result, the couplings that appear in CC and NC interactions are not independent, but related by the underlying electroweak structure. This ensures that the two interactions provide complementary probes of new physics.

The starting point of the EFT framework is to enumerate possible operators that are relevant in describing a physical process. Not all operators parametrize Standard Model interactions, and many may not even be well-motivated. Tracking the dimensionality of operators is a key feature of EFT. For energies below a characteristic energy scale  $\Lambda$ , one can consider any operator of dimension  $d > 4$  that enter in an expansion of the interaction Lagrangian in powers of  $\Lambda^{d-4}$ . The familiar  $V - A$  picture, which parametrizes the standard charged and neutral vector bosons, is described by dimension six operators.

After identifying the relevant operators, the low-energy, quark-level effective NC and CC Lagrangians can be written as [1]:

$$\mathcal{L}_{\text{NC}}^{\text{eff}} = \frac{G_F}{\sqrt{2}} \sum_q (\bar{\nu}\gamma^\mu(g_V^\nu - g_A^\nu\gamma_5)\nu)(\bar{q}\gamma^\mu(g_V^q - g_A^q\gamma_5)q), \quad (2.5)$$

$$\mathcal{L}_{\text{CC}}^{\text{eff}} = \frac{G_F}{\sqrt{2}} (\bar{\nu}\gamma^\mu(1 - \gamma_5)e)(\bar{u}\gamma^\mu(1 - \gamma_5)d). \quad (2.6)$$

Here, neutrino indices have been suppressed and for the CC interaction the particular quark flavors are chosen to correspond to beta decay. In general, only  $u, d$ , and  $s$  quarks are considered. Coefficients of the effective theory have been replaced with couplings of the

Standard Model,

$$g_V^\nu = \frac{1}{2}, \quad g_A^\nu = \frac{1}{2},$$

$$g_V^q = T_3^q - 2Q^q \sin^2 \theta_W = \begin{cases} \frac{1}{2} - \frac{4}{3} \sin^2 \theta_W, & q = u, c, t, \\ -\frac{1}{2} + \frac{2}{3} \sin^2 \theta_W, & q = d, s, b. \end{cases} \quad (2.7)$$

Where  $T_3^q$  is the weak isospin of the gauge group,  $Q$  is the electromagnetic charge, and  $\theta_W$ , often called the weak mixing angle, is a parameter of the Standard Model. The effective Lagrangians are factored into leptonic currents and quark currents.

This is the effective weak theory to first order in perturbation theory. Higher order corrections, in particular electromagnetic radiative corrections, can be parametrized as corrections to the couplings described at the quark level.

## 2.2 Constructing Nuclear Currents

Connecting the quark-level effective operators to nucleon and nuclear interactions, proceeds in two steps: first to nucleons, then to nuclei.

### 2.2.1 Nucleon-Level Currents

Quark currents are projected onto nucleon states using phenomenological form factors that encode the effects of non-perturbative Quantum Chromo Dynamics (QCD). For a nucleon  $\eta = n, p$  and quark flavor  $q$ , the vector and axial currents are parametrized as

$$\langle \eta(p_f) | \bar{q} \gamma^\mu q | \eta(p) \rangle = \bar{u}(p_f) \left[ F_1^{q,\eta}(q^2) \gamma^\mu + i \frac{F_2^{q,\eta}(q^2)}{2m_N} \sigma^{\mu\nu} q_\nu \right] u(p), \quad (2.8)$$

$$\langle \eta(p_f) | \bar{q} \gamma^\mu \gamma_5 q | \eta(p) \rangle = \bar{u}(p_f) \left[ G_A^{q,\eta}(q^2) \gamma^\mu \gamma_5 + G_P^{q,\eta}(q^2) q^\mu \gamma_5 \right] u(p), \quad (2.9)$$

where  $q = p_f - p$  is the four-momentum transfer,  $u(p)$  is the nucleon spinor, and  $F_{1,2}$  and  $G_{A,P}$  are the Dirac, Pauli, axial, and pseudoscalar form factors [5, 6]. This Lorentz-covariant form,  $\bar{u}(p_f) \Gamma^\mu u(p)$ , is the same spinor structure that appears in fermionic S-matrix elements,

with the form factors encoding the composite internal structure of the nucleon.

In the limit  $q^\mu \rightarrow 0$ , the Dirac form factor  $F_1^{q,\eta}(0)$  reduces to the quark charge of the nucleon and the axial form factor  $G_A^{q,\eta}(0)$  to its axial charge. In this regime, the currents become momentum-independent and describe the static weak charges of the nucleon. The CC and NC currents differ in their flavor structure: NC currents involve a sum over quark flavors weighted by their weak couplings, while CC currents connect  $u \leftrightarrow d$  quarks in beta decay.

### 2.2.2 Nuclear Currents and Matrix Elements

In the non-relativistic limit, where momentum transfers  $q$  are small compared to the rest mass of the nucleus  $M$ , the nuclear current can be expressed as a sum over all nucleon currents [7]. This approximation treats each nucleon as a point like target and neglects many-body interactions. It is a reasonable approximation below  $\sim 100$  MeV, and a standard regime for the work of this thesis. Here, the nuclear current can be written in the rest frame of the nucleus as:

$$J_{\text{nucleus}}^\mu(\vec{q}) = \sum_{k=1}^A e^{i\vec{q}\cdot\vec{r}_k} J_k^\mu(\vec{q}), \quad (2.10)$$

where  $k$  labels nucleons,  $\vec{r}_k$  is the position of the  $k$ -th nucleon, and  $J_k^\mu(\vec{q})$  is the single-nucleon current. Corrections to the static limit can be expressed as an expansion in powers of  $q/M$  [8, 9]:

$$J_{\text{nucleus}}^\mu(\vec{q}) = \sum_{n=1}^A e^{i\vec{q}\cdot\vec{r}_k} \sum_{k=1} \left[ g_V V^{(k)\mu} + g_W W^{(k)\mu} - g_A A^{(k)\mu} - g_P P^{(k)\mu} \right], \quad (2.11)$$

where  $V^{(k)}$ ,  $W^{(k)}$ ,  $A^{(k)}$ , and  $P^{(k)}$  are general non-relativistic operators of a nucleon inside the nucleus which is labeled by index  $n$ . The momentum dependent couplings  $g_V$ ,  $g_W$ ,  $g_A$ , and  $g_P$  are linear combinations of the nucleon-level form factors. The non-relativistic expansion in powers of  $|\vec{q}|/M$  is summarized in Table 2.1, showing which components dominate at the

zero momentum limit and which terms enter as first order corrections in the recoil expansion.

**Table 2.1:** Non-relativistic reduction of the nuclear currents to order  $k = 1$  in  $q/M$ . Time and spatial components are defined separately. Three-momentum  $\vec{q} = \vec{p} - \vec{p}'$ ,  $\vec{Q} = \vec{p} + \vec{p}'$ , and  $\vec{\sigma}$  is the Pauli spin matrix. This table is reproduced from [8].

	$k = 0$	$k = 1$
$V_0^{(k)}$	1	0
$W_0^{(k)}$	0	$-(1/2M) \mathbf{q} \cdot (\mathbf{q} - i\boldsymbol{\sigma} \times \mathbf{Q})$
$A_0^{(k)}$	0	$(1/2M) \boldsymbol{\sigma} \cdot \mathbf{Q}$
$P_0^{(k)}$	0	$-(1/2M) q^0 \boldsymbol{\sigma} \cdot \mathbf{q}$
$V^{(k)}$	0	$(1/2M)(\mathbf{Q} - i\boldsymbol{\sigma} \times \mathbf{q})$
$W^{(k)}$	$-i\boldsymbol{\sigma} \times \mathbf{q}$	$-(1/2M)q^0(\mathbf{q} - i\boldsymbol{\sigma} \times \mathbf{Q})$
$A^{(k)}$	$\boldsymbol{\sigma}$	0
$P^{(k)}$	0	$-(1/2M) \mathbf{q} \mathbf{q} \cdot \boldsymbol{\sigma}$

The nuclear matrix element for a transition  $|\alpha\rangle \rightarrow |\beta\rangle$  is

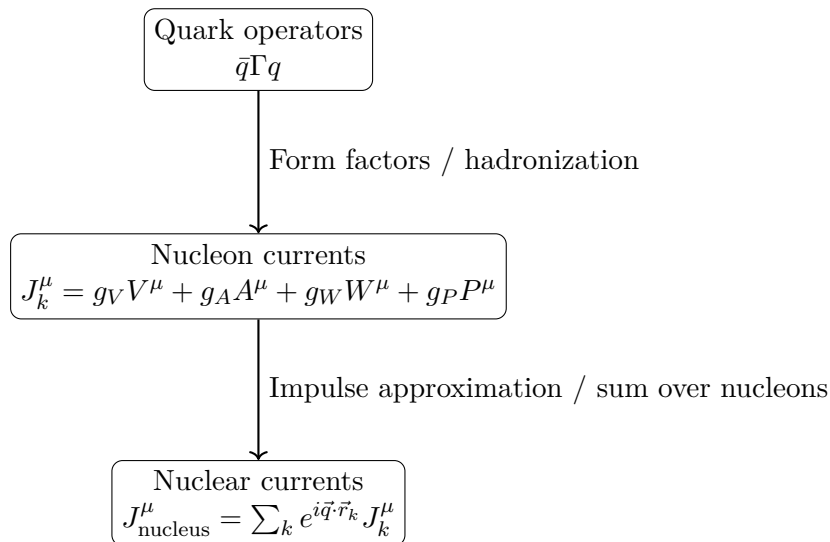
$$\langle \beta | J_{\text{nucleus}}^\mu(\vec{q}) | \alpha \rangle = \sum_{k=1}^A \langle \beta | e^{i\vec{q} \cdot \vec{r}_k} J_k^\mu(\vec{q}) | \alpha \rangle, \quad (2.12)$$

where  $J_k^\mu$  denotes the non-relativistic single-nucleon current. Here,  $|\alpha\rangle$  and  $|\beta\rangle$  are many-body wavefunctions of the nucleus expressed in coordinates of the individual nucleons,  $\vec{r}_k$ . The exponential factor  $e^{i\vec{q} \cdot \vec{r}_k}$  arises from the Fourier transform of the spatially distributed nucleon currents in the non-relativistic limit. In this way, nucleons are treated as localized within the extent of the nuclear wavefunction.

In practice, this expression can be reorganized into a multipole expansion [5, 10], which separates contributions according to their angular momentum and parity. The result connects the single-nucleon form factors to observable nuclear transition amplitudes, allowing non-perturbative effects to be systematically included within the EFT framework.

Figure 2.2 illustrates schematically the construction of nuclear currents from quark-level operators. Quark currents are projected onto nucleons using hadronic form factors, and nucleon currents are summed over the nucleons in the nucleus to form the nuclear current. Coherent contributions arise when the nucleon currents add constructively across the nucleus

(how this leads to the familiar  $N^2$  scaling in NC processes is shown in Chapter 4), while sub-dominant nuclear responses can enter as interferences with specific Lorentz structures.



**Figure 2.2:** Schematic of the construction of nuclear currents from quark-level operators. Quark currents are projected onto nucleons using hadronic form factors, and nucleon currents are summed over the nucleus.

At low momentum transfer, the leading contributions come from the vector  $g_V$  and axial  $g_A$  currents, while weak magnetism ( $g_W$ ), induced pseudoscalar ( $g_P$ ), and higher-order recoil corrections enter at subleading order. These suppressed terms can be systematically included to quantify theoretical uncertainties, consistent with the  $q/M$  expansion.

Corrections to the impulse approximation can be expressed as many-body currents which arise from meson-exchange or correlations between nucleons, that modify the total nuclear current beyond a simple sum of single-nucleon contributions [11]. These corrections affect both NC and CC interactions, particularly in spin-dependent and axial-vector channels [5], and can be parametrized as effective quenching of nuclear form factors. In this thesis these effects are still sub-dominant to experimental precision — though future studies of the low energy weak interaction can include these higher order corrections.

Low-energy experiments that test the weak interaction require a language to capture the effects of hadronic and nuclear physics. The familiar recipe summarized here connects the single-nucleon hadronic form factors to the experimental observables, such as total

amplitudes and Lorentz structure of interactions. This is the feature of the EFT picture most relevant to this work: by parameterizing the Standard Model in this basis calculations are tractable and deviations can be easily modeled, allowing precision neutrino experiments to probe physics, both old and new.

# Chapter 3

## $2\nu\beta\beta$ : A Case Study in Effective Theories of the Weak Interaction

In this work, double beta decay serves as a theoretical example of how EFT techniques provide a unified language connecting low- and high-energy physics by systematically encoding the effects of high-energy dynamics into the parameters of the low-energy theory. This same framework underlies the main results of this thesis in neutrino scattering, presented in later chapters, providing a prescription for how precise measurements of low-energy observables can be used to infer the presence of new interactions arising from physics at higher energy scales. This chapter therefore serves as a concrete test case for the methodology used throughout this thesis.

Beta decay has played a central role in the development of theory of the weak interaction. It is the canonical charged-current weak process. Historically, beta decay experiments led to the positing of the small neutral particle now known as the neutrino [12]. Today, it still serves as one of the most powerful probes of the Standard Model [13]. Two neutrino double beta decay ( $2\nu\beta\beta$ ) is a related process where a nucleus decays through the conversion of two neutrons to two protons. It is a second-order transition in the weak perturbation and, as a result, it is one of the rarest processes on earth.

The process of  $2\nu\beta\beta$  was first proposed by Maria Goeppert Mayer in 1935 [14], and went undiscovered for decades due to its long half-lives [15]. Because of this rarity, the experimental precisions of the differential rates of  $2\nu\beta\beta$  has lagged behind the precision of single beta decay experiments [16]. However, with the growing global program in the search for neutrinoless double beta decay ( $0\nu\beta\beta$ ), experimental capabilities in the precise measurements of  $2\nu\beta\beta$  are steadily improving. In particular, the low backgrounds present in these experiments, combined with the distinct kinematics of the process make it a potentially

complementary experimental approach to studying the weak CC interaction. Already, this has led to an increased interest in using this process as a probe of the Standard Model [17, 18].

The ability to use this rare process as a probe of physics beyond the Standard Model follows the recipe laid out in Chapter 2: within the context of an effective theory precise predictions can be made that parametrize the uncertainty encoded in the non-perturbative nuclear physics while simultaneously parameterizing any potential new physics. This strategy requires that the theoretical uncertainties be well understood, so as not to confuse the new physics for the intricacies of the microscopic physics of the nucleus.

This chapter<sup>1</sup> summarizes some of the formalism involved in precision calculations of the  $2\nu\beta\beta$  differential rate observable in the laboratory. It goes over some of the standard approximations that have sufficed up to now in experiments and summarizes some novel contributions in the next-to-leading order (NLO). In this framework this chapter covers how exotic physics might enter into this picture and how precise calculations are necessary to untangle the two.

### 3.1 Double Beta Decay

Double beta decay is a nuclear transition where two neutrons are converted into two protons, through the emission of two neutrinos and two beta particles. For the purposes of this work a simplifying assumption is made to consider only ground state to ground state transitions with zero total angular momentum. This is the dominant mode observed experimentally through which metastable isotopes can transition through double beta decay. This process is described by:

$$0^+(E_i) \rightarrow 0^+(E_f) + e^-(E_{e1}) + e^-(E_{e2}) + \bar{\nu}_e(E_{\nu1}) + \bar{\nu}_e(E_{\nu2}), \quad (3.1)$$

---

<sup>1</sup> This includes a few contributions to a publication that appeared in JHEP Volume 2025, article number 82, (2025).

where  $E_{e_1}, E_{e_2}, E_{\nu_1}, E_{\nu_2}$  are the energies of the final state leptons, respectively, and  $E_i$  and  $E_f$  are the energies of initial and final state nuclei.

The dominant contributions to the decay rate are well-studied in the literature [8, 19–21]. The differential decay rate can be written in terms of the leptonic phase space factor and nuclear transition amplitude  $|C_{2\nu}|$ ,

$$dW_{2\nu} = 2\pi \sum_{i,j} \sum_{\text{spin}} |C_{2\nu}|^2 \delta(E_{e_1} + E_{e_2} + E_{\nu_1} + E_{\nu_2} + E_f - E_i) \frac{d\mathbf{p}_1}{(2\pi)^3} \frac{d\mathbf{p}_2}{(2\pi)^3} \frac{d\mathbf{k}_1}{(2\pi)^3} \frac{d\mathbf{k}_2}{(2\pi)^3} \quad (3.2)$$

The differential rate depends on a sum over the participating neutrino flavors, which in the Standard Model includes the three light species. The kinematic terms,  $\mathbf{p}_1$ ,  $\mathbf{p}_2$ ,  $\mathbf{k}_1$ , and  $\mathbf{k}_2$  are the three momenta of the two electrons and two neutrinos, respectively. The remaining sum over spins accounts for all possible helicity configurations of the final-state particles, ensuring the rate reflects the total, unpolarized decay probability.

### 3.1.1 Transition Amplitude

The transition amplitude can be expressed as a contraction of the nuclear currents with the usual leptonic currents,

$$C_{2\nu} = \frac{1}{2} \left( \frac{G_F \cos \theta_c}{\sqrt{2}} \right)^2 U_{ei} U_{ej} \sum_n \int d\mathbf{x} d\mathbf{y} \langle 0^+ | J^{\nu\dagger}(\mathbf{y}) | n \rangle \langle n | J^{\mu\dagger}(\mathbf{x}) | 0^+ \rangle [1 - P(\mathcal{N}_i, \mathcal{N}_j)] \\ \times [1 - P(e_1, e_2)] \frac{\bar{e}_{p_2, s_2}(\mathbf{y}) \gamma_\nu (1 - \gamma_5) \mathcal{N}_{j, k_2, s_2}^c(\mathbf{y}) \bar{e}_{p_1, s_1}(\mathbf{x}) \gamma_\mu (1 - \gamma_5) \mathcal{N}_{i, k_1, s_1}^c(\mathbf{x})}{E_{e_1} + E_{\nu_1} + E_n - E_i}. \quad (3.3)$$

Here,  $G_F$  is the Fermi constant,  $\theta_c$  is the Cabibbo angle, and  $U_{ei}$  are matrix elements that encode the strength of flavor changing weak interactions.

The leptonic wavefunctions  $\bar{e}_{p,s}(\mathbf{x})$  and  $\mathcal{N}_{i,k,s}^c(\mathbf{x})$  represent the outgoing electron and neutrino spinors evaluated at the positions of the nucleonic currents inside the nucleus. In the plane-wave approximation, they take the form of Dirac spinors multiplied by their

spacetime phase factors,

$$\bar{e}_{p,s}(\mathbf{x}) = \bar{u}_{p,s} e^{i\mathbf{p}\cdot\mathbf{x}}, \quad \mathcal{N}_{i,k,s}^c(\mathbf{x}) = v_{i,k,s} e^{i\mathbf{k}\cdot\mathbf{x}}. \quad (3.4)$$

The integrals over  $\mathbf{x}$  and  $\mathbf{y}$  sum over the possible spatial locations of the two weak interaction vertices inside the nucleus, which are nonzero only where nucleons participate in the decay.

The operators  $P(e_1, e_2)$  and  $P(\mathcal{N}_i, \mathcal{N}_j)$  antisymmetrize the amplitude under exchange of identical leptons. Finally, the sum over intermediate nuclear states  $n$  runs over all virtual states connecting the initial and final even-even nuclei, and the energy denominator reflects these virtual states.

This expression illustrates the factorization into a purely nuclear part and a purely leptonic part. The nuclear matrix elements depend on the nuclear structure, while the leptonic currents determine the kinematic of the emitted electrons and neutrinos.

After anti-symmetrization, the transition amplitude can be expressed as four separate contractions which run over the sums of intermediate states. This contraction depends on the single-nucleon current, which can be read off the non-relativistic expansion discussed in Equation 2.2.2. In particular, in analogy to vertex corrections in Quantum Electro Dynamics, the leading order terms are related to the symmetries of the interaction — giving the familiar vector and axial currents — whereas the NLO are linear in  $q/m_n$  (where  $m_n$  is the mass of decaying neutrons). Considering only leading order terms for now, this current reduces to:

$$J^\mu(\mathbf{x}) = (J_V^0, \mathbf{J}_A). \quad (3.5)$$

Contributions for intermediate transitions decompose into the time-like vector (Fermi) transitions and the space-like axial (Gamow–Teller) transitions. After contracting with the leptonic currents, the expression can be decomposed into the time-like and space-like

components into terms such as:

$$L_1^\mu L_1^\nu J^{\mu\dagger} J^{\nu\dagger} = L_1^0 L_1^0 J_0^V J_0^V - L_1^i L_1^j J_i^A J_j^A ,$$

where the index 1 refers to an anti-symmetrized term under the exchange of identical leptons. Up to isospin symmetry breaking, the doubly Fermi transitions are negligible and the nuclear current reduces to the purely space-like, doubly Gamow–Teller transition.

The nuclear currents depends on intermediate sums of nuclear matrix element (NME) which are often collected in the literature into two terms:  $M^K$  and  $M^L$ . Simplifying leptonic kinematic terms are introduced:

$$\epsilon_K = (E_{22} - E_{11})/2, \quad \epsilon_L = (E_{12} - E_{21})/2 \quad (3.6)$$

in terms of  $E_{ij} = E_{ei} + E_{\nu j}$ . Using these, the nuclear current terms can be re-expressed as:

$$M_{GT}^K = \frac{m_e}{g_A^2} \sum_n \left( \frac{1}{E_n - \frac{1}{2}(E_i + E_f) - \epsilon_K} \langle 0^+ | \sum_k \vec{\sigma}_k \tau_k^+ | 1_n^+ \rangle \cdot \langle 1_n^+ | \sum_l \vec{\sigma}_l \tau_l^+ | 0^+ \rangle \right. \\ \left. + \frac{1}{E_n - \frac{1}{2}(E_i + E_f) + \epsilon_K} \langle 0^+ | \sum_k \vec{\sigma}_k \tau_k^+ | 1_n^+ \rangle \cdot \langle 1_n^+ | \sum_l \vec{\sigma}_l \tau_l^+ | 0^+ \rangle \right), \quad (3.7)$$

where  $m_e$  is the electron mass and  $g_A$  is the nucleon axial-vector form factor.  $M^L$  can be obtained by exchanging  $\epsilon_K \rightarrow \epsilon_L$ .

Written in this way it can be shown that the amplitude reduces to the purely spacelike contractions so that

$$C_{2\nu} \sim M_{GT}^K L_{11}^i L_{22}^i - M_{GT}^L L_{12}^i L_{21}^i, \quad (3.8)$$

where the leptonic currents have been re-expressed as tensors. After performing the traces that arise from the unpolarized decay, and integrating over the lepton phase space the

amplitude  $C_{2\nu}$  can be shown to be:

$$C_{2\nu} = \frac{g_A^4}{3} ((M^K)^2 + (M^L)^2 + (M^K M^L)). \quad (3.9)$$

To now this has followed the standard derivation of the decay amplitude in the literature, retaining only the terms that arise in the leading order of the non-relativistic reduction of the nuclear currents.

### 3.1.2 Nuclear Structure and Closure Approximation

From here, the prediction of experimental observables depends on the terms  $M^{K,L}$  which depend on large numbers of NME in the sum over intermediate states, making detailed calculations daunting. In the literature simplifying assumptions for computing double beta decay rates rely on tricks to limit the number of inputs from nuclear theory. Being able to control the nuclear theory uncertainties is key in hoping to make precise predictions.

The sum over virtual intermediate nuclear states is usually simplified by invoking the closure approximation:

$$\sum_n \frac{|n\rangle\langle n|}{E_i - E_e - E_\nu - E_n} \longrightarrow \sum_n \frac{|n\rangle\langle n|}{E_i - E_e - E_\nu - \langle E_n \rangle}, \quad (3.10)$$

where  $E_i$  is the energy of the initial nucleus,  $E_n$  denotes intermediate-state energies, and  $\langle E_n \rangle$  is an average excitation energy chosen to allow for a powerful simplification of factoring the denominator out of the intermediate sum in rate normalized spectra. Indeed, since typical excitation energies are  $\mathcal{O}(10)$  MeV, and the lepton energies are  $\mathcal{O}(1)$  MeV this is taken to be a reasonable assumption.

In Reference [21] this method is extended by considering a Taylor expansion in kinematic ratios of the denominator as follows. The two terms that appear in the symmetrized expressions for  $M_{GT}^K$  and  $M_{GT}^L$  contain energy denominators shifted by  $\pm\epsilon_{K,L}$ . Writing the

excitation energy of an intermediate state relative to the average of the ground states as

$$\tilde{A}_n \equiv E_n - \frac{1}{2}(E_i + E_f), \quad (3.11)$$

the pair of denominators in the symmetrized combination can be combined algebraically. For a single intermediate-state contribution proportional to  $M_n$  one encounters the combination

$$\frac{1}{\tilde{A}_n - \epsilon} + \frac{1}{\tilde{A}_n + \epsilon} = \frac{2\tilde{A}_n}{\tilde{A}_n^2 - \epsilon^2} = \frac{2}{\tilde{A}_n} \frac{1}{1 - (\epsilon/\tilde{A}_n)^2}, \quad (3.12)$$

where  $\epsilon$  denotes either  $\epsilon_K$  or  $\epsilon_L$  as appropriate.

Expanding the final factor as a geometric series in  $(\epsilon/\tilde{A}_n)^2$  (valid when  $|\epsilon| < |\tilde{A}_n|$ ) gives an even-power expansion,

$$\frac{1}{\tilde{A}_n - \epsilon} + \frac{1}{\tilde{A}_n + \epsilon} = \frac{2}{\tilde{A}_n} \sum_{m=0}^{\infty} \left( \frac{\epsilon}{\tilde{A}_n} \right)^{2m}. \quad (3.13)$$

Multiplying by the common prefactors (including the electron mass  $m_e$  and the nuclear structure factor  $M_n$ ) and summing over intermediate states yields the expansion

$$M_{GT}^{K,L} = m_e \sum_n M_n \frac{2}{\tilde{A}_n} \sum_{m=0}^{\infty} \left( \frac{\epsilon_{K,L}}{\tilde{A}_n} \right)^{2m} = \sum_{m=0}^{\infty} M_{GT}^{(-2m-1)} \left( \frac{\epsilon_{K,L}}{2m_e} \right)^{2m}, \quad (3.14)$$

where the nuclear-energy-weighted coefficients are defined as

$$M_{GT}^{(-2m-1)} \equiv 2 m_e (2m_e)^{2m} \sum_n \frac{M_n}{\tilde{A}_n^{2m+1}}, \quad (3.15)$$

so that the power counting is written in the form used in the literature. The  $M_n$  are the nuclear matrix elements given by  $\langle 0^+ | \sum_k \vec{\sigma}_k \tau_k^+ | 1_n^+ \rangle \cdot \langle 1_n^+ | \sum_l \vec{\sigma}_l \tau_l^+ | 0^+ \rangle$ , and all nuclear dependence is contained in the coefficients  $M_{GT}^{(-2m-1)}$ .

For the purposes of this work, the goal is to emphasize a case study in the factorization of nuclear uncertainties in the search of physics beyond the Standard Model. In this case it suffices to illustrate this point using only the leading order terms in the leptonic expansion

— that is, the standard closure approximation.

Written in this way the transition amplitude reduces to

$$C_{2\nu} \sim g_A^4 \sum_n M_n . \quad (3.16)$$

This result is a powerful simplifying assumption as the nuclear uncertainties have factored completely. For isotopes which undergo  $2\nu\beta\beta$  the lepton-energy combinations  $|\epsilon_{K,L}|$  are bounded by  $Q/2 \sim \mathcal{O}(\text{MeV})$ , whereas typical intermediate-state excitation energies  $\tilde{A}_n$  are  $\mathcal{O}(10\text{--}20 \text{ MeV})$ . As a result, the expansion parameter  $|\epsilon/\tilde{A}_n|$  is small, and the series converges rapidly; numerically the leading correction is at the few-percent level in most cases [21, 22].

### 3.1.3 Differential Rates and Uncertainties

Substituting the series (3.14) for both  $M^K$  and  $M^L$  into the squared amplitude  $C_{2\nu}$  and performing the required phase-space integrations produces an expansion of the observable rate in even powers of the lepton-energy combinations  $\epsilon_{K,L}$ . At leading order one recovers the usual closure result proportional to  $[M_{GT}^{(-1)}]^2$ , while the 2<sup>nd</sup>- and higher-order corrections are suppressed by powers of  $(\epsilon/\langle\tilde{A}\rangle)^2$ .

This leading order total rate can be expressed as

$$\Gamma = \frac{\ln 2}{\tau_{2\nu}} = g_A^4 \left( M_{GT}^{(-1)} \right)^2 G_0^{2\nu}, \quad (3.17)$$

where the leading term  $G_0^{2\nu}$  contains the phase-space integral over the electron and neutrino kinematics, and constants of the theory.

Expressed as a differential rate with respect to the most relevant experimental observable, the total electron energy  $\varepsilon$ , this can be written as

$$\frac{d\Gamma}{d\varepsilon} = g_A^4 \left( M_{GT}^{(-1)} \right)^2 \frac{dG_0^{2\nu}}{d\varepsilon}. \quad (3.18)$$

Importantly, the rate normalized differential expression

$$S(\varepsilon) = \frac{1}{\Gamma} \frac{d\Gamma}{d\varepsilon}, \quad (3.19)$$

at the leading order in the leptonic energy expansion has factored out all uncertainties in the NME. To here this has been a summary of existing calculations in the literature for the observable electron energy in  $2\nu\beta\beta$  [8, 20, 21]. The main results of the EFT framework have been to systematically parametrize and control the introduction of nuclear physics uncertainties in a low-energy weak test. To this end, it is possible to start considering how new physics beyond the Standard Model might manifest in the low-energy regime and perturb this picture. However, as experimental precision increases it is important to consider as well the expected Standard Model terms beyond the simplifying assumptions made so far.

### 3.2 NLO Corrections in Nuclear Current

In the non-relativistic reduction of the nuclear currents there is a NLO correction from induced weak magnetism (WM). In general, this term is a small correction, linear with  $q/m_n$ . Described here is the contribution from WM. This can be interpreted as a small perturbation in the nuclear current, linear in recoil order, with a vector current term proportional to  $g_M$

$$J_N^\mu = g_V v^\mu - 2g_A S^\mu + \left[ \frac{ig_M}{m_N} \epsilon^{\mu\nu\alpha\beta} v_\alpha S_\beta q_\nu \right], \quad (3.20)$$

where  $g_M = 1 + \kappa_1 = 4.7$  is the magnetic coupling,  $v^\mu = (1, \vec{0})$  and  $S^\mu = (0, \frac{1}{2}\vec{\sigma})$  are the nucleon velocity and spin. As a clarifying point, the four-vectors used here  $v^\mu$  and  $S^\mu$  serve as book-keeping terms that encode the space-like and time-like components of the nuclear currents.

The resulting amplitude can now be expressed as:

$$\begin{aligned}
C_{2\nu} = & \frac{g_A^4}{3} \left\{ [(M_{GT}^K)^2 + (M_{GT}^L)^2 + M_{GT}^K M_{GT}^L] \right. \\
& \times \left[ 1 - \frac{2g_M}{3m_N g_A} \left( Q + 2m_e - (E_{e1} + E_{e2}) \frac{2E_{e1}E_{e2} - m_e^2}{E_{e1}E_{e2}} \right) \right] \\
& \left. \right\} + \dots
\end{aligned} \tag{3.21}$$

Here  $Q = E_{11} + E_{22} - 2m_e$  and the dots denote higher-order terms.

The corresponding decay rate can be calculated using the fully explicit phase space integrals:

$$\begin{aligned}
\Gamma = & \frac{(G_F V_{ud})^4}{8\pi^7 m_e^2} \int_{m_e}^{E_i - E_f - m_e} dE_{e1} \int_{m_e}^{E_i - E_f - E_{e1}} dE_{e2} \int_0^{E_i - E_f - E_{e1} - E_{e2}} dE_{\nu 1} \\
& \times E_{e1} p_{e1} E_{e2} p_{e2} E_{\nu 1}^2 E_{\nu 2}^2 \times F(E_{e1}, Z_f) \times F(E_{e2}, Z_f) \times C_{2\nu},
\end{aligned} \tag{3.22}$$

where the substitution is used  $E_{\nu 2} = (E_i - E_f - E_{11} - E_{e2})$ . Also, here, the phase space integrals depend on commonly used corrections  $F$  which are parametrized by the lepton energy and the final nuclear charge  $Z_f$ . These corrections, called Fermi functions (or sometimes Coulomb corrections) encode the electronic screening effects on the available phase space of the outgoing electrons.

This prescription, and further corrections including those from [21], are considered in [22]. Here only the corrections from WM are highlighted as a tractable example of increasingly precise corrections in the Standard Model formulation and how these corrections can alter the observable kinematics through modified phase space integrands. In this way it is possible to calculate the rate normalized differential spectra of interest with negligible introduction of nuclear matrix uncertainties.

Using this result — which is a novel calculation of the differential rate of  $2\nu\beta\beta$  — it is possible to now think of the stated goal of systematically parametrizing physics beyond the Standard Model.

### 3.3 Double Beta Decay beyond the Standard Model

An important way to parametrize new physics beyond the Standard Model is to consider effective operators that are non-standard, that is outside of the quark level  $V - A$  theory, and that distort the charged-current interactions. Several options are possible [13, 23], but it is interesting to focus here on tensor interactions of left-handed neutrinos to illustrate a specific point.

In the conventions of Ref. [13], this interaction is given by the Lagrangian

$$\mathcal{L}_{BSM} = -\frac{4G_F}{\sqrt{2}} V_{ud} \left\{ \epsilon_T \bar{e}_R \sigma^{\mu\nu} \nu_L \bar{u}_R \sigma_{\mu\nu} d_L \right\}. \quad (3.23)$$

The operator in Eq. (3.23) has a mapping to the Standard Model Effective Field Theory (SMEFT) [24]. In particular,  $\epsilon_T$  corresponds to semileptonic four-fermion operators in SMEFT. This interaction can be constrained in low-energy  $\beta$  decays [13].  $\epsilon_T$  can also be constrained at the LHC [13]. Following Reference [25], the 90% CL on this parameter are:

$$-1.2 \cdot 10^{-3} < \epsilon_T < 1.1 \cdot 10^{-3}, \quad (3.24)$$

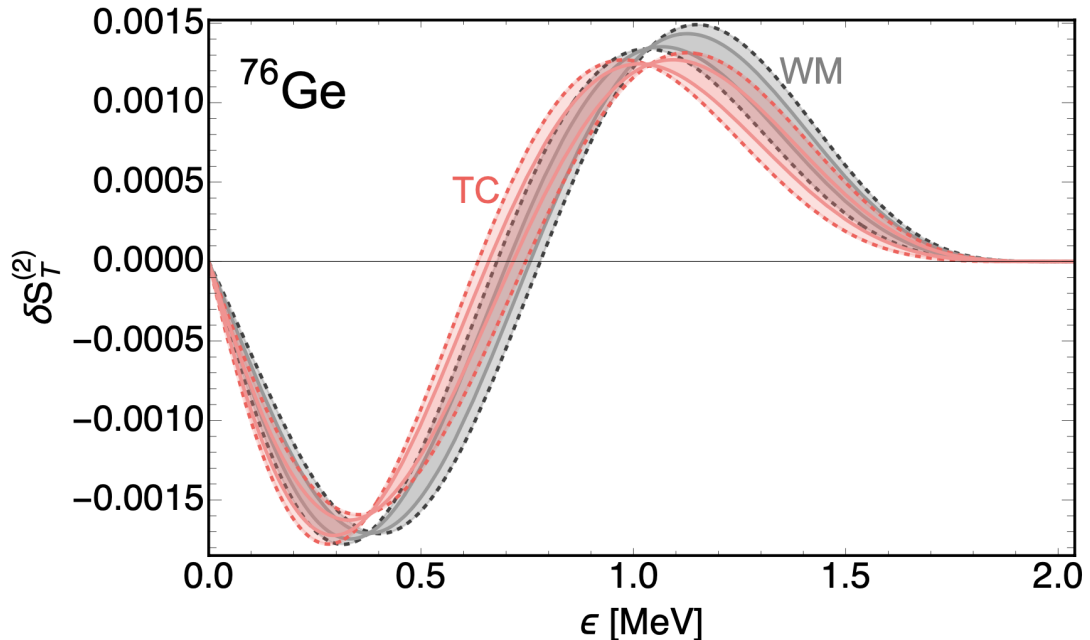
Following a similar procedure outlined above in the contraction of nuclear and leptonic currents, the leading order corrections to the transition amplitude are given by:

$$C_{2\nu}^{\text{BSM}} = \frac{8g_A^4}{3} \left\{ \epsilon_T \frac{g_T}{g_A} \left[ (M_{GT}^K)^2 + (M_{GT}^L)^2 + M_{GT}^K M_{GT}^L \right] \frac{m_e (E_{e1} + E_{e2})}{E_{e1} E_{e2}} \right\},$$

where the isovector tensor charge used is [26]

$$g_T = 0.989(34). \quad (3.25)$$

It is perhaps not surprising that the result in Equation 3.3 resembles the structure of Equation 3.21 as the Lorentz structure of the induced WM resembles the tensor current Lorentz structure. The amplitude of the distortion from WM is fixed by the lepton energies

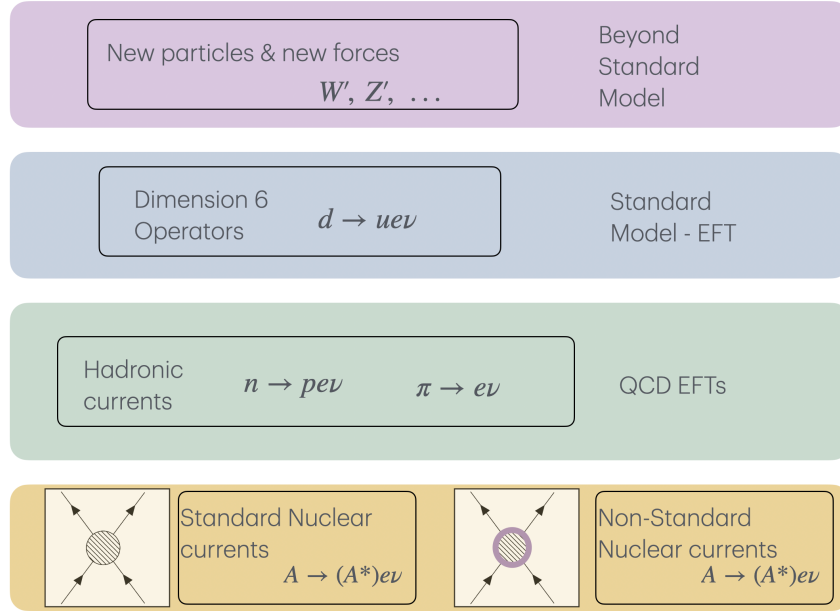


**Figure 3.1:** The induced distortions in the standard differential rate induced by the the BSM tensor current (TC) for  $\epsilon_T = -0.0014$  and weak magnetism (WM) for  $^{76}\text{Ge}$ . The shaded bands indicate choices of different nuclear matrix elements. Depending on the sign of  $\epsilon_T$  the direction of this distortion reverses. Figure reproduced from [22].

and size of the nucleus, while the amplitude of the distortion from the hypothetical tensor current depends on the coupling of the new physics to the Standard Model particles.

In Reference [22] this comparison was investigated explicitly and can be seen in Figure 3.1. What is emphasized in this result is that in the search for new physics which induces a low-energy effective tensor-current, there are scenarios where the new physics can mimic the spectral distortions from the subtle corrections of the expected nuclear currents.

In this way, weak magnetism is a natural “background” correction at the  $10^{-3}$  level, to this category of search for exotic physics. This makes precision calculations of laboratory weak processes as important as precision experiments.



**Figure 3.2:** Schematic illustration of the effective field theory hierarchy linking new high-energy interactions to measurable low-energy observables. Each layer integrates out heavy degrees of freedom, producing a set of effective operators that govern nuclear and leptonic processes such as double beta decay and CE $\nu$ NS.

### 3.4 Lessons for Other Weak Tests

In general, a main focus of this thesis is how increasingly precise laboratory tests of low-energy weak interactions can be used to search for physics beyond the Standard Model. The theoretical framework of EFT makes this a particularly tractable computational problem as new physics can be parametrized into effective operators in straightforward calculations. In the most promising tests, the inputs of the non-perturbative nuclear physics are factored into self-contained pieces with controllable uncertainties. Figure 3.2 illustrates schematically how information flows between energy scales in this framework, which the new physics at high energies manifests as corrections to the low-energy observables.

This chapter, and indeed the rest of the work in Reference [22], emphasize the importance of increasingly precise theoretical calculations to this end. The distortions from the new physics can appear through a variety of signatures in the low-energy regime, determined by the Lorentz structure of the effective operators. In cases where this new physics is subtle,

it is important to understand the corrections that can arise alongside these distortions from the detailed handling of the parametrized nuclear physics. This is the success of the EFT framework — the ability to simultaneously parametrize both new physics and difficult-to-calculate, non-perturbative effects in the same formalism.

The rest of this thesis focuses on an experimental program that uses neutrino scattering as a uniquely promising probe of the Standard Model. While the main thrust outlines the experimental methods and successes of neutrino scattering in constraining the Standard Model, this case study lays the blueprint for applications of experimental data in searches for new physics in the language of EFT.

# Chapter 4

## Coherent Elastic Neutrino Nucleus Scattering

In 1974 experimental evidence which indicated the existence of neutral currents in neutrino interactions was mounting [27]. That same year, a particular neutral current process involving neutrino interactions was proposed in a short, elegant paper that was inspired by an analogy to electron scattering [28].

The process was simple to describe and related the amplitudes of coherent forward scattering of low energy neutrinos off of a nucleus to key model parameters of the weak neutral current. This process, now referred to as CEvNS, remains one of the most powerful probes of the physics of the weak interaction.

For over forty years, access to this probe of fundamental physics was restricted by the technological barrier preventing the observation of the low-energy neutrino-induced nuclear recoils. The original author refers to the suggestion itself as an act of hubris, pointing out three major experimental concerns: the rates of interaction, the ability to detect the small signals, and the large backgrounds which could swamp the neutrino signals.

This intriguing process in the neutrino sector went largely unnoticed for decades, save a few curious individuals [29, 30]. All the while, the finer details of the weak interaction and neutrino phenomenology were continuously developed. Then, in 2017, the first experimental confirmation of the theorized process fulfilled its long-standing promise of being a uniquely sensitive probe of fundamental physics. But what makes CEvNS such a good probe? The common answer is that it has precise predictions in the Standard Model. This chapter outlines the formalism of the theory, what it means for this process to be well-predicted, and some technical challenges to realizing these predictions in the laboratory.

## 4.1 Theoretical Formalism

CEvNS is an elastic scattering process between a neutrino and a nucleus mediated by a  $Z^0$  boson. The interaction takes place for neutrino energies  $\lesssim 50\text{MeV}$ . At these energies, the formalism of effective theories is a useful way to parametrize both the Standard Model interactions and any potential deviations from the Standard Model.

The overall theoretical uncertainty in the prediction is small. Indeed, the hadronic and nuclear structure form factors enter at the per cent level. As a result, experimental measurements serve as a broadband search for new physics which can modify the interaction in numerous ways.

In the literature there are many precise calculations of the CEvNS cross-section [5, 10, 31]. Here, some of the key features of the calculations are summarized with particular attention to some of the standard simplifying assumptions and the level of precision achievable.


### 4.1.1 Calculation

In the relevant terrestrial neutrino scattering experiments, calculations are performed in the laboratory frame where the nucleus is at rest. Following the notation in [10], the incoming neutrino is described by the initial four-momentum  $k = (E_i, \vec{k}_i)$ . The nucleus is described by the initial four-momentum  $p = (M_A, \vec{0})$ , where  $m_A$  is the rest mass of the nucleus. The scattered neutrino has four-momentum  $k' = (E_f, \vec{k}_f)$  and the recoiling nucleus has four-momentum  $p' = (M_A + T, \vec{p}_A)$ , where  $T$  is the recoil energy of the nucleus. Allowing  $q^2 = |\vec{k}_f - \vec{k}_i|^2$  to be the square of the three-momentum transfer, it is possible to express the recoil energy  $T = q^2/2M_A$ . Schematically, this is represented in Figure 4.1.1.

Differential rates for particular observables can be extracted from the following form [10]

$$d^6\sigma = \frac{1}{|\vec{v}_i - \vec{v}_A|} \frac{m_i m_f}{E_i E_f} \frac{d^3\vec{k}_f}{(2\pi)^3} \frac{M_A}{M_A + T} \frac{d^3\vec{p}_A}{(2\pi)^3} (2\pi)^4 \sum_{fi} \overline{|\mathcal{M}|^2} \delta^{(4)}(k_i + p_A - k_f - p'_A), \quad (4.1)$$

where the summation runs over the initial and final spin states, and the squared amplitude



**Figure 4.1:** Diagram of the tree-level CEvNS interaction in the effective theory in the laboratory frame. An incoming neutrino scatters off of a nucleus at rest. The recoiling nucleus acquires energy  $T = q^2/2M_A$ , where  $q$  is the three-momentum transfer and  $M_A$  is the rest mass of the nucleus.

can be expressed as the full contraction of the lepton and nuclear tensors

$$\sum_{fi} |\mathcal{M}|^2 = \frac{G_F^2}{2} L_{\mu\nu} W^{\mu\nu}. \quad (4.2)$$

Following the prescriptions outlined in Chapter 2, it is possible to relate the nuclear tensor  $W^{\mu\nu}$  to a series of successive projections. The nuclear current is written as the projection over the hadronic current,

$$\mathcal{J}^\mu = \langle \mathcal{N}(p_f) | \eta (\gamma^\mu (1 - \gamma_5)) \eta | \mathcal{N}(p_i) \rangle \quad (4.3)$$

$$= \bar{\mathcal{N}} (\mathcal{F}_V^\eta(q^2) \gamma^\mu + \mathcal{G}_A^\eta(q^2) v^\mu) \mathcal{N}, \quad (4.4)$$

where  $\eta$  is a nucleon state, either a proton or a neutron, and  $\mathcal{F}_V(q^2)$  and  $\mathcal{G}_A(q^2)$  are the vector and axial-vector nuclear structure form factors respectively. The four-vector  $v^\mu = (0, \vec{\sigma})$  represents the Lorentz structure of the spin-dependent axial-current to leading order in recoil expansion.

Considering only the leading order terms in the non-relativistic reduction of the nuclear currents, it is possible to express the full differential rate. While in principle it is possible to consider the differential rate with respect to the recoil angle, CEvNS experiments to-date

are only sensitive to the differential rate with respect to the total recoil energy  $T^1$ . This is written as

$$\frac{d\sigma}{dT} = \frac{G_F^2 M}{2\pi} \left[ (G_V + G_A)^2 + (G_V - G_A)^2 \left(1 - \frac{T}{E_\nu}\right)^2 - (G_V^2 - G_A^2) \frac{MT}{E_\nu^2} \right] \quad (4.5)$$

$$G_V = g_V^p Z \mathcal{F}_V^p(q^2) + g_V^n N \mathcal{F}_V^n(q^2) \quad (4.6)$$

$$G_A = g_A^p (Z_+ - Z_-) \mathcal{G}_A^p(q^2) + g_A^n (N_+ - N_-) \mathcal{G}_A^n(q^2). \quad (4.7)$$

In the literature, a few more common simplifying assumptions are made. First, the separate proton and neutron dependent structure factors are also taken to be equal. The nuclear structure form factors for the vector and axial vector components are also taken to be equal. This allows for the definition of the weak nuclear charge  $Q_W$  to be

$$Q_W = g_V^p Z + g_V^n N, \quad (4.8)$$

so that the vector current and axial-vector current terms can be re-written as

$$G_V = Q_W \mathcal{F}(q^2) \quad (4.9)$$

$$G_A = [g_A^p (Z_+ - Z_-) + g_A^n (N_+ - N_-)] \mathcal{F}(q^2). \quad (4.10)$$

It is possible to see that the axial-axial term is sub-dominant and, beyond the vector-vector terms, only the interference term contributes to leading-order. It is possible then to re-write the differential rate into a spin-independent (SI) component and a spin-dependent (SD) component as

$$\frac{d\sigma}{dT} = \frac{G_F^2 m_A}{4\pi} \left(1 - \frac{m_A T}{2E_\nu^2} - \frac{T}{E_\nu}\right) Q_W^2 |\mathcal{F}(q^2)|^2 + \frac{G_F^2 m_A}{4\pi} \left(1 + \frac{m_A T}{2E_\nu^2} - \frac{T}{E_\nu}\right) \mathcal{F}(q^2), \quad (4.11)$$

<sup>1</sup> There are interesting efforts to develop experiments that are sensitive to the recoil angle [32].

where the SI vector component is proportional to  $Q_W^2$  and the SD interference term is proportional only to the combined spin of the nucleus. If considering a nucleus whose ground state is a  $0^+$  spin state, or for heavy nuclei where the  $Q_W^2$  term dominates, it is possible to finally recover the familiar expression

$$\frac{d\sigma}{dT} = \frac{G_F^2 m_A}{4\pi} \left( 1 - \frac{m_A T}{2E_\nu^2} - \frac{T}{E_\nu} \right) Q_W^2 |\mathcal{F}(q^2)|^2. \quad (4.12)$$

### 4.1.2 CEvNS as a Probe

Due to the simple form that CEvNS takes in Eq. 4.1.1, it is possible to probe a broad range of physics through measurements of the differential rate across nuclear targets. Two examples here are illustrated to highlight the ability of this process to probe both Standard Model physics and physics beyond the Standard Model.

#### Nuclear Structure Form Factors

The nuclear structure form factors which appear in Eq. 4.1.1 encode the distribution of nucleons inside of the nucleus. In practice they are the Fourier transform of the neutron distribution given by:

$$\mathcal{F}(q^2) = \frac{1}{Q} \int e^{-i\mathbf{q}\cdot\mathbf{r}} \rho(\mathbf{r}) d^3r, \quad (4.13)$$

where  $\mathbf{r}$  is the position of nucleons (in this case neutrons),  $\rho$  is the spatial distribution, and  $Q$  is the normalizing charge of the total distribution.

Under assumptions of spherical symmetry there are common parametrization of the neutron distributions [33, 34]. More recently there have been more explicit calculations of the microscopic interactions of the nucleus to model the neutron distributions [35].

As the differential rates depend on these form factors, precise measurements of CEvNS can reveal important information about the nuclear structure of target nuclei.

## Non-Standard Interactions

The theoretical framework of EFT allows CEvNS to be a powerful probe of physics beyond the Standard Model. A popular example is to consider a new, neutral current whose mediator is a new heavy boson [36, 37]. In these models, it is possible to distort the standard vector couplings through the form:

$$\mathcal{L}^{\text{NSI}} = -2\sqrt{2}G_F \sum_{\alpha,\beta=e,\mu,\tau} \bar{\nu}_\alpha \gamma^\mu \nu_\beta \sum_{f=u,d} \varepsilon_{\alpha\beta}^{fV} (\bar{f} \gamma_\mu f), \quad (4.14)$$

where  $\varepsilon$  parametrizes the strength of the deviation from the standard coupling. If this non-standard neutral-current is related to an undiscovered particle of mass  $m_X$  much greater than the electroweak scale, and coupling  $g_X$  then the distortion parameter  $\varepsilon$  is related to the ratio  $m_X/g_X$ . This way, the effective four fermion interaction parametrizes searches for physics which can be, in general, beyond the reach of the largest colliders. Precision measurements of the differential rate can be sensitive to these subtle deviations and complement more direct attempts at producing this mediator in the laboratory.

### 4.1.3 Corrections

Throughout this summary of the CEvNS derivation, many simplifications were made. The central question of this work is to compare experimental measurements in the weak interaction to predictions in this framework. To this end, it is necessary to understand the error introduced in the prediction from some of the major simplifications.

#### Radiative Corrections

Up to now, only the tree-level amplitude was considered in the differential rate. There are calculations in the literature of the impacts of electromagnetic radiative corrections to the differential rate [38].

Qualitatively, one of the most important phenomena this correction introduces neutrino flavor dependence in the differential rate. This is because the one-loop diagrams that enter

the calculation at the next to leading order depend on the fermion masses, and thus are different for the species of neutrinos.

In the framework of EFT, Reference [31] uses a similar four-fermion interaction picture, where they consider exclusively a  $0^+$  nucleus, to arrive at a corrected differential cross-section

$$\frac{d\sigma_{\nu\ell}}{dT} = \frac{G_F^2 M_A}{4\pi} \left(1 - \frac{T}{E_\nu} - \frac{M_A T}{2E_\nu^2}\right) \left[Q_W^2 \mathcal{F}_W(Q^2) + \frac{\alpha}{\pi} (\delta^{\nu\ell} + \delta^{\text{QCD}}) F_{\text{ch}}(Q^2)\right]^2, \quad (4.15)$$

where  $F_W(q^2)$  is the weak nuclear form factor and  $F_{ch}(q^2)$  is the electromagnetic nuclear form factor. The corrections  $\delta^{\nu\ell}$  and  $\delta^{\text{QCD}}$  encode the dynamics of the flavor dependent and hadronic electromagnetic radiative corrections.

Part of the flavor-dependent term  $\delta^{\nu\ell}$  can be interpreted as the neutrino charge radius, which is related to the first moment of the modified weak nuclear form factor, and, overall, the corrections are of the order  $\mathcal{O}(1\%)$  [31].

### **Recoil Order Corrections**

Similar to the work outlined in the charged-current interaction, the nuclear currents can be decomposed in a non-relativistic expansion that include induced Lorentz structures beyond the typical  $V - A$  interactions. In the most general expression, the nuclear current will contain: induced weak-magnetism induced pseudo scalar, higher order axial, and higher order vector corrections. At next-to-leading order the most relevant correction is weak magnetism at the  $\mathcal{O}(0.1\%)$  level. Higher order terms come in at a smaller level.

### **Hadronic Current Corrections**

In the construction of the hadronic currents, only the lightest quark degrees of freedom were considered. Strange quark contributions are next-leading-order corrections in the hadronic form factors. These corrections can be written as [38]

$$g_V^p = g_V^p(0) + \delta. \quad (4.16)$$

In general, this effect is of the order of  $\mathcal{O}(0.1\%)$ . This can be accounted for entirely by a shift in the hadronic form factors.


Other corrections to the hadronic currents can be considered. The general recipe for constructing hadronic currents given in Chapter 2 involved projections over sums of valence quarks. It is possible to consider higher-order corrections from strange quark contribution which perturb the hadronic form factors [39]. These corrections depend on the size of the nucleus and can be at most of order  $\mathcal{O}(1\%)$  for lighter nuclei.

#### 4.1.4 Challenge 1: Interaction Rates

As Freedman points out in his original paper, as a weak neutral current process the total interaction rates are small. To illustrate this point, it is useful to consider a few concrete examples. For typical neutrino energies and a few common target nuclei argon (Ar), cesium (Cs), and germanium (Ge) it is possible to calculate the differential rates in recoil energy. The specific neutrino energies used in this example range from 1 to 50 MeV and follow a common distribution from a  $\pi$  decay at rest source described in further detail in Chapter 5. Here, the Klein-Nystrand nuclear form factor model is assumed [34] for the calculation.

Figure 4.2 shows graphically the differential recoil rates for each target. Integrating the total differential rates for each target gives the following total rates:  $\langle \sigma_{Ar} \rangle \sim 1.3 \times 10^{-39} \text{ cm}^2$ ,  $\langle \sigma_{Ge} \rangle \sim 5.2 \times 10^{-39} \text{ cm}^2$ ,  $\langle \sigma_{CsI} \rangle \sim 16.1 \times 10^{-39} \text{ cm}^2$ . As a general comparison, the typical interaction rates for the comparatively common neutron scattering is  $10^{-24} \text{ cm}^2$  [40].

This feature of low rates of interaction is not unique to CEvNS among peer-neutrino interactions. Indeed, it is a defining feature of the weak nuclear force and is responsible for the long half-lives of weak-decaying matter. Compared to other low-energy neutrino interactions, CEvNS benefits from an enhancement in the rate, from the weak nuclear charge  $Q_W^2 \propto N^2$ . This feature is evident in the strong amplification in rate in heavier nuclei, where the coherence is most prominent.



**Figure 4.2:** An example of differential rates of CEvNS on three target nuclei. The black line shows the recoils on Ge, red line shows the recoils on Ar, and blue line shows the recoils on Cs. The kinematic endpoint of the recoil spectra depends on the mass of the nucleus. Figure reproduced from [41].

#### 4.1.5 Challenge 2: Recoil Energy Thresholds

The second major obstacle Freedman points out is the difficulty in detecting the feeble signals when they do occur. In Figure 4.2 it is possible to see that to even be able to observe these signals it is necessary to detect nuclear recoil energies on the scale of  $\mathcal{O}(\text{keV})$ .

From equation 4.1.1 it is possible to see that the maximum recoil energy is given by

$$T_{\max} = \frac{2E_{\nu}^2}{2E_{\nu} + m_A}. \quad (4.17)$$

As a result, the heavier nuclei which benefit from the rate enhancement as penalized in recoil energy.

To characterize the challenge of observing recoil energies on this scale it is necessary to consider, briefly, detection mechanisms. In the broad field of radiation detection there are different strategies to reconstruct energy deposited in the form of small nuclear recoils. It is worth emphasizing that in CEvNS the quantum state of the nucleus is unperturbed and the only mechanism for observing the nuclear recoil energy are the secondary forms of radiation and their subsequent impacts on a material. Popular mechanisms include scintillation [42, 43], heat [44, 45], and ionization [46, 47], though more exotic detection mechanisms which rely on phase transitions or material defects also exist [48].

For nuclear recoils in an excellent scintillating material, such as room-temperature sodium doped cesium iodide (CsI[Na]), the detectable light yield from a typical nuclear recoil in a CEvNS experiment is on the order of  $\mathcal{O}(10)$  photo-electrons).

#### 4.1.6 Challenge 3: Backgrounds

The final experimental challenge that Freedman considers is the presence of background radiation that can mimic and swamp the desired signal of neutrino-induced nuclear recoils. The particular background that Freedman considers is the possibility of neutron-induced nuclear recoils dominating in an experimental setting.

As seen above, neutron scattering rates are comparatively abundant. The observed recoils between neutrons and neutrinos are difficult to distinguish. Without additional kinematic information, such as the ability to reconstruct the recoil angles as well as energy, it is in fact impossible. Neutrons can also be present in large quantities in the experimental settings aimed at measuring CEvNS.

## 4.2 Summary

Following the theoretical formalism outlined in Chapter 2 the details of the calculation of differential rates of the CEvNS interaction have been summarized. Where possible, particular attention was given to simplifying assumptions in the literature and the magnitude of the careful corrections that have been studied at length. In general, the sources of theoretical uncertainties in the prediction come from the inputs of non-perturbative QCD and the microscopic physics of intra-nucleon and intra-nucleus interactions. These are summarized into form factors which can introduce as much as  $\mathcal{O}(1\%)$  uncertainty in the prediction.

This level of precision, along with the broad class of mechanisms that can contribute to or distort the predicted differential rates makes CEvNS an excellent probe for new physics. Different applications of this interaction were highlighted, such as the search for exotic interactions and the study of nuclear structure.

The experimental challenges in observing this interaction were outlined. To benefit from the rich physics program enabled by CEvNS measurements, these challenges must be overcome.

# Chapter 5

## The COHERENT Collaboration at the Spallation Neutron Source

In 2017, a little more than forty years after its original proposal, CEvNS received its first experimental verification by the COHERENT collaboration. In the eight years since this discovery there have been a half-dozen successive measurements around the world on different targets and using different sources of neutrinos. As a historical matter, it is worth asking: what can explain this forty-year drought, which finally broke to give way to the present wealth of measurements? Much of the credit belongs to the decades-long, careful and methodical development of radiation detection techniques in the related field of dark-matter direct detection.

The second leg of the answer to this historical question is related to the problem of identifying an advantageous source of neutrinos. The discovery of CEvNS by the COHERENT collaboration was enabled by the fruitful exploitation of a uniquely good neutrino production facility, the Spallation Neutron Source (SNS). This chapter summarizes the key features of the neutrino production at the SNS and the success to date of the COHERENT collaboration in using this source to study neutrino scattering.

### 5.1 The Spallation Neutron Source

The SNS, located at Oak Ridge National Lab (ORNL), is a Basic Energy Science (BES) user facility aimed at producing an intense, pulsed source of neutrons to enable materials research and technology development. The facility has been in operation since 2006.

The SNS consists of three main components to achieve the desired neutron production: a linear accelerator, a proton accumulator ring, and a target station. In the linear accelerator component,  $H^+$  ions are accelerated to 1.3 GeV. Inside the accumulator ring, protons are

bunched into pulses with a characteristic width of approximately  $\sim 400$  ns. Once bunched, the proton pulses reach the target station at a repetition rate of 60 Hz. At these energies and at this rate, the SNS is able to deliver an average power of 1.8 megawatt (MW).

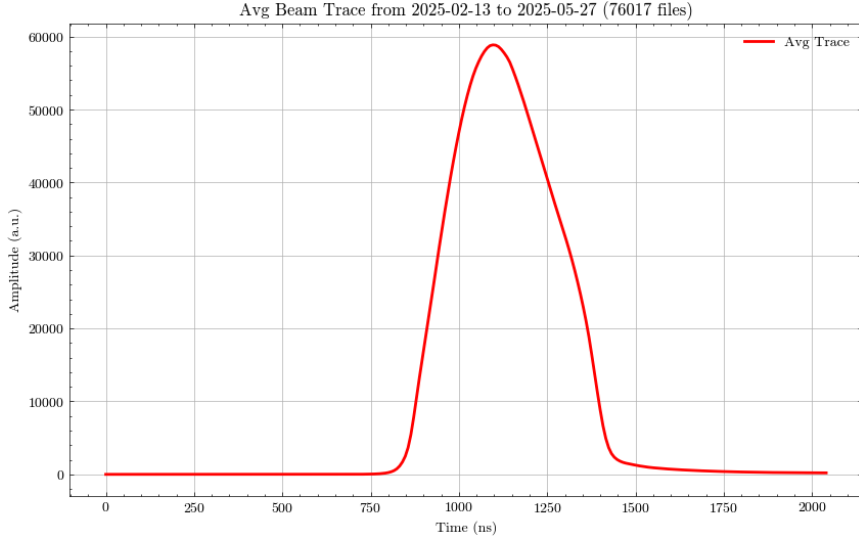
It is important to highlight here that the energies and power quoted here represent the current capabilities of the facility. The work presented in this thesis straddles both sides of a major SNS upgrade, referred to as the Proton Power Upgrade [49]. Prior to this upgrade protons were delivered with 1.0 GeV energy and power around 1.3 MW. There are plans to continue increasing the beam power during operation to 2.0 MW in the years to come.

Inside the target station protons strike a liquid mercury target. At these proton energies, collisions with a mercury nucleus can result in several different interaction processes. The dominant process is called “spallation”, and refers to a sequence of events that take place inside of a nucleus resulting in the emission of many neutrons. Spallation is a process mediated by the strong nuclear force, which, in contrast to weak interactions, happens nearly instantaneously. As a result, the timing distribution of the creation of neutrons is characterized entirely by the timing distribution of the protons striking the target — that is, the width of the proton bunch. An average of the measured current from proton bunches is shown in Figure 5.1, highlighting this timing structure.


At the SNS, the neutrons created in the target are moderated down in energy by heavy water shielding. This broadens both the energy and timing distributions of neutrons produced, which are then guided towards user experiments in beam-lines and experiment halls. Unfortunately, the neutrons are still diffusely present in the facility during operations.

## 5.2 $\pi$ Decay at Rest

In addition to neutron spallation, protons striking the mercury can produce pions inside the target. At these energies, the majority of the charged pions produced in the mercury lose their energy to ionization before coming to rest inside the target. While the negatively charged pions overwhelmingly capture on mercury nuclei, the positively charged pions will decay at rest according to the decay scheme:



**Figure 5.1:** A recent example of the average current measured from protons on target at the SNS. The characteristic width of the beam pulses is approximately 400 ns. This narrow pulse is one of the central features that makes the SNS an ideal facility for neutrino science.



The process of hadronization which creates the pion, and the subsequent ionization that bring the pion to a rest before decaying, are all fast compared to time-scales relevant at the SNS. While the half-life of the pion is 26 ns, there is a continuous production of pions during the several hundred nanoseconds of bunched proton current. This timing structure produces an isotropic prompt source of muon neutrinos. Due to the kinematics of the two-body decay the muon neutrinos are mono-energetic with energy 29.9 MeV.

The muons produced from the pion decay also come to a stop in the mercury before decaying at rest into a positron, electron neutrino, and muon anti-neutrino. Muon decay has a characteristic half-life of 2.2  $\mu$ s, and as a result, the neutrinos produced from the muon decay are delayed. The resultant neutrino energy distributions are governed by the kinematics of the three-body decay. The range of neutrino energies is then from 0 to 52.85 MeV, or half the rest mass of the muon.

The decays described here are well approximated by the analytic functions:

$$\Phi_{\nu_\mu}(E_{\nu_\mu}) = \frac{2m_\pi}{m_\pi^2 - m_\mu^2} \delta\left(1 - \frac{2E_{\nu_\mu}m_\pi}{m_\pi^2 - m_\mu^2}\right), \quad (5.2)$$

$$\Phi_{\nu_e}(E_{\nu_e}) = \frac{192}{m_\mu} \left(\frac{E_{\nu_e}}{m_\mu}\right)^2 \left(\frac{1}{2} - \frac{E_{\nu_e}}{m_\mu}\right), \quad (5.3)$$


$$\Phi_{\bar{\nu}_\mu}(E_{\bar{\nu}_\mu}) = \frac{64}{m_\mu} \left(\frac{E_{\bar{\nu}_\mu}}{m_\mu}\right)^2 \left(\frac{3}{4} - \frac{E_{\bar{\nu}_\mu}}{m_\mu}\right). \quad (5.4)$$

There are several sub-dominant contaminants to this otherwise highly pure source of  $\pi$ -DAR neutrinos. Inside of the target, not all negatively charged pions are captured on mercury. Less than 1% will also decay at rest, creating a charge inverted decay chain. Less than 1% of the positively charged pions will also decay in flight before coming to rest. The kinematics smeared by the energy of the pion, neutrinos produced from this decay can have energies as high as several hundred MeV.

The collaboration has modeled the neutrino flux production at the SNS extensively [50] to account for these contaminants and also to estimate the normalization of the neutrino flux. While the decay processes that produce the neutrinos are very well understood, the hadronization processes from intra-nuclear cascades depend on (once again) non-perturbative inputs from the strong force. Modeling the pion production is done through the use of phenomenological nuclear cascade models which have large uncertainty on their normalization. As a result, while the energy and timing distributions of neutrinos at the SNS is well predicted, the total normalization is known only to within ten percent.

Figure 5.2 shows the result of the neutrino flux modeling at the SNS. The energy spectra extends to relatively large energies, though the dominant contributions from  $\pi$ -DAR have an average neutrino energy of around 30 MeV. The timing distribution is shown after convolution with the measured current of the protons on target, to account for the extended pion production. It is possible to see the flavor decomposition clearly in time, where it is often considered as a prompt and delayed component.

There are several features of this neutrino source that specifically enabled the discovery of CEvNS. Most prominently, the timing structure of the neutrino arrival allows for a remarkable



**Figure 5.2:** Dominant components of the neutrino flux probability distribution function from a Geant 4 simulation performed by Rebecca Rapp [51]. The projection in time (left) is the instantaneous neutrino production before any convolutions with the proton trace.


reduction in backgrounds. Neutrinos are only produced in coincidence with proton current on target. By considering only 10  $\mu$ s windows around each of the 60 proton bunches that arrive in one second, there is an effective reduction in the steady-state environmental backgrounds by a factor of 1,667. Meanwhile, over 90% of the total neutrino flux is retained in this window, as can be seen in Figure 5.2.

### 5.3 Neutrino Alley

To date, all of the detectors deployed by the COHERENT collaboration have their home in a sub-basement service corridor dubbed Neutrino Alley. Neutrino Alley is a narrow corridor located approximately 20-30 m away from the mercury target. Figure 5.3 shows a cross-sectional engineering drawing of the facility and the location of Neutrino Alley relative to the SNS target.

In between the Neutrino Alley and the target are several meters of concrete and gravel backfill that provide orders of magnitude of reduction of beam-related neutrons.

The neutrino flux at a typical location in Neutrino Alley, a baseline of 20 m, can be estimated from existing modeling of the neutrino flux [50]. Over the course of a typical year




**Figure 5.3:** (Left) Engineering drawing of the target. (Right) Cross-sectional view of the SNS target station that houses the mercury target. Highlighted in cyan is the concrete monolith, and in purple is a sliced view of Neutrino Alley. In between the two is gravel backfill. This Figure is reproduced from [51]

of operation at the SNS (after the PPU this now means 1.3 GeV and 1.8 MW, though this will continue to increase), there is a total flux of approximately  $3 \times 10^{10} \nu/\text{cm}^2$ .

Considering Neutrino Alley’s proximity to an abundant neutrino source, the sharp timing structure of the beam, and the surprisingly low level of beam-related neutron backgrounds, it is a privileged niche for the study of CEvNS. There is a modest engineering price to pay for these advantages. Every detector deployed to Neutrino Alley must fit within narrow physical constraints to allow for open access in the facility. Use of cryogenically cooled detectors is also constrained by the small spaces available. The detectors deployed to this location have changed many times of the collaborations decade of operation. Figure 5.4 shows the current schematic of detectors operating.

## 5.4 Backgrounds

Despite the background reduction from the timing of structure of the beam, environmental steady-state backgrounds must still be considered. In particular, cosmogenic sources can still create important levels of backgrounds. Overburden, typically expressed in units of




**Figure 5.4:** Schematic showing the currently (and soon-to-be) operating neutrino detectors in the COHERENT collaboration. Not all detectors are searches for CEvNS. Several detectors deployed are part of the growing effort to measure inelastic scattering processes. The Ge-Mini detector system, the subject of this work, is shown in the center of the Alley.

meters water-equivalent (m.w.e.), can vary throughout the facility. All the locations in the facility are shallow, around 10 m.w.e.. Other standard uncorrelated backgrounds include gamma rays from natural levels of environmental radio-isotopes.

While these uncorrelated, steady-state backgrounds are greatly reduced by the timing structure of the beam, it is still necessary to carefully consider the correlated, or beam-related, backgrounds. The most prominent beam-related backgrounds come from the fast neutrons produced in spallation. These neutrons, often referred to as beam-related neutrons to distinguish from neutrons produced from cosmic backgrounds, are produced in coincidence with the prompt neutrino component. In sufficient quantity, this neutron background can spoil any attempt to look for neutrino signals.

Inside the SNS, construction material can provide natural shielding of neutrons, while corridors and line of site access to the target area can create natural funnels for diffuse neutron propagation. While some work has been done to model the propagation of neutrons




**Figure 5.5:** The flux of neutrons was measured at several locations in the SNS that are relevant to existing and potential neutrino measurements. The shaded regions correspond to the approximate areas of concrete/gravel backfill and correspond to the greatest neutron shielding. The locations of existing CEvNS measurements are highlighted with vertical bands. This figure is reproduced from [51].

in the facility, the greatest success has been found with extensive measurement campaigns that map the variable neutron density throughout the facility. Figure 5.5, reproduced from the doctoral work of a COHERENT collaborator [51], shows some of these efforts to map out the neutron flux in a particular corridor inside the facility. The areas with the lowest neutron flux in turn have the lowest backgrounds.

## 5.5 Existing Measurements

The COHERENT collaboration has made several measurements of low-energy neutrino scattering, which include both the first measurements of CEvNS on different targets, but also include several inelastic neutrino scattering processes. These existing measurements are briefly summarized.



**Figure 5.6:** The full on-beam CsI[Na] dataset is shown as a residual against the best fits of the CEvNS signal (decomposed by flavor) and neutron backgrounds. Figure is reproduced from [42].

### 5.5.1 CEvNS

#### CsI

In 2017 the COHERENT collaboration announced the first observation of CEvNS using a 14.6 kg CsI[Na] detector [41]. The CsI[Na] detector was deployed to an optimal location in Neutrino Alley, at a short baseline of 19.3 m and boasting the lowest neutron backgrounds that the collaboration could find. This first result gave a  $6.7\sigma$  rejection of the no-CEvNS hypothesis, with a best fit of  $134 \pm 22$  signal events.

In 2021 the collaboration followed-up this result with an improved measurement [42]. Combined, these results were a  $11.6\sigma$  rejection of the no-CEvNS hypothesis. The results of the full measurement, displayed as the number of counts as function of the arrival time of the neutrinos is shown in Figure 5.6.

This discovery ushered in an era CEvNS measurements at the SNS.

## Ar

In 2021 the COHERENT collaboration announced a measurement of CEvNS on argon using a detector-system called CENNS-10, which housed 24 kg of atmospheric argon [43]. This scintillation based detector was located in a less advantageous location, 27 m away from the target and flooded with larger neutron backgrounds. The smaller nucleus compared to CsI, also means a significant reduction in the rate of CEvNS. Despite this the detector system was able to observe  $159 \pm 43$  neutrino events which contended with  $553 \pm 34$  beam-related neutron events and  $3,131 \pm 23$  steady-state background counts. This result was a  $3.9\sigma$  rejection of the no-CEvNS hypothesis.

This result marked an important step for the collaboration in probing different targets.

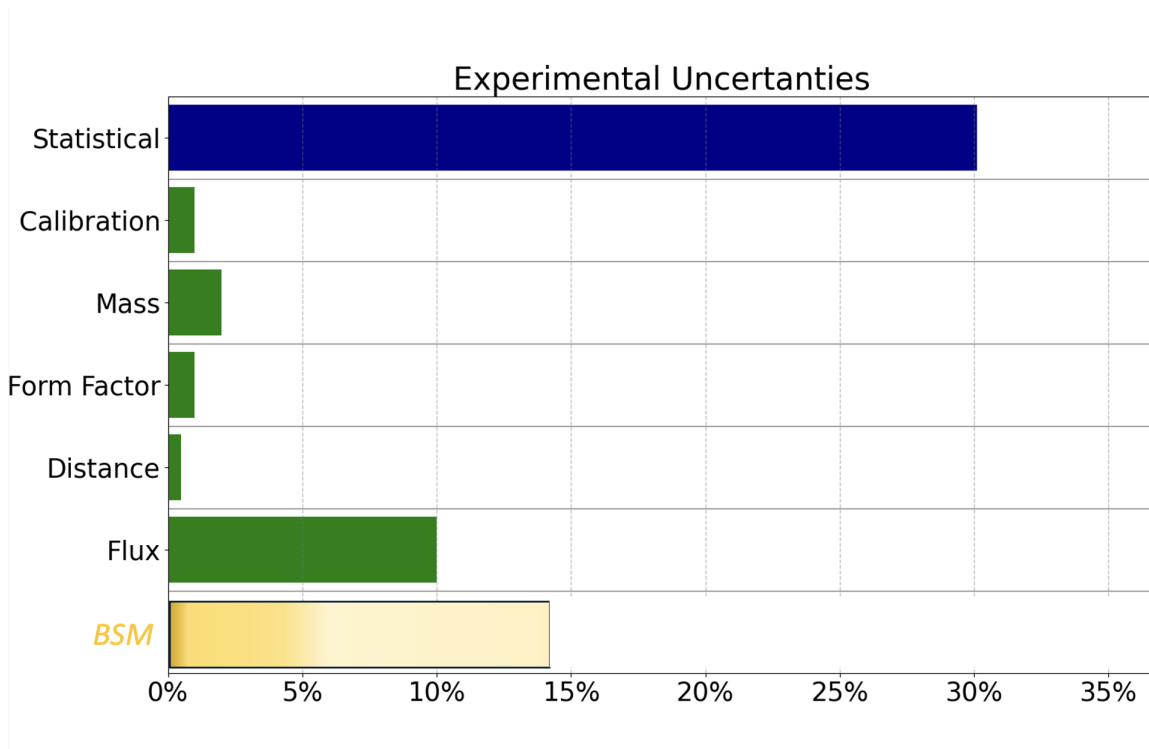
## Ge

In 2023 the collaboration announced a measurement of CEvNS on germanium using an array of High Purity Ge (HPGe) detectors [46]. This work, and its follow-up measurement, form the experimental basis of this thesis.

## 5.6 Experimental Sources of Uncertainty

To conclude this chapter it is important to review these existing measurements at the SNS and consider the level of precision achieved by each result. In particular, what is emphasized here is uncertainty on measurements of the total, flux-averaged cross-sections. Figure 5.7 highlights the experimental error budget in the latest COHERENT CEvNS measurement on germanium. This first result on germanium, as will be discussed in the text, was statistically limited. The ability to use these measurements of CEvNS to explore the physics of the weak interaction depends, broadly speaking, on the precision of the measurements.

Experimental uncertainties are broken into statistical uncertainties and systematic uncertainties. Both contribute to the overall sensitivity of these measurements, and to their potential application to the many physics goals outlined in chapter 4. The total uncertainty



**Figure 5.7:** Graphical representation of the error budget for the first germanium result by COHERENT. The statistical uncertainties is shown in blue, and dominant. Systematic uncertainties are shown in green and are mostly comprised of the flux normalization uncertainty. Physics beyond the Standard Model is graphically represented as a graded bar in yellow, representing the reduction of available parameter space in new models constrained by increasingly precise measurements - or, alternatively, waiting to be discovered at some unknown precision.

of CsI[Na] detector system on the CEvNS cross-section of 14%. The argon detector system achieved an uncertainty of 30% and the first germanium result had an uncertainty of 35%. All of these are, of course, far larger than the ultimate promise of the percent level uncertainty in the theoretical predictions. It is important to understand the current gap in experimental precision and theoretical precision and how this work aims to address this gap.

Statistical uncertainty is a clear limitation of current CEvNS experiments. The statistical uncertainty of measurements depends on the counting statistics of both the signal and the background. Experiments need to balance both exposure and background-level to maximize the ability to make precise measurements. This work in particular shows how improvements in both exposure and background-level can have dramatic impacts on the sensitivity of a

CEvNS experiment to physics models of interest.

Within the category of systematic uncertainties, there are several sources that can contribute. The largest source of systematic uncertainty in the germanium result is the uncertainty on the total neutrino production at the SNS, at 10%. This uncertainty is shared by all experiments deployed to the SNS, and must be addressed in future measurements.

The collaboration strategy to mitigate this uncertainty is to use an alternative neutrino detection mechanism to simultaneously constrain the neutrino flux uncertainty [52]. A heavy-water-based CC neutrino detector is deployed to measure the similarly well-predicted charged current cross-section on deuterium. The goal is to use this auxiliary measurement at the SNS to reduce the uncertainty on the flux down to three percent.

There are several sub-dominant sources of experimental uncertainty that also contribute the precision of the cross-section measurement. These sources of uncertainty are often detector-specific and reflect the current technological limitations of each measurement. For the germanium detectors this includes the uncertainty in total active mass of the detector, that is, how many nuclear targets are available in the scattering experiment. This uncertainty is only controlled at the percent level.

In summary, the experimental precision in CEvNS is currently trailing the theoretical precision. Indeed, CEvNS predictions are comfortably in the NLO, which can give single percent level theoretical uncertainty or better. Meanwhile, the best experimental measurements are still in the double digit percent uncertainty. This thesis demonstrates a promising path to narrow this gap and achieve an experimental precision worthy of careful calculations.

# Chapter 6

## The Ge-Mini Detector System


The Ge-Mini system was developed to advance precision measurements of CEvNS using High Purity Ge (HPGe) detector technology. Germanium detectors have long been a workhorse in the rare-event searches of particle physics. With these detectors it is possible to achieve low intrinsic backgrounds and sub-keV energy thresholds — all within a monolithic kilogram scale package. These advantages make germanium detectors an attractive vehicle for precision studies of the weak interaction.

The particular type of HPGe used in this work is a point-contact germanium detector. In the late 1980s scientists at Lawrence Berkeley National Lab (LBNL) were searching for ways to lower the energy thresholds of germanium detectors to detect feeble nuclear recoils from the “missing” dark matter of the universe [53]. The goal was simple: use small point contact electrodes to reduce electronic noise obscuring low energy signals. Decades later in the mid 2000s another scientific collaboration, COGENT, would develop an improved version of this technology (taking advantage of developments in low noise electronics) to look for nuclear recoils from both neutrinos and dark matter [54]. Finally, in 2023 the COHERENT collaboration succeeded in using a matured version of this same simple idea to observe neutrino-induced nuclear recoils on germanium.

This chapter summarizes the detectors used in this system, the design of the supporting apparatus, and its commissioning as an experiment.

### 6.1 Inverse Coaxial Point Contact Detector

The Ge-Mini detector system employs eight Inverse Coaxial Point-Contact (ICPC) detectors. These are p-type point-contact detectors with a particular geometry that allows for a kilogram scale crystals while still retaining the low-noise characteristics of the first point-contact



**Figure 6.1:** Schematic of a HPGe WELL<sup>TM</sup> detector made by Mirion Technologies. The ICPC detector is a type of WELL detector. The detector is housed in a low-background copper cryostat. Figure reproduced from [56].

detectors. By using a cylindrical bore-hole through the center face of the top of the crystal, it is possible to achieve strong electric fields throughout the geometry and successfully deplete the bulk of the crystal at operating voltages of a few kilovolts. A schematic of one of these crystals is shown in Figure 6.1. The ICPC is the same technology used in the Large Enriched Germanium Experiment for Neutrinoless  $\beta\beta$  Decay (LEGEND). Key features of the technology and its characteristics are summarized in [55].

### 6.1.1 Operating Principle

HPGe detectors are semiconductors. Two electrodes define the semiconductor: a p+ type point contact electrode on the center face of the bottom surface and an n+ contact formed by a thin, diffused layer of lithium on the outer surface of the crystal. Operated at cryogenic temperatures the crystal is an effective insulator. Under the application of significant voltage, typically two to four thousand volts for large detectors, free charge carriers are depleted. In the p-type detectors used in this experiment, electrons are guided along electric field lines to the n+ contact and the positively charged ions, or holes, are guided towards the p+ contact.

The inverse coaxial geometry, along with the large crystal size, creates long drift paths

to the point-contact. It can take up to several microseconds for charge carriers to traverse these long drift paths.

In the pursuit of reconstructing small signals from nuclear recoils, what must be detected are the quanta produced by energy deposits. For scintillating materials these quanta are the photons produced in scintillation, whereas for semiconductors the electron-hole pairs produced are the quanta. An advantage of semiconductor technology is that the ionization energy is small and so the number of quanta produced is large. Where a scintillating detector may count only a few photo-electrons (charge avalanches kicked off by a single photon) for a 1 keV nuclear recoil [57], a HPGe detector will drift around 50 electron-hole pairs [58]. What's more, the ionization yield is highly linear and as a result there is a reliable calibration from charge measured to energy deposited in the crystal.

However, the technical challenge of measuring individual charge carriers creates a need for charge-amplifying circuits, which are commonly used in the read-out of energy signals from HPGe detectors. In general, these charge-amplifying circuits do not permit exact measurements of charge yield. At the lowest energies, the resolution of a detector is dominated by electronic noise from the charge readout circuit.

### 6.1.2 Signal Readout

Reading out the drifted electric charge without incurring significant electronic-noise is a key feature of low threshold germanium detectors. The Ge-Mini detectors use charge sensitive pre-amplifiers located close to the point contact to minimize stray capacitance. In this pre-amplifier, the point-contact is connected to a field-effect transistor (FET) whose drain is fed back to a small capacitor  $C_f$ . When ionization radiation deposits charge  $Q$  inside the crystal, this charge drifts to the point-contact. When it reaches near the point-contact it induces a current that accumulates on the capacitor, producing a voltage  $V$  proportional to the charge. This is a standard in charge sensitive pre-amplification that allows for the integration of charge to be measured as voltage. This type of charge integration is slow, though not as slow as the drift-times themselves within the crystal.

A feature specific to the types of pre-amplifiers used in the Ge-Mini detector system is the lack of continuous resistive feedback. Instead, the Ge-Mini detectors use transistor reset pre-amplifiers (TRP), which continuously integrate the capacitor  $C_f$  until periodically discharging it. This discharge incurs a small dead-time as the pre-amplifier resets, but, in exchange, the overall electronic noise of the circuit is reduced by the elimination of additional capacitance from a feedback circuit [59] [60].

As has been emphasized, electronic noise is of the utmost importance, and determines the ultimate energy threshold. The equivalent noise charge (ENC) of the signal readout can be decomposed into [60]:

- Series noise, which is proportional to the total input capacitance.
- Parallel noise, which can arise from leakage currents.
- 1/f noise, which contributes at low frequencies.

The revolutionary success of point-contact germanium detectors is that the small p+ contact introduces minimal series noise by keeping the capacitance down to the pF level. For these detectors, it is typical to achieve an ENC full-width half-maximum (FWHM) of 50 electrons at zero energy, enabling thresholds of 500 eV or better.

Finally, the signals from the preamplifier are digitized and recorded so that they can be processed in off-line analysis.

### 6.1.3 Ge-Mini Detectors

The Ge-Mini detector system uses eight ICPC detectors, each with a mass of approximately 2 kilograms. The detectors, manufactured by Mirion Technologies Inc. [61], were supplied with specifications for mass, expected operating voltage, and electronic noise. All detectors were delivered with an expected energy resolution at zero energy of 150 eV FWHM or better. The details are summarized for each detector in Table 6.1.

Variations in specification are all typical of the specialized manufacturing process. The term active mass is an important parameter to consider. Each detector has a small so-called

**Table 6.1:** Manufacturer specifications for the mass of the Ge-Mini ICPC detectors. Manufacturer values for resolution changed with operating conditions and required in-situ characterization.

Detector	Total Mass (kg)	Pulser Resolution (eV)
Ge21	2.215	99
Ge22	2.191	96
Ge23	2.215	76
Ge24	2.222	95
Ge25	2.215	99
Ge26	2.209	140
Ge27	2.214	124
Ge28	2.220	84

dead layer on the lithium-diffused n+ contact where charge collection is incomplete. Inside the remaining active mass, charge from ionization is fully collected at the point-contact. In a CEvNS experiment this active mass defines the total number of nuclear targets than can be considered in the calculation of scattering rates.


## 6.2 Detector System Design

An engineering drawing of the detector apparatus is shown in Figure 6.2. The detector system is composed of an inner core that houses the eight ICPC detectors and the cryogenics necessary for operation. Surrounding this inner core are layers of passive and active shielding designed to mitigate backgrounds. These various components are briefly summarized.

### 6.2.1 Cryogenic Operation

The ICPC detectors are contained in individual low-background cryostats. The crystals are electrically isolated but thermally connected via a copper cold-finger to a common 180 liter liquid nitrogen Dewar that is topped-off daily during operation.

For the sensitive, low-noise detectors the liquid nitrogen (LN2) fills are turbulent periods marked by sharp increases in noise. In the decomposition of noise components, this is




**Figure 6.2:** A schematic of the Ge-Mini detector system (right) with key components labeled. The ICPC detectors are cryogenically cooled in a shared liquid nitrogen Dewar. Lead, plastic, and copper as used as passive shielding of environmental backgrounds. An active muon veto surrounds the system. An inset of the detector cryostats is also shown (left), highlighting the inverse coaxial geometry and small point contact readout.

contribution, often called microphonic noise, is parallel. During operations, the level of liquid nitrogen in the Dewar is recorded and used to identify daily refills during analyses.

Safely operating a cryogenic experiment at the SNS requires careful monitoring, and oxygen sensors surround the experiment.

### 6.2.2 Shielding

There are several layers of shielding that surround the detector cryostats, which serve to reduce the impact of environmental backgrounds. Working from the inside out, nearest the detector cryostats is an inner layer of high density poly ethylene (HPDE), which serves as neutron shielding. Surrounding this inner HPDE is a two-inch high-purity copper box. Copper is an effective shield of gamma rays while maintaining radio-purity to reduce backgrounds inside the shielding package. Surrounding the copper box is another four inch



**Figure 6.3:** Photograph of the Ge-Mini detector system during assembly. The ICPC are housed in copper cryostats, which are surrounded by layers of HPDE, copper, and lead. The visible side panels of the muon-veto surround the shielding package.

layer of HPDE shielding. The large amount of plastic inside the shield is crucial to mitigate neutron backgrounds which mimic the desired neutrino-induced recoils. Finally, surrounding this layer of neutron shielding is a lead castle that is four inches thick. The lead is an effective shield for environmental gammas inside the accelerator facility. An image of the detector cryostats with an incomplete shielding assembly is shown in Figure 6.3.

### 6.2.3 Muon Veto

Surrounding the passive shielding package of the Ge-Mini detector system is a final layer of active shielding in the form of two-inch plastic scintillator that serve as a veto of cosmic ray induced backgrounds. The plastic scintillator is segmented into large side panels which are instrumented with two photomultiplier tubes (PMTs), and one large top panel instrumented with six PMTs.

Active muon vetoing is a crucial part of achieving the necessary levels of backgrounds. The large amounts of lead and copper near the detectors are a source of muon-induced

neutron production. Some fraction of the neutrons produced inside the shielding make their way into the germanium detectors. While the neutron shielding inside the shielding package plays an important role in mitigating neutrons, identifying events in the germanium that are in coincidence with signals in the plastic scintillator significantly reduces backgrounds overall in the experiment, while incurring minimal dead-time in operation.

## **6.3 Commissioning and Campaign 1**

Commissioning of the Ge-Mini detector system began in the spring of 2022. The author, along with fellow graduate students James Browning (North Carolina State University (NCSU)) and Emma van Nieuwenhuizen (Duke), assembled the shielding package, developed the cryogenic operations, installed the detectors, and began the first cool down by the winter of the same year 6.4 .

The first data taking campaign, dubbed Campaign 1, did not produce data suitable for use in a CEvNS search. Instead this data-taking period, which ran for a few months in the early parts of 2023, served to develop the data acquisition, storage, and processing pipeline.

### **6.3.1 Slow Controls**

The experimental operation required several components of monitoring and apparatus control. The HPGe detectors were positively biased using MPOD [62] high voltage supply. The PMTs were operated using similar, negative bias high voltage supply. Low noise high voltage supplies are important for the performance of the experiment.

The liquid nitrogen level of the common detector Dewar was monitored and recorded, along with a redundancy of four thermocouples placed at fixed heights inside of the Dewar.

### **6.3.2 Data Acquisition**

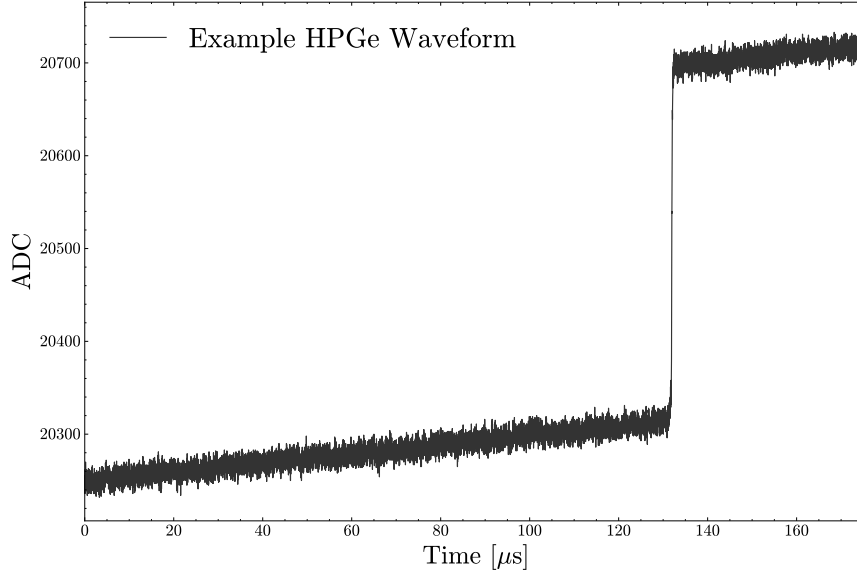
The data-collection scheme for the relevant of energy and timing signals was established during the Campaign 1 data-taking period. This served as the basis of the two subsequent data-taking periods described in this work.



**Figure 6.4:** A photograph of the Ge-Mini detector system, sitting inside of Neutrino Alley, on the day of the first cool down, on November 28, 2022. The system was in a state of partial assembly. The copper box, which surrounds the detector cryostats, is fully assembled. Visible inside the partially assembled shielding are the four-inch layers of plastic and lead shielding, as well as the active muon veto system.

The goal of the data acquisition is to reconstruct the energy in the HPGe detectors in coincidence with protons hitting the mercury target at the SNS. While other rare-event searches in neutrino physics or dark matter must continuously record signals from their detectors in hopes of observing a signal, CEvNS searches at the SNS benefit from only looking for neutrino signals in narrow windows around the regular proton collisions.

The strategy used to achieve this goal is digitization of the raw, continuous voltage from the HPGe detectors in coincidence with the beam. The energy signals from the HPGe detectors are digitized in a SIS3316 digitizer which are manufactured by Struck Innovative Systeme [63]. These digitizers use 16-bit analog to digital converter (ADC) to sample signals



**Figure 6.5:** Digitized waveform for a HPGe detector. Due to the lack of resistive feedback in the preamplifier, the waveform has a characteristic slope proportional to the leakage current of the detector. The step-like component of the waveform has amplitude proportional to the total charge  $Q$  collected from a radiation event. Electronic noise is present throughout the waveform and introduces finite resolution.

at 125 MHz. This produces waveforms that record a history of the voltage at a given time. An example of a digitized waveform is shown in Figure 6.5.

The SNS accelerator operators use a variety of electrical signals to monitor the performance of the beam used. Two of these signals are broadcasted to Neutrino Alley and are crucial to the operation of the experiments as they are used to synchronize data-taking to the presence of neutrinos. One signal is called event-39 (evt-39). This signal is broadcasted at 60 Hz from the accelerator facility and is tied to a magnetic kick that sends proton bunches to the target. The other signal is called event-61 (evt-61). This signal is a binary logic signal that indicates the presence of protons for a given magnetic kick. Combined, these signals indicate exactly when to look for neutrinos and if one can safely expect them to be there.

Operationally, the evt-39 signal is used to trigger the recording of HPGe waveforms externally, while the evt-61 signal is digitized separately internally. Here, external triggering refers to recording a waveform of what was happening in an energy signal based on an external logic signal, in this case the timing signal from the SNS, regardless of the energy of

the signal. It is contrasted to internal triggering, which is the recording of energy signals based on an internal logic, typically based on an energy threshold that is met. It is an incredible advantage to be able to trigger a low-energy neutrino experiment externally on a signal that indicates the potential presence of neutrinos.

The arrival time of neutrinos is well measured relative to the timing structure of these SNS signal cables.

There are several supporting signals that are also recorded. Each detector has an associated logic signal which indicates when pre-amplifier resets occur. These signals are recorded through an internal trigger using the same SIS3316 digitizer. The ability to associate in time, and ultimately reject, events in the HPGe detector with this logic signal is an important step in the preparation of relevant data. Energy signals from the muon-veto system must also be recorded through internal triggers, to perform a similar association and rejection.

In summary, in the Ge-Mini detector system there are up to eight total HPGe detector energy signals. Each energy signal has an associated TRP-reset logic signal. There are fourteen supporting muon-veto channels, which carry the photomultiplier tube (PMT) energy signals. These are recorded across two Struck SIS3316 [64] digitizers with a synchronized clock.

The digitizers are operated using the computer software Rdigdaq, developed by Jason Newby and James Matta supported by the radiation imaging group at ORNL. This software controls the setting of relevant registers in the operation of the digitizers as well as managing the file-writing system that must keep up with as much as 50 MB/s of raw data during full-operation.

### **6.3.3 Data Processing**

Data collected during regular operations was stored on the ORNL physics division computing cluster. The author and fellow graduate student James Browning developed an data processing software, JPROC, to perform the analyses of this experiment. JPROC is based

on the Julia computer programming language and uses common algorithms in radiation detection digital signal processing to reconstruct key features of HPGe events in the waveforms collected.

# Chapter 7

## Campaign 2: First Light<sup>1</sup>

From June 18, 2023 to August 15, 2023 the Ge-Mini detector system collected its first physics data. A total neutrino fluence of  $1.22 \times 10^{14} \nu/\text{cm}^2$  was recorded. On November 16 of that same year, the COHERENT collaboration unblinded the results of the first measurement of CEvNS on germanium at the SNS. In total,  $20.6_{-6.3}^{+7.1}$  neutrinos were detected. Due to the low backgrounds present in the search, a  $3.9\sigma$  significant rejection of the no-CEvNS hypothesis was achieved.

As a historical aside, there have been several notable examples of discoveries in neutrino physics made with countably few neutrinos. The discovery of the tau neutrino was possible with exactly four events [65]. The Homestake experiment which drove a central question in neutrino physics, did so reporting around only a dozen events per year [66]. In this light, the first observation of CEvNS on germanium, possible with only 20 events, is in respectable company.

The goal of this work is not just observation, but ultimately precision — the work established in this data-taking campaign laid the groundwork for subsequent measurements. This chapter summarizes the key features of the analysis that made this observation possible. It covers the treatment of HPGe waveforms, features of the data analysis, and the statistical result.

### 7.1 Detector Characterization

The eight COHERENT HPGe detectors are referred to throughout the text by the abbreviation Ge- followed by the last two digits of their manufacturer serial number. The eight detectors range from Ge-21 to Ge-28 sequentially. During this data-taking period five of the eight total detectors were used. An experimental challenge that has dogged this effort is

---

<sup>1</sup> *Published in Phys. Rev. Lett.* **134**, 231801 (2025)

excessive leakage currents due to vacuum failures of the individual detector cryostats. The three detectors not used in this analysis were inoperable during this time.

The five detectors used were Ge-21, Ge-23, Ge-25, Ge-26, and Ge-28. They were characterized in-situ leading up to the measurement. The evaluations consisted of measures of electronic noise, energy dependent energy resolution and drift-time characterization.

**Table 7.1:** Summary of HPGe detectors used during Campaign 2. Active mass is estimated from manufacturer specifications. Electronic noise is given as the full-width half-maximum of a gaussian noise peak at the start of the campaign.

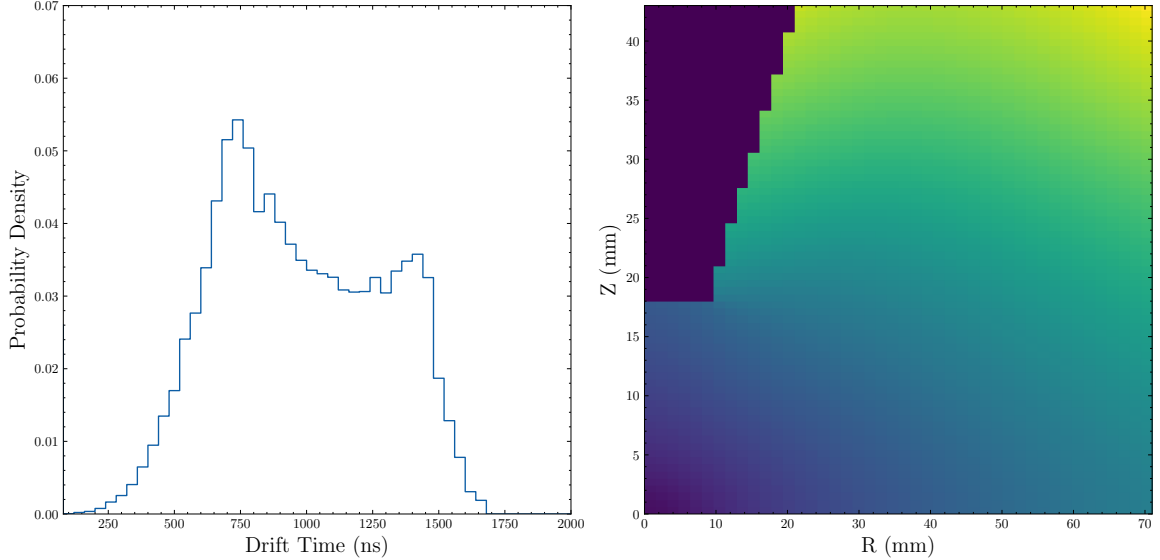
Detector	Active Mass (kg)	Pulser FWHM (eV)	Fano Factor
Ge-21	$2.13 \pm 0.04$	116(1)	6.8(3)
Ge-23	$2.13 \pm 0.04$	114(1)	7.1(4)
Ge-25	$2.13 \pm 0.04$	145(1)	6.4(3)
Ge-26	$2.13 \pm 0.04$	152(1)	6.8(3)
Ge-28	$2.14 \pm 0.04$	135(1)	6.4(3)

As the leakage currents of the detectors evolved over the course of the data-taking campaign the electronic noise degraded slightly. This effect was seen to have a negligible impact on the thresholds of the detectors, which was far above the noise.

Energy dependent energy resolution of HPGe detectors is often modeled, at sufficiently low energies, by the equation

$$\sigma(E) = \sqrt{\sigma_{\text{noise}}^2 + 2.96FE}, \quad (7.1)$$

where 2.96 eV is the band-gap of high purity germanium, and  $F$ , the Fano factor, is a measure of statistical dispersion the production of charge carriers. From Equation 7.1 it is possible to see that for small energy  $E$ , the resolution is dominated by the electronic noise resolution  $\sigma_{\text{noise}}$ . The energy resolution was characterized using calibration sources prior to the start of the data-taking campaign. The results of the characterization are summarized in Table 7.1.



**Figure 7.1:** (Right) calculations of drift-times of a COHERENT ICPC HPGe, Ge-28. (Left) The resultant calculated distributions of drift-times from a uniform sampling of the crystal. The particular shape of the drift-time distributions depends on the particular geometry of the crystal and are on the order microseconds.

Detector drift-times were evaluated using a numerical drift-time calculation software called SigGen [67], which takes as input the geometry and material composition of HPGe detectors. The calculations are matched to in-situ characterization measurements to account for practical variations in electric field strengths. As a consistency check, two different characterization measurements were used. One was a dedicated drift-time measurement using coincident gammas in the thorium decay chain, while the other was an in-situ measurement using coincidences in the germanium with tagged cosmic rays in the active muon-veto. The results were found to be consistent. Figure 7.1 shows an example of a drift-time SigGen calculation, and a representative drift-time distribution after the data-driven correction.


## 7.2 Data Collection

The Campaign 2 dataset was collected between June 18 and August 15 of 2023. During regular operations, the SNS delivers a pulsed 60 Hz proton beam, which creates a fixed number of neutrinos per Protons On Target (POT). For the Campaign 2 dataset an exposure

of  $1.96 \times 10^{22}$  POT were recorded. For the beam conditions during this data collection period, an expected total of  $0.288 \pm 0.029$  decay-at-rest neutrinos are produced per proton.

### 7.2.1 Triggering Scheme

While the neutrino production at the SNS follows a 60 Hz timing structure, the Ge-Mini data acquisition system was configured to trigger the recording of HPGe detector waveforms externally at 120 Hz. This produced, for each detector, one waveform coincident with every beam pulse, and a second waveform out of phase with the beam, 3 ms later. These sets of measurements, referred to as on-beam and off-beam throughout the text, enable a continuous measurement of the steady-state background. In the simplest counting experiment, the off-beam triggers can be subtracted from the on-beam triggers and the remaining excess can be attributed to beam-correlated physics, and ultimately neutrinos. A schematic illustrating the triggering scheme is shown in Figure 7.2




**Figure 7.2:** Timing diagram of the 120 Hz triggering scheme. SNS beam pulses arrive at 60 Hz (orange). For each pulse an on-beam trigger (green) is recorded in coincidence, and an off-beam trigger (blue) is recorded 3 ms later to sample steady-state backgrounds. 3 ms is much larger than the time scale of any potential beam-correlated backgrounds.

### 7.2.2 Exposure Summary

There were several technical challenges that caused interruptions in data collection. Humidity issues, related to the environmental constraints of operating in Neutrino Alley, led to a

mid-run pause to rehabilitate detector electronics. Individual live-times were variable. The total exposure calculated was dependent on the cumulative integration of each detector combined.

The resulting cumulative exposure is  $2.09 \times 10^{23}$  POT · kg. The calculation of exposure depends on the average beam-power as a function of time, reported by the SNS, and the fractional live-time after all analysis cuts (described further in Section 7.4). Figure 7.3 shows the cumulative exposure as a function of time over the data-taking period for each detector.



**Figure 7.3:** Cumulative beam exposure recorded over the duration of Campaign 2 as a function of date. Different detectors are individually shown as live-time for each detector was different.

### 7.2.3 Calibration Runs

Regular energy calibrations were performed to track gain stability and measure backgrounds with higher statistics. It is important to emphasize that while the external triggering scheme has two major advantages – it is a zero threshold trigger and comes with a three order of magnitude background reduction – it has the disadvantage of providing a lower statistical

measure of the steady-state backgrounds. Because of this, the calibration runs are useful complementary datasets.

A summary of calibration runs is provided in Table 7.2.

**Table 7.2:** Summary of weekly calibration runs during Campaign-2 data-taking.

Date	Duration (h)	Detectors List
June 20	7.7	Ge-21, Ge-23, Ge-25
June 27	23.1	Ge-21, Ge-23, Ge-25, Ge-28
July 03	32.4	Ge-21, Ge-23, Ge-25, Ge-28
July 11	16.1	Ge-23, Ge-25, Ge-28
July 25	14.5	Ge-21, Ge-23, Ge-25, Ge-26, Ge-28


#### 7.2.4 Data Selection

Independent of the activity inside of the HPGe detectors, a data collected through the external triggering scheme, referenced in Figure 7.2, was only considered during periods of stable beam operation. By recording the SNS reported beam power and energy, it was possible to identify stable periods of operation.

Beam energy was found to be stable at 1.05 GeV, however beam stability cuts resulted in approximately 5% loss in exposure. Figure 7.4 shows the beam power of the course of the data-taking period and the 0.5 MW cut value used to identify stable operations.

#### 7.2.5 Blinded Analysis

To avoid potential biases in the development of the waveform analysis and subsequent cuts, a blinding procedure was developed to separate on-beam and off-beam triggers. Throughout the development of the analysis, the on-beam triggers were temporarily removed from the data processing pipeline. Exclusively the off-beam triggers were considered throughout the development of the waveform analysis routines. Only after collaboration review were the on-beam triggers unblinded.




**Figure 7.4:** Average measured beam power versus time during Campaign 2, illustrating the time structure and stability of the beam delivery.

### 7.3 HPGe Waveform Analysis

When the data acquisition was triggered, whether externally through SNS timing signals or internally during calibration runs, waveforms were produced through digitizer sampling of preamplifier output at 125 MS/s. When triggered, the SIS3316 digitizers are configurable to report a buffered pre-trigger history, as well as the post-trigger record. For this measurement the germanium detector channels were configured to produce total record lengths of 22,000 samples, corresponding to 176  $\mu\text{s}$ . Of the 22,000 samples, 15,000 were pre-trigger history. Reconstructing the pre-amplifier output with near optimal performance depends on accurate measures of baseline charge.

A representative waveform from detector Ge-21 is shown in Fig. 7.5, highlighting the pre-trigger region and the signal arrival region. In particular, the data acquisition configuration ensures that when the SNS beam is on target, the neutrino signal events occur in the signal arrival region highlighted.



**Figure 7.5:** Example waveform showing the pre-trigger baseline and the signal-arrival region. The vast majority of the waveforms recorded, like this one, contain no detectable charge from ionization. For on-beam triggers, the expected neutrino arrival signal is highlighted. Charge arrival in the waveform is dependent on the drift-times of each detector.

Considering the externally triggered data used in the CEvNS search, charge collection caused by steady-state background radiation events can and do occur at any point throughout the recorded waveforms. In this analysis, events are only considered if they do not contain spurious events in the pre-trigger region. Due to the low-background nature of this experiment, this incurs negligible dead-time.

To perform the CEvNS search, it is necessary to reconstruct the key features of the several hundred million recorded waveforms. Of particular interest, for each triggered event, is the question: was a recorded charge collection above a particular energy threshold? And if so, when did it occur, relative to the protons striking the mercury target? For the overwhelming majority of recorded events there will be no charge collection of any interest – this is the power of only looking for events in coincidence with the beam – but to demonstrate this without bias it is necessary to algorithmically handle every waveform nonetheless.

The algorithmic treatment of waveforms, often referred to as digital signal processing (DSP), is well-studied in the case of low noise germanium detectors [60]. In particular, the problem of minimizing the error of the measurement of a charge  $Q$  in the presence of noise has well-defined optimal solutions in the literature. This work follows these standard

solutions closely. This section outlines the treatment of waveforms focusing on the energy estimation of events, and the pulse-onset estimation of those same events.

### 7.3.1 Pre-processing

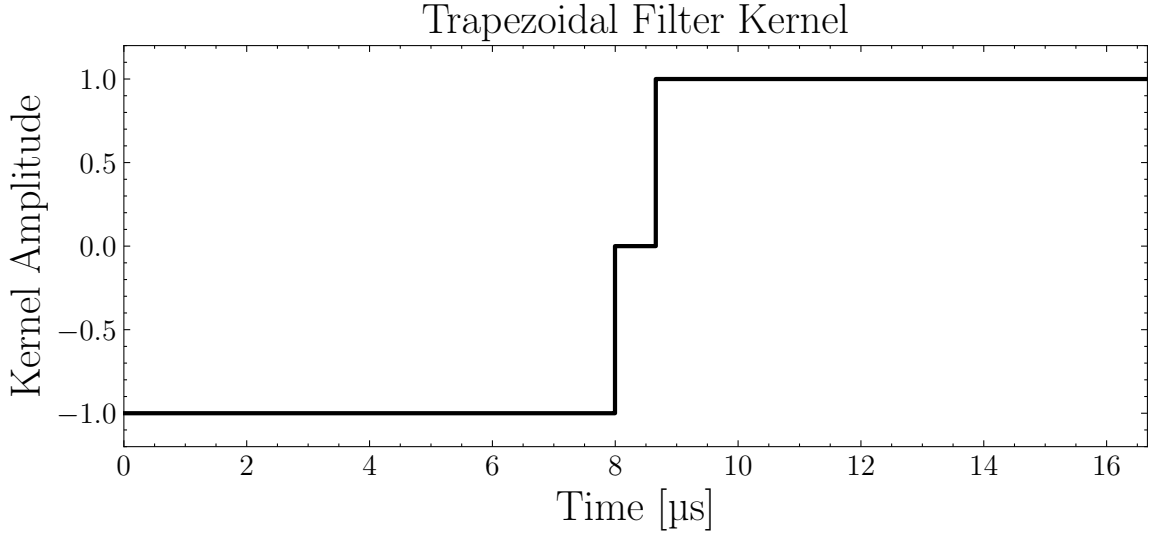
The pre-amplifiers used in this experiment do not have resistive feedback, and so, signal pulses from charge collection sit on a linearly increasing rise from the leakage currents of the detector. In this analysis, waveform baselines are subtracted by performing linear regressions to the first 13,000 samples, and subtracting this line from the entire waveform. Long baseline estimation regions are particularly important in the presence of low-frequency noise that can degrade the energy estimation.

### 7.3.2 Energy Estimation

The goal of the energy estimation is to reconstruct the total charge  $Q$ , which is proportional to the energy released in the form of ionization within the diode. In the presence of both signal and pre-amplifier noise, there is a well-established theory for the least-squares (optimal) measurement of  $Q$ . Under white-spectral noise and finite baseline estimation time, the optimal linear filter is a finite cusp-shaped weighting function [59]. For a diode with a known calibration to electron-recoils this is reported in units of energy as the ENC.

In practice, the digital implementation of these ideas takes the form of a linear time-invariant (LTI) filter applied to the sampled preamplifier waveform. The trapezoidal filter is a convenient approximation to the optimal cusp: it is defined by the discrete convolution of the waveform with a finite-length kernel such that rises linearly, remains flat, and falls linearly back to zero. This filter is nearly optimal in terms of noise performance, yet straightforward to implement and commonly used in radiation detection. The maximum of the output of the convolution is the estimated energy.

The configurable parameters of the trapezoidal filters are the rise time and flat time of the kernel. In particular the rise time, which is an effective integration window, has a large impact on the performance of the ENC resolution. An example of the convolution kernel



**Figure 7.6:** An example kernel of the energy-estimating trapezoidal filter shown in the time domain. Rise times, visible as the periods where the kernel is equal to  $\pm 1$ , are optimized for each detector.

and the resultant output is shown in Figure 7.6.

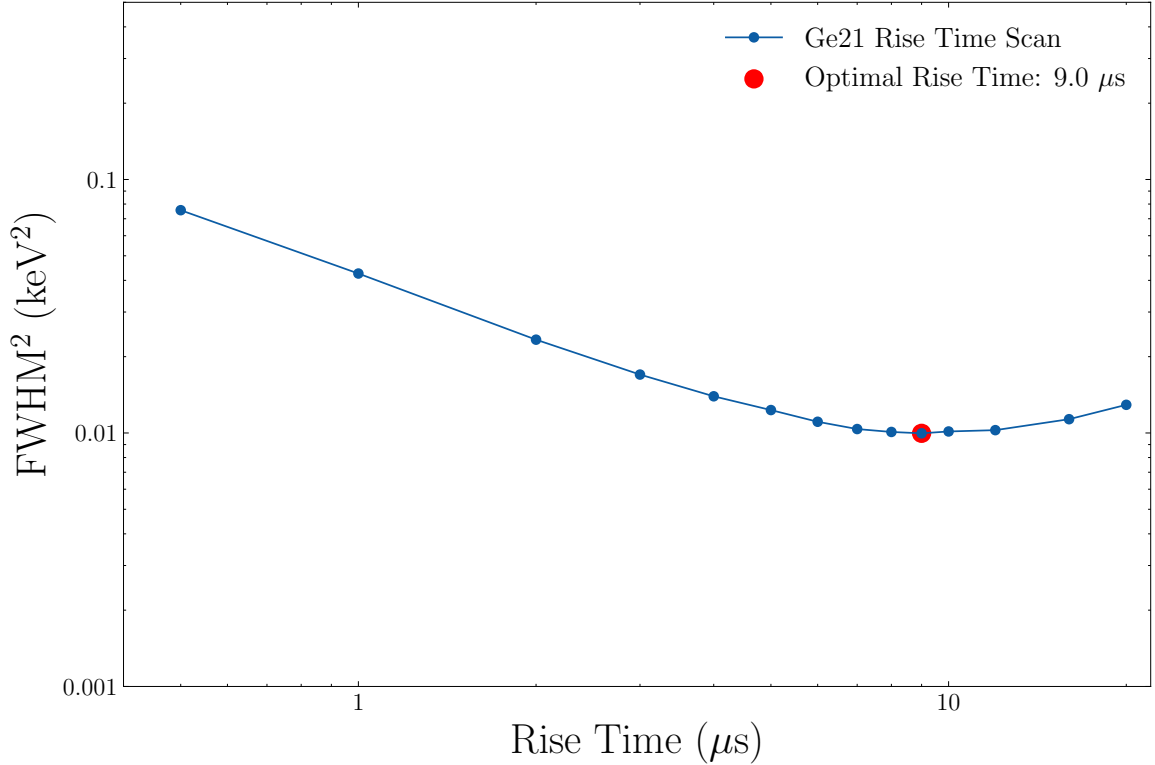
Since the performance of the trapezoidal filter depends on the choice of filter parameters and the particular noise conditions of each detector, a data-driven artificial waveform simulation is used to find optimal parameters. In particular the following procedure is used to find detector dependent parameters:

1. Analytic waveform templates are generated based on a sigmoid function that approximates the expected charge collection profile of a typical signal event.

$$S(x) = \frac{1}{1 + e^{-ax}},$$

where  $a$  parametrizes the rise-time.

2. Measured noise-only waveforms are harvested from the in-situ externally triggered data, and are added to the analytic waveform templates, producing realistic simulated signals of known amplitudes.
3. A range of trapezoidal filter rise times are applied to the sample of simulated waveforms.



**Figure 7.7:** An example of a rise time optimization performed on Ge-21 after the data-taking period.

For a given rise time, the application of the trapezoidal filter produces a gaussian distribution around the expected amplitude. The rise time producing the lowest FWHM for the resulting distribution is the optimal rise time.

An example of energy resolution optimization for a detector is shown in Figure 7.7.

As the leakage currents evolved throughout the dataset, filter parameters were re-evaluated daily. There was no impact of any minor variation on the analysis. Table 7.3 summarizes the filter parameters used, and the ENC resolution expressed as the FWHM in keV of each detector on the last day of data-taking.

### 7.3.3 Timing Extraction

A feature that is specific to CEvNS searches using neutrinos produced in an accelerator facility is the need to determine not just the energy of an event, but the time at which it

**Table 7.3:** Optimal trapezoidal filter parameters and corresponding electronic noise equivalent (ENC) resolution on the final day of data-taking.

Detector	Rise Time ( $\mu\text{s}$ )	Flat Time ( $\mu\text{s}$ )
Ge-21	8.5	0.67
Ge-23	4.0	0.67
Ge-25	4.5	0.67
Ge-26	6.0	0.67
Ge-28	8.0	0.67

occurred. The powerful background rejection possible at the SNS depends on the ability to correlate events precisely to the timing of protons hitting the mercury.

In HPGe detectors the time between a recoil that causes ionizing radiation and the charge collection near the point-contact depends on the drift-times in the detector, on the time-scale of microseconds.

The characteristic rise commonly observed in waveforms, like that shown in Fig 7.8, is related to the arrival of charge near the point-contact (and the resultant accumulation of charge on the preamplification capacitor). The rising edge of the waveform signal does not indicate the *creation* of the ionized charge - and yet, it is the creation time that is specifically needed to be associated to neutrino production.

In fact, from the shape of low energy HPGe waveforms alone it is impossible to determine the exact time of that the ionizing radiation. Instead, the strategy used in this experiment is to identify the onset of the pulse as it is seen in the electronics readout. The specific feature in the waveform that is estimated, referred to as  $t_{50}$  in the analysis, is the time at which the recorded waveform has reached half of the estimated energy.

This  $t_{50}$ , which is not the creation time of the ionized charge, is the only thing that can be reconstructed and is the key timing feature of every waveform considered. In the search for neutrinos recoiling off of germanium nuclei, and producing ionization, the drift-time distributions, described in Section 7.1, must be used to predict when signal-related pulse onsets can occur.

The timing reconstruction in Campaign 2 followed a custom procedure adapted from

a related algorithm used in the literature [68]. The goal was to determine characteristic timing points ( $t_{10}$ ,  $t_{50}$ ,  $t_{90}$ ) for small-amplitude pulses in the presence of electronic noise.

An asymmetric trapezoidal filter was first applied to each waveform. This fast rising filter provided a balance of suppressing high-frequency noise and preserving the rising edge of the original waveform. Starting from the maximum of the filtered waveform, the algorithm moved backward in samples to locate the first sample falling below 50% of the reconstructed energy: ( $t_{50,1}$ ). Then the iteration continued backwards in samples to find the 10% crossing: ( $t_{10}$ ). From  $t_{10}$  it then moved forward in samples to find a second 50% crossing ( $t_{50,2}$ ). The final  $t_{50}$  was defined as the mean of these two estimates,

$$t_{50} = \frac{1}{2} (t_{50,1} + t_{50,2}). \quad (7.2)$$


In the case where the filtered waveform did not exceed 50% of the reconstructed energy,  $t_{50}$  was set to  $t_{50,1}$ .

While sufficiently effective for an observation of CEvNS, this algorithm required several ad-hoc checks and relied on the energy reconstruction to guide the timing search. Ultimately, the performance of this algorithm was the limiting factor in the determination of an energy threshold in the analysis. Subsequent analyses replaced it with an improved technique approach that offered better performance precision more statistically well-defined timing estimator. This is discussed in Chapter 8.

This implementation was not detector dependent. In particular every detector used the same kernel for the asymmetric filter

$$h[n] = \begin{cases} +1, & 0 \leq n < 5, \\ 0, & 5 \leq n < 18, \\ -1, & 18 \leq n < 268, \end{cases}$$

A figure highlighting the convolution outputs of respective trapezoidal filters used in



**Figure 7.8:** HPGe waveform in black with reconstructed energy of 1.5 keV. This is the lowest energy waveform considered in the analysis. Shown in orange is the output of the energy-optimized trapezoidal filter. The maximum of the filter is highlighted as the energy reconstructed for this event. In green is the asymmetric trapezoidal filter used for timing reconstruction. The identified  $t_{50}$  is highlighted. Figure is adapted from [46].

both the energy estimator and onset finding are shown in Figure 7.8 for a waveform at the energy threshold considered in this analysis: 1.5 keV.

### 7.3.4 Reconstruction Uncertainty and Thresholds

The accuracy of both the energy and timing estimators is limited by electronic noise. The data-driven waveform simulation described in Section 7.3.2 was used to evaluate the bias and resolution of the feature-reconstruction as a function of true signal energy.

For each detector, a large sample of artificial waveforms were generated ( $\sim 20,000$ ), each with a known energy and timing. The simulated waveforms were processed with the same optimized energy algorithm and custom timing algorithm used in the analysis. The reconstructed values were compared to the known inputs. This procedure yielded the mean distributions of residual reconstruction – effective measures of both bias and variance – as a

function of true energy.

The results confirmed an excellent energy reconstruction, comparable to the underlying noise limit present in the pre-amplifier electronics. Across all detectors the energy estimator remains essentially unbiased, with fixed performance in each detector, down to approximately 300 eV.

In contrast, the timing estimator showed a rapidly broadening distribution of reconstructed  $t_{50}$  below 1.5 keV. The distribution also showed significant biases below this energy threshold. This behavior is consistent across detectors and reflects the fundamental limitation of the asymmetric trapezoidal timing algorithm in the low-signal regime.

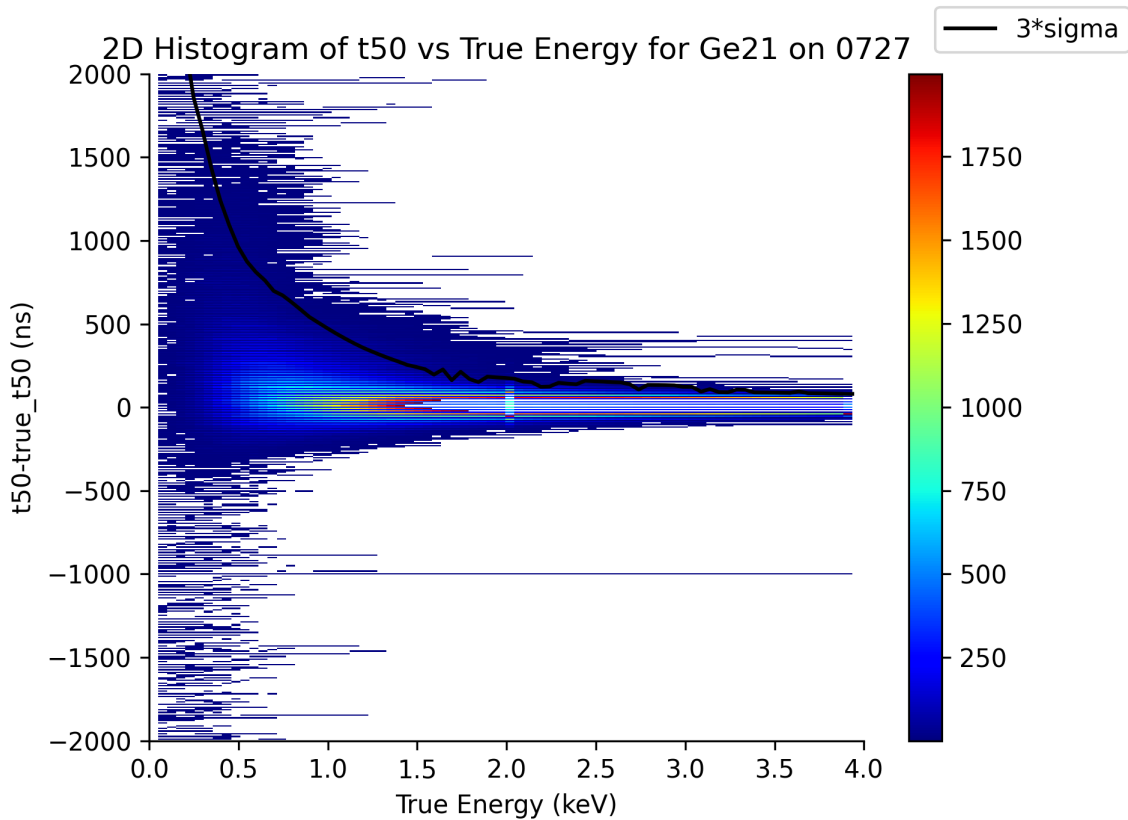
Representative results for time estimator reconstruction study are shown in Figures 7.9. This illustrates the degradation below 1.5 keV, and the narrow identification above pulse onsets above this threshold.

## 7.4 Data Quality and Cuts

Out of the hundreds of millions of waveforms recorded, the vast majority contain no measurable charge collection. To re-emphasize a key feature of measurements at the SNS: the remarkably low-background nature of this experiment is enabled by only looking for physics events in a narrow coincidence with the beam.

For the events that are reconstructed above the 1.5 keV threshold, the majority do not originate from a real ionization in the diode. Instead, many events are reconstructed in the energy region-of-interest by the indiscriminate trapezoidal filters used, despite not being, in a sense, real physics events. These may be vibrational noise induced by a turbulent LN2 fill, or a regular pre-amplifier reset event that tricks the energy estimation algorithm.

Luckily, the waveforms associated to these spurious events are topologically distinct from the so-called physics events that constitute the CEvNS search (be they signal or background events). Therefore, the analysis depends on a series of cuts to filter events down to this physics-only set. It is important that this procedure of pruning events originating from spurious noise does not inadvertently throw out any physics events. In the technical language,



**Figure 7.9:** Timing reconstruction residuals and resolution ( $t_{50}$ ) as a function of true energy for Ge-21, illustrating the rapid growth of timing uncertainty near the analysis threshold. Original figure courtesy of James Browning.

the data-cleaning cuts used in this analysis were ensured to be nearly 100% efficient in retaining signal. In general the cuts chosen were either based on uncorrelated auxiliary timing information, or very conservative and targeted towards obvious pathological features. Due to the low-background nature of this experiment it was possible to manually inspect every true radiation event in the entire dataset to ensure the 100% signal acceptance.

Finally, once trimmed to a physics-only dataset, the application of the active muon-veto further reduced the backgrounds to the final level used in the statistical analysis. This section outlines the simple and effective data cleaning cuts used, the dead-times they incurred, and a description of the active muon-veto.

### 7.4.1 Timing Based Cuts

Several cuts were based on removing pathological events using coincidences with timing information from axillary signals. Since these are independent of pulse-shape, and uncorrelated to the beam, they cannot have any impact on signal-acceptance.

#### Pre-amplifier Reset

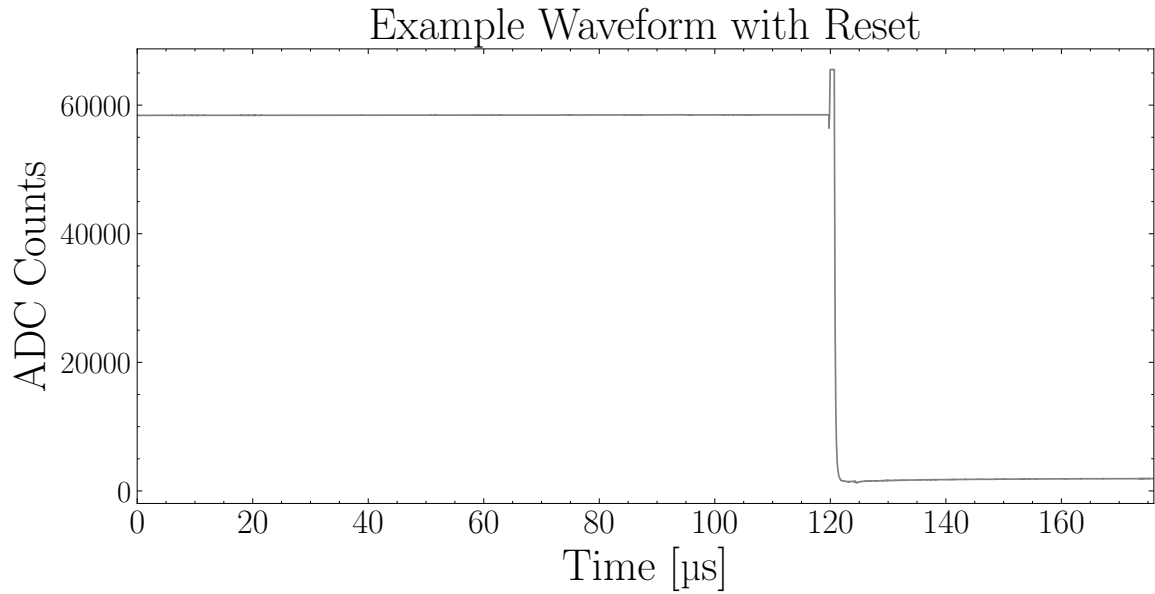
The pre-amplifier used in the front-end electronics of the HPGe detectors used do not contain a feedback resistive element. As a result the charge accumulated from either radiation events or from leakage currents does not dissipate continuously. Instead, when a specific charge (voltage) is reached, the a discharge (reset) is triggered. The dead-time associated with the restoration of the pre-amplifier was measured to be less than 200  $\mu$ s.

An example of a pre-amplifier reset and subsequent baseline restore is shown in Figure 7.10. The relative frequencies of these resets depends on the leakage current of each detector. The front-end electronics of the HPGe detectors used have a configurable transistor-transistor logic (TTL) signal associated with the pre-amplifier reset. These signals were digitized internally, independent of the SNS timing signals, to track and reject these pre-amplifier resets, as well as the 200  $\mu$ s window after each reset.


#### LN2 Fills

The Dewar that houses all HPGe detectors was refilled daily with LN2. These fills lasted approximately forty five minutes and during this time detectors were sensitive to vibrations from the boil-off and liquid flow. The number of non-physics events reconstructed in the energy region of interest during these fills was large. An example of such a waveform is shown in Figure 7.11.

As a result, all events occurring during a LN2 fill were thrown out. Depending on the detector sensitivity to the vibrations induced, additional time was added after fills to the rejected time periods. A table summarizing the time after each fill is shown in Table 7.4.



**Figure 7.10:** Example pre-amplifier reset event. The integrating capacitor is discharged (with an accompanying drop in baseline voltage of several volts). The pre-amplifier resets to normal operating conditions over  $200\mu\text{s}$



**Figure 7.11:** An example of an event recorded during a liquid nitrogen Dewar fill. Events recorded during these periods show large baseline fluctuations that bias energy reconstruction and are not correlated to physics events in the crystals.

**Table 7.4:** Amount of time the LN2 refill exclusion window is extended after fills have concluded for each detector.

Detector	Additional Time (m)
Ge-21	0
Ge-23	10
Ge-25	30
Ge-26	0
Ge-28	0

### Noisy Period Removal

Apart from the the expected LN2 fill periods, there were additional small periods of time of recorded microphonic noise with no known correlation. There were also isolated periods of abnormally high TRP reset rates recorded.

#### 7.4.2 Waveform Based Cuts

Another broad category of data quality cuts is related to the specific topology of waveforms. There were identifiable categories of pathological waveforms are not related to charge collection in the diode, but that are reconstructed by the energy-estimating algorithm in the energy region of interest. While they were degenerate with signal-like events in energy dimension, simple, additional parameters were found to lift this degeneracy and remove these events without affecting the signal-like events.

### Saturated Events

There were several types of events which either positively or negatively saturated the configured dynamic range of the analog to digital converter on the SIS3316 digitizer. Often these had their origins in pre-amplifier resets that did not have an accompanying TTL signal.

Any recorded waveform that saturated the dynamic range of the digitizer was rejected.

### Trapezoidal Filter Minimum (Cross-Talk)

One large type of pathological event was due to cross-talk inside of the digitizer between the reset TTL signals from one detector and the energy channel in another detector. During the analysis of this data-set, the fact that this was cross-talk was not known. The category of events was mistakenly assumed to be some type of transient noise unrelated to other detectors. Later analyses (discussed further in Chapter 8), improved the understanding of this behavior. Instead during this analysis a waveform based strategy was used to identify the key features of the inductive cross-talk, and reject waveforms in this manner.

An example of this cross-talk is shown in Figure 7.12. The strategy that was taken to reject these waveforms was to cut on the minimum of the energy-estimating convolution. This cut, referred to in the analysis as the trapezoidal filter minimum, ensured that after baseline subtraction no waveform had a negative-going pulse which was present in these cross-talk events.


The particular values chosen in this cut are summarized in Table 7.5, and were validated with manual inspection of every waveform in the data-set to ensure that no physics-like events were removed during the blinded analysis.

**Table 7.5:** Optimized cut values for the minimum of the trapezoidal filter convolution output to reject cross-talk events..

Detector	Trap Minimum
Ge-21	-6
Ge-23	-8
Ge-25	-10
Ge-26	-7
Ge-28	-6

### 7.4.3 Muon Veto

Finally, apart from the removal of pathological waveforms from the dataset, there are physics-like events that are recorded in coincidence with tagged events in the active muon-veto



**Figure 7.12:** An example of a detector undergoing a pre-amplifier reset (top) and the induced cross-talk (bottom). The width of the bi-polar cross-talk is related to the width of the TTL signal sent by the resetting detector.

around the HPGe detectors. These events have their origin in ionizing radiation from cosmic rays directly, from secondary particles from electromagnetic showers, or from muon-produced neutrons produced inside the shielding material itself.

In the case of muon-produced neutrons, the time between the muon passage through the material and the neutron recoil in the germanium can be on the order of microseconds due to neutron scattering within the shielding material. Dedicated simulations were performed using the computer software MCNP [69], to determine appropriate coincidence windows for the analysis veto applied later in software. It was determined that  $10 \mu\text{s}$  was sufficiently large to contain effectively all coincidences.

To simplify the application of the veto, and to provide an easily calculable dead-time, a

fixed veto window was used for every waveform in the analysis. That is, for every single HPGe waveform recorded, if there was any triggered event in the muon-veto system between waveform samples 11,200 and 20,200, it was rejected. This was the case regardless of the presence of any detectable charge collected. This window ( 73  $\mu$ s long) was significantly larger than the coincidence window determined by simulation. It incurred a slightly excessive dead-time penalty.

#### 7.4.4 Dead-Time

In summary, a series of cuts were applied to arrive at a low background of pure physics events, without introducing non-trivial efficiency for the signal-like population. All cuts used were either based on uncorrelated timing signals – such as the active muon-veto or pre-amplifier reset logic signals – or were based on disqualifying waveform topologies that were uncorrelated to the presence of ionizing radiation.

These cuts had a non-negligible impact on the recorded live-time, or exposure. The percent of events rejected was calculated directly and folded into all final exposure calculation. This is summarized in Table 7.6

**Table 7.6:** Percentage of triggers remaining after each successive quality cut. The final column shows the cumulative effect of all cuts.

Detector	Bad evt61 Cut	TRP Reset	Muon-Veto	Trap-Min + Beam Ops	All Cuts
Ge21	96.83	96.13	88.30	88.24	86.02
Ge23	97.24	94.35	87.04	86.98	79.40
Ge25	97.24	95.40	88.02	85.53	76.76
Ge26	92.02	89.74	81.97	81.11	79.61
Ge28	97.32	96.71	89.20	89.08	86.76
Average	96.59	94.96	87.44	86.70	81.76

## 7.5 Backgrounds

The statistical analysis is an estimate of the number of CEvNS events observed in coincidence with protons bombarding the target at the SNS. This signal arrives in the presence of background, and so these backgrounds must be simultaneously estimated. The two major backgrounds that must be considered in this analysis are steady-state backgrounds and beam-related backgrounds.

### 7.5.1 Steady-State Backgrounds

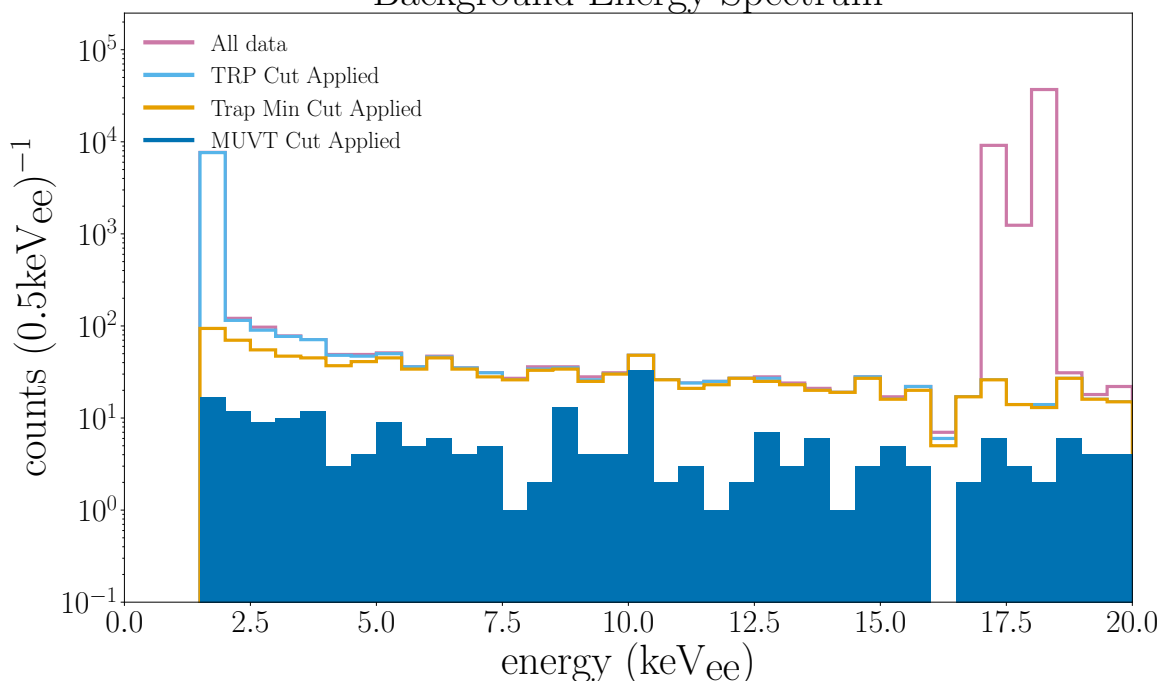
The steady-state backgrounds can be well measured in the absence of protons on target. In this analysis there are two measurements of the steady-state background. The first comes from the one-to-one, externally triggered dataset, that accompanies every on-beam trigger. These triggers, referred to as the off-beam dataset, are analyzed in the blinded analysis to develop all cuts described in Section 7.4. The externally triggered background dataset, after all cuts, is shown in Figure 7.13.

The other important measure of the steady-state backgrounds are the weekly calibration runs. These provide high-statistics measures of the shape of the background distribution. There is an important distinction between the two. The externally triggered dataset is collected with zero threshold in energy at the digitizer level. An energy threshold of 1.5 keV is applied in the analysis. The internally triggered calibration data is collected with energy thresholds at the digitizer level. This threshold introduces – for the calibration data only – an unknown trigger efficiency.

### 7.5.2 Beam-Related Backgrounds

Beam Related Neutrons (BRNs) originate from the SNS target and propagate through the facility to reach Neutrino Alley. Because they are produced in time coincidence with the beam pulse and can deposit nuclear-recoil energies within the CEvNS region of interest, they represent a potential irreducible background.

## Background Energy Spectrum



**Figure 7.13:** Campaign 2 externally triggered background spectra integrated over the full data-taking period. A series of histograms are shown after successive cut application. Data cleaning in the form of the pre-amplifier reset rejection (TRP Cut) and the cross-talk cut (Trap Min Cut) are shown. The final background dataset shown in blue is used in the statistical analysis as the sample of the steady-state background. Original figure courtesy of James Browning.

To quantify their potential impact, a dedicated MCNP simulation was performed using the Ge-Mini Campaign 2 shielding configuration. Two source hypotheses were considered: neutrons generated at the concrete wall adjacent to the detectors, and neutrons entering from the downstream hallway. The input neutron spectra were motivated by prior neutron flux measurements performed by the collaboration. In particular, using an Eljen [70] liquid scintillator cell in the Ge-Mini location.

The two source configurations predict similar low-energy recoil spectra. When scaled to the Campaign 2 exposure, the total expected Beam Related Neutron (BRN) contribution was found to be

$$N_{\text{BRN}} = 0.67 \pm 0.34.$$

The quoted uncertainty includes the  $\sim 20\%$  measured flux normalization error.

### 7.5.3 Background Model

The statistical analysis requires a the background probability density function (PDF). For this analysis it is derived entirely from the internally triggered calibration data collected. This choice provides a high-statistics sample that is representative of the steady-state distribution.

A potential concern is the reduction of trigger efficiency in self-triggering mode at low energy. Dedicated comparisons to the externally triggered (off-beam) data revealed that such efficiency losses affected two detectors above the 1.5 keV threshold: Ge-25 and Ge-26. Consequently, the final background PDF is constructed using only the calibration data from Ge-21, Ge-23, and Ge-28, which showed no measurable efficiency loss.


Figure 7.14 displays the agreement, within statistics, between the externally triggered data (black) and the rescaled internally triggered data (green). Also showed is the expected CEvNS signal (blue), discussed in further detail in Section 7.6.

The region around prominent peaks are excluded from the analysis. This is due to the apparent disagreement in the internally triggered data used to generate the PDF and the externally triggered off-beam spectrum. There is different radiogenic activity present in each detector. The use of only some of the detectors for the PDF, but all of the detectors in the statistical analysis led to a poor modeling of the peaks. This ultimately has negligible impact on the sensitivity of the measurement, evidenced by the small signal present in this excluded energy range.

## 7.6 Signal Prediction

### 7.6.1 Overview

The objective of this section is to describe the prediction of the observable CEvNS signal in reconstructed energy and time. This prediction provides the two-dimensional PDF used in the statistical analysis. It requires combining the CEvNS differential cross-section, the SNS neutrino flux, and detector effects.



**Figure 7.14:** A comparison of the internally triggered calibration data background is shown as data points. The externally triggered background dataset is shown as a yellow histogram. Agreement within statistic is observed. The CEvNS prediction over this exposure is overlaid. Figure is adapted from [46]

### 7.6.2 Flux Averaged Rates

In general, the two-dimensional differential rate is written as an integral over neutrino energies relevant in the SNS neutrino flux, described in Chapter 5, as the following:

$$\frac{d^2 R}{dT dt} = N_{\text{tgt}} \int_{E_{\nu}^{\text{min}}}^{E_{\nu}^{\text{max}}} \phi(E_{\nu}, t) \frac{d\sigma}{dT}(E_{\nu}, T) dE_{\nu}, \quad (7.3)$$

where  $N_{\text{tgt}}$  are the number of nuclear targets, in this case the number of germanium atoms in the active mass,  $\phi$  is the distribution function of neutrino energies propagated to the germanium targets with units  $\nu/\text{cm}^2$ , and  $d\sigma/dT$  is the differential rate described in Chapter 4.

As has been emphasized throughout this work, the differential rate is predicted with small theoretical uncertainty. The largest source of theoretical uncertainty lies in the

choice of nuclear structure form factor. The COHERENT collaboration uses as a base prediction the Klein-Nystrand model described in Section 4.1.2 with diffractive neutron radius  $R_n = 1.2 A^{1/3}$ , where  $A$  is the atomic mass of the target isotope.

Importantly, while the differential rates do not depend on neutrino flavor at tree-level, electromagnetic radiative corrections introduce flavor dependent differences in the form of neutrino charge radii. For this reason it is of physical interest to consider the differential rates for each flavor exactly. However, for the sake of simplicity the text neutrino flavor indices are suppressed in this text.

### 7.6.3 Quenching Factor

The observable energy in a germanium detector is the *electron-equivalent* energy  $T_{ee}$ , which corresponds to the amount of ionization charge produced by a nuclear recoil of true energy  $T$ ; sometimes in the literature this is also expressed as  $T_{nr}$  to denote the nuclear recoil energy. Because a fraction of the total recoil energy is lost to dissipation mechanisms that are not ionization,  $T_{ee}$  is smaller than  $T$ . The ratio

$$Q(T) \equiv \frac{T_{ee}}{T} \tag{7.4}$$


is often referred to in the literature as the quenching factor.

This ionization yield from nuclear recoils has been extensively studied [71]. There is a phenomenological model of the ionization yield that follows the Lindhard model described as:

$$Q_L(T, k) = \frac{k g(\epsilon)}{1 + k g(\epsilon)},$$

where  $k$  is a dimensionless parameter determined from experiment and  $g(\epsilon)$  is the Lindhard ionization function of the dimensionless energy  $\epsilon$ , which depends on recoil energy  $T$ . In this analysis, the model parameters taken from the CONUS collaboration measurement [58], so that  $k = 0.157 \pm 0.004$ .

The energy threshold in this analysis is a threshold in ionization energy measured, not the



**Figure 7.15:** Measured germanium quenching factor as a function of nuclear recoil energy. The curve shows the Lindhard model with best-fit parameter  $k$ . Original figure and data appears in [58]. The vertical red line marks 7.5 keV, the recoil energy threshold of this analysis. The modified figure is from Janina Hakenmuller.

nuclear recoil energy. Under this transformation, and with this model parameter, the nuclear recoil energy threshold is in fact 7.5 keV. Figure 7.15 shows a collection of measurements of the quenching factor in HPGe detectors with the equivalent energy threshold for this analysis highlighted. In this energy range the quenching factor data agrees very well with the empirical model used and the uncertainty it introduces on the CEvNS measurement is small.

In addition to the mean quenching factor, Lindhard’s model predicts an energy-dependent spread in the electron-equivalent yield for a fixed nuclear-recoil energy. This “quenching smearing factor” (QSF) represents the stochastic partition of recoil energy between ionization and other channels of energy dissipation. In this analysis, the magnitude of this broadening

is below the dominant detector resolution [72].

To obtain the observable spectrum in  $T_{ee}$ , the nuclear-recoil differential rate of Eq. (7.3) is transformed according to

$$\frac{d^2 R}{dT_{ee} dt} = \int_0^\infty dT \frac{d^2 R}{dT dt} \mathcal{R}_Q(T_{ee} | T), \quad (7.5)$$

where  $\mathcal{R}_Q(T_{ee} | T)$  is the quenching response function. For an ideal deterministic mapping – that is, no QSF – this reduces to

$$\frac{d^2 R}{dT_{ee} dt} = \frac{1}{Q(T)} \frac{d^2 R}{dT dt} \Big|_{T=T(T_{ee})}.$$

With this transformation it is possible to arrive at an energy range relevant for this analysis. While the recoil energy range considered was from 7.5 keV<sub>nr</sub> to the kinematic endpoint of the recoil spectrum, the new electron-equivalent energy range is [1.5, 20] keV<sub>ee</sub>, where the subscripts “nr” and “ee” refer to the nuclear recoil energy and electron equivalent energy respectively.

#### 7.6.4 Detector Smearing and Efficiencies

The idealized two-dimensional recoil rate of Eq. (7.5) must be further transformed with the finite detector resolution, timing uncertainties, and selection efficiencies.

##### Energy Resolution

The reconstructed electron-equivalent energy  $T_{ee}$  is subject to finite detector resolution, with energy-dependent energy resolution modeled by a gaussian distribution function, dependent on the specifics of each detector. In particular, the resolution of a function of energies is characterized using internally triggered calibration data, with width  $\sigma_E(T_{ee})$ .

## Onset ( $t_{50}$ ) Smearing

The reconstructed pulse-onset also carries a timing uncertainty from the  $t_{50}$  finding algorithm. For this analysis it was decided to approximate this as a fixed gaussian distribution with fixed width of  $\sigma_t \approx 100$  ns. While this is not correct, it simplified the numerical work of producing the PDFs and had a small effect in the shape of the final distributions, too small to be noticed in this first measurement. Later analyses have corrected this.

## Detector-Dependent Drift-Times

The detector dependent drift-times must also be folded into the signal prediction. That is, the charge carriers from neutrino-induced nuclear recoils must drift from the interaction site in the crystal to the readout point-contact. This introduces a probabilistic delay of true interaction times. The drift-time distributions, as described in Section 7.1, are drift-time calculations matched to characterization data. These empirical drift-time probability density functions  $P_{\text{drift}}(t)$  are convolved with the onset smearing to produce the final timing model.

## Efficiency Functions

Two simple, efficiency windows are applied in energy and time:

- **Energy efficiency:** a step function equal to 1 for  $1.5 < T_{ee} < 8.5$  keV and  $11 < T_{ee} < 20$  keV, and 0 elsewhere.
- **Timing efficiency:** a step function equal to 1 for  $-4 < t < 36$   $\mu\text{s}$ , where  $t$  is measured relative to the external proton-on-target timing signal.

As an aside, the fact that the signal acceptance is 100% in the energy and time region of interest is due to a few things. First, the external triggering scheme ensures that the data is recorded with no losses due to a trigger efficiency, as can be the case in some CEvNS analyses with a self-triggering scheme [41]. Second, the low-background nature of the experiment ensures that the vast majority of waveforms recorded have no events in them. As a result the likelihood of timing efficiencies due to pile-up, related to the Poisson statistics of the

steady-state background, are vanishingly small. Finally, all cuts applied in the analysis were ensured to have full signal acceptance.

### 7.6.5 Signal PDF

The signal PDF produced by the successive convolutions of the ideal recoil spectrum with the smearing functions described, along with the multiplication of the trivial efficiency windows, yields a final, two-dimensional PDF used in the statistical analysis.

In this analysis, the unity-normalized PDF has no free parameters that distort the shape. Future analyses may consider to parameterize the shape of the signal PDF through alternative physics models or to parametrize the nuclear structure form factor. Here, however, the only free parameter considered is the total number of counts observed.

## 7.7 Expected Counts

The total expected counts depends on the signal prediction that has been described here, including all of the smearing and thresholding effects, as well as the overall dimensionful normalization of the differential rate equations used. In particular it depends on the total number of neutrinos that pass through the total number of sensitive germanium nuclei.

This integrated live-time is 10.22 GWhkg, and the resultant Standard Model expectation is 35.1 in the signal region of interest.

A discussion of systematic uncertainty on the Standard Model prediction is necessary and summarized in Table 7.7. The dominant source of systematic uncertainty – that is, an uncertainty in the prediction – comes from the neutrino flux uncertainty of 10 percent.

## 7.8 Statistical Analysis

Following the Particle Data Group convention [73], the likelihood for  $N_{\text{on}}$  observed events, when the total number of events itself depends on the parameters of interest, is written as an extended likelihood:

**Table 7.7:** Summary of relative uncertainties on CEvNS signal rate. The uncertainties considered here are on the amplitude of the signal prediction and all shape distorting effects are considered sub-dominant and ignored in this analysis.

Parameter	Relative Uncertainty
Total neutrino flux (pion production)	10%
Distance to target ( $19.1 \pm 0.1$ ) m	0.5%
Quenching factor ( $k = 0.157 \pm 0.004$ )	$\ll 1\%$
Quenching smearing ( $0.024 \pm 0.002$ )	negligible
Form factor (Klein, Helm)	$\sim 1\%$
Active mass ( $96.3 \pm 2.0$ )%	2%
Energy resolution	negligible
Energy scale calibration	1–2%

$$\mathcal{L}(\nu_s, \nu_b) = \underbrace{\frac{(\nu_s + \nu_b)^{N_{\text{on}}} e^{-(\nu_s + \nu_b)}}{N_{\text{on}}!}}_{\text{Poisson probability for total counts}} \underbrace{\prod_{i=1}^{N_{\text{on}}} \frac{\nu_s P_s(E_i, T_i) + \nu_b P_b(E_i, T_i)}{\nu_s + \nu_b}}_{\text{contribution from individual events}}.$$

Similarly, the extended likelihood for observing  $N_{\text{off}}$  is described by:

$$\mathcal{L}(\nu_b) = \frac{(\nu_b)^{N_{\text{off}}} e^{-\nu_b}}{N_{\text{off}}!} \prod_{i=1}^{N_{\text{off}}} \nu_b P_b(E_i, T_i),$$

where:

- $E_i, T_i$  are the energy and timing of the  $i$ -th event,
- $P_s(E, T)$  and  $P_b(E, T)$  are normalized signal and background PDFs,
- $\nu_s$  and  $\nu_b$  are the expected number of signal and background events in the ROI.

The total likelihood is simply the product of the two likelihoods:

$$\mathcal{L}_{\text{Full}} = \mathcal{L}_{\text{On}} \times \mathcal{L}_{\text{Off}}. \quad (7.6)$$

This is what is meant by fitting the off-beam and on-beam data simultaneously. In doing so the amplitude of the steady-state background is constrained by an additional

measurement, absent of the presence of signal. In practice this, the maximization of the likelihood is impractical, and instead the negative logarithm of the likelihood is minimized. In this work this is done through the software package `iminuit` [74].

The likelihood fit yields

$$N_{\text{CEvNS}} = 20.6_{-6.3}^{+7.1} \text{ signal-like events.}$$

To emphasize, the beam-related neutron background expectation was  $N_{\text{BRN}} = 0.67 \pm 0.34$ , and was not included in the background PDF.

The no-CEvNS hypothesis was rejected at  $3.9\sigma$ , confirmed via toy Monte Carlo studies.

A complementary counting analysis over an optimized region of interest. In particular the energy window was  $[1.5, 8.5]$  keV and the timing window was  $[0, 8]$   $\mu\text{s}$ . This gave a consistent result of  $N_{\text{CEvNS}}^{\text{counting}} = 21.0 \pm 7.8$  signal-like events.

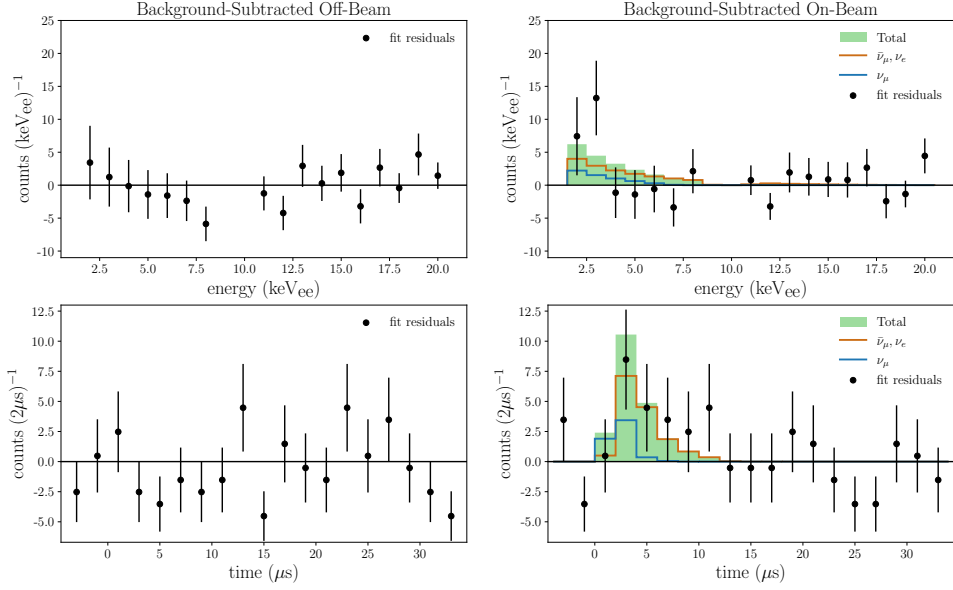
Figure 7.16 shows the background-subtracted energy and timing distributions for the on-beam data.

## 7.9 Discussion


Accounting for both statistical and systematic uncertainties, the measured signal agrees with the Standard Model prediction within  $1.95\sigma$ . This is seen in the profiled log-likelihood of the observed signal in Figure 7.17.

Having performed a detailed analysis and an exhaustive list of cross-checks, this discrepancy is likely attributable to a statistical under-fluctuation in the signal in a relatively small dataset. Nonetheless, the measurement shows great promise. The low-background nature of the experiment means that with only 20 counts it is possible to clearly separate signal from background.

There were several places with significant room for improvement in the analysis, but clearly from this result, the ultimate conclusion was simply that more data were necessary. Chapter 8 covers the novel contributions of this work stemming from a new, improved



**Figure 7.16:** Off-beam and on-beam data after subtracting the best-fit amplitude of the steady-state background PDF. Projections of the two-dimensional distribution are shown in energy and time. A statistically significant excess was observed. Figure adapted from [46].



**Figure 7.17:** The profiled log-likelihood of the amplitude of the signal distribution. The best-fit number of signal counts is  $1.9 \sigma$  below the Standard Model expectation. Figure adapted from [46].

measurement performed with an larger dataset.

# Chapter 8

## Campaign 3: Towards Precision

The dataset that informed the observation described in Chapter 7 was collected in a race against a long beam shutdown at the SNS. Going into the fall of 2023, the SNS entered into a one year outage to prepare the accelerator complex for the planned Proton Power Upgrade (PPU). Having just finished the first result on germanium, it was not possible to immediately collect more data.

This was a convenient time to reconfigure and improve the experiment. Campaign 2 revealed persistent leakage current in detectors, and the experimental set-up needed a few operational adjustments to remain compliant with the environmental constraints in Neutrino Alley. So throughout 2024, an effort was made to rehabilitate detectors at Duke and work on the shielding enclosure.

Detectors were shuttled back and forth from Oak Ridge to Durham so that cryostat vacuums could be serviced at Duke. Several detectors had persistent vacuum issues, related to the large, low-background cryostats used in this experiment. In the end, not all detectors were successfully rehabilitated at Duke by the time the SNS had completed its upgrade — though enough to begin collecting neutrino data once again.

In the fall of 2024, the experiment was recommissioned, and starting in February of 2025, Campaign 3 began in earnest. Data-taking continued for three and a half months into the end of May, which resulted in the most sensitive measurement of CEvNS to-date.

This chapter follows a similar structure to Chapter 7, as the key features of the CEvNS measurement are largely the same. Details of this measurement that are unchanged are briefly summarized, while more emphasis is placed on the significant improvements that were made.

In particular, there were two major goals in the effort of improving the analysis:

- To lower the analysis threshold from 1.5 keV to 0.5 keV.

- To lower the backgrounds by applying pulse shape discrimination to remove non-neutrino related surface events.

These improvements were achieved. In combination with the significantly larger exposure, this ushered in an era percent level precision measurements of CEvNS.

## 8.1 Data Collected

The Campaign 3 dataset was collected between February 15 and May 27 of 2025. Following the PPU the SNS average power delivered started at 1.7 MW, and ramped to 1.8 MW on April 16. The PPU also introduced a new operating regime with a proton energy of 1.3 GeV.

For the Campaign 3 dataset an exposure of  $4.68 \times 10^{22}$  POT were recorded. This was over twice the total number of POT recorded in Campaign 2.

### 8.1.1 Triggering Scheme


The key features of the triggering scheme were unchanged from Campaign 2. Figure 7.2 illustrates the arrival of on-beam and off-beam triggers that are used to search for neutrinos and sample the steady-state backgrounds present.

Waveforms are recorded at 120 Hz, externally triggered on timing signals from the SNS which indicate the presence of neutrinos.

A minor adaptation was made in the relative phase of the waveforms, with respect to the timing signals of the SNS, to increase the amount of pre-trigger baseline recorded, and to improve the resolution of the electronic noise. This is illustrated graphically in Figure 8.1. Where, previously, the expected neutrino arrival was at sample number 13,701 out of 22,000 samples, now the expected neutrino arrival was at sample number 14,975. As a result the total baseline estimating window went from  $105.6 \mu\text{s}$  to  $112 \mu\text{s}$ .

### 8.1.2 Exposure Summary

The combination of the detector rehabilitation and the work on the shielding during the outage resulted in much more stable operation and data-collection. All detectors used in



**Figure 8.1:** An example waveform from the externally triggered off-beam dataset. This event illustrates the main regions of the waveform, as they are used in the analysis, the pre-trigger baseline estimation window, the signal arrival region, and regions not analyzed to avoid edge effects and inefficiencies in reconstruction. The main adaptation from Campaign 2 is a relative phase shift to increase the pre-trigger window. For on-beam triggers, the expected start of neutrino arrival is highlighted.

this analysis were operating for the same amount of time.


A total exposure of 28.83 GWhkg was recorded before all data selection and analysis cuts. Similar to the Campaign 2 analysis a series of cuts reduce this down to a final exposure of 23.65 GWhkg.

### 8.1.3 Calibration Studies

There were, again, regular internally triggered calibration studies performed. These confirmed a high degree of gain stability, and were an important measure of the steady-state background.

### 8.1.4 Data Selection

With the increased exposure, beam stability acceptance was more strict in this analysis than in the Campaign 2 analysis. Beam periods were only considered in the analysis if the SNS reported average beam power was 80% of the target power at that time. This is illustrated



**Figure 8.2:** Average recorded beam power delivered during Campaign 3 as a function of time. The beam power was mostly stable, allowing for quality data recording. A ramp in power occurred on April 16, 2025 from 1.7 MW to 1.8 MW.

in Figure 8.2. Finally, periods of rapid change, defined as a change larger than 0.3 MW in roughly one minute, were removed.

Similarly, the beam energy was found to be highly stable at 1.3 GeV. Beam stability cuts incurred a 5% dead-time — similar to Campaign 2.

### 8.1.5 Blinded Analysis

Finally, the analysis was blinded in a similar manner to Campaign 2. All analysis routines were developed exclusively on the off-beam and calibration datasets. This was particularly important with the introduction of new analysis routines.

## 8.2 HPGe Waveform Analysis

Aside from the relative phase of the neutrino arrival in the waveform, the key features of the waveform digitization are unchanged from Campaign 2. Waveforms have total record lengths of 22,000 samples, corresponding to 176 s.

The two major differences in the waveform analysis are an improvement in the onset ( $t_{50}$ ) finding algorithm, and the development of a background reduction technique using pulse shape discrimination (PSD).

### 8.2.1 Pre-processing

In this analysis, the baseline fitting window is the first 14,000 samples. No other aspect of the baseline subtraction routine has changed.

### 8.2.2 Energy Estimation

The essential features of the energy estimation are also unchanged from Campaign 2. The leakage currents were greatly improved by the detector rehabilitation, and proved more stable throughout the data-taking period; therefore, a single filter optimization was performed, based on noise measured during the middle of the data-taking period.

The results, and trapezoidal filter parameters are summarized in Table 8.1. All the rise-times used in the trapezoidal filters applied are significantly longer than in Campaign 2 — and closer to the manufacturer specification. The longer optimal rise-times are characteristic of a decrease in parallel noise from reduced leakage currents.

### 8.2.3 Timing Extraction

The bespoke  $t_{50}$  finding algorithm used in Campaign 2 was the limiting factor in the determination of the 1.5 keV threshold. To lower the threshold to 0.5 keV it was necessary to reconsider this approach.

After testing several digital signal processing techniques, the analysis adopted an efficient

**Table 8.1:** The optimal rise times used to reconstruct the energy of HPGe signals in each detector is reported alongside the resulting ENC. The optimal rise times are characteristically longer in Campaign 3 compared to Campaign 2 due to the improved leakage currents reducing the contributions of parallel noise.

Detector	Rise Time ( $\mu\text{s}$ )	Pulser FWHM (eV)
Ge21	9.0	98(1)
Ge25	10.0	174(1)
Ge26	8.0	130(1)
Ge28	12.0	101(1)

and statistically well-defined method: the matched filter. Matched filters are a powerful linear filter for finding the onset of an expected signal in the presence of noise [75].

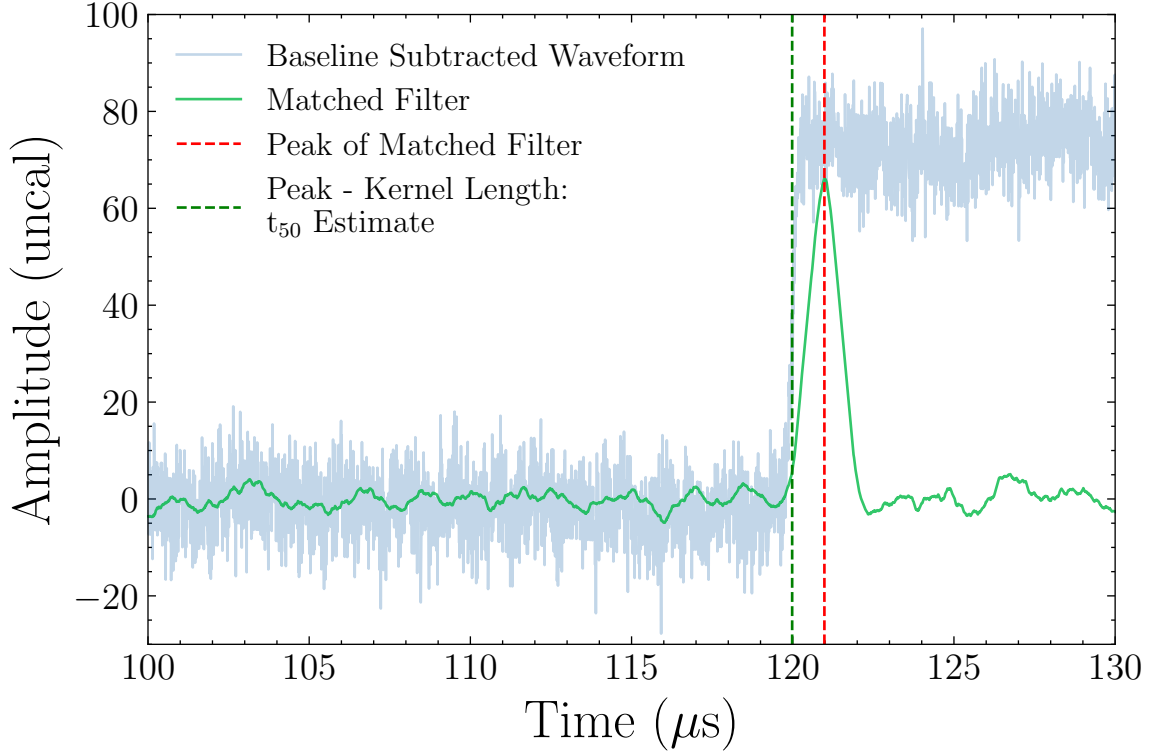
Matched filters are, like trapezoidal filters, Linear Time Invariant (LTI) systems. But, instead of shaping the signal for energy estimation, they correlate the waveform with a predefined signal template to maximize a statistic that identifies the most probable time of signal onset.

To estimate the signal onset, or  $t_{50}$ , a simple, bipolar signal template was chosen. The discrete template  $s[n]$  is defined as:

$$s[n] = \begin{cases} -1, & 0 \leq n < 125 \\ 0, & n = 125 \\ 1, & 126 \leq n < 251, \end{cases} \quad (8.1)$$

where each sample corresponds to the digitization sampling of 8 ns. Therefore, the negative and positive regions span 1  $\mu\text{s}$ . It turns out this is equivalent to a (relatively fast) trapezoidal filter with nearly no flat time. This is effectively identical to what is sometimes called a “triangle” filter in some analog shapers.

By way of developing some intuition for this filter, the vanishingly small “flat”-time ensures that the convolution is particularly sensitive to the rising edge of the waveforms considered. Longer flat-times were empirically found to introduce more error in the reconstruction of




**Figure 8.3:** A HPGe waveform (blue) and matched filter output (green) are shown. The red dashed line marks the filter maximum, while the green dashed line indicates the corresponding  $t_{50}$  estimate, defined as the maximum minus the kernel length. The matched filter yields an optimal estimate of the signal onset in the presence of noise.

the  $t_{50}$  parameter.

The matched filter output is obtained by convolving this kernel with the baseline-subtracted waveform. The onset parameter  $t_{50}$  is then defined as the index of the maximum of this convolution, offset by the kernel’s rise time (125 samples in this implementation). Figure 8.3 illustrates this procedure for a representative waveform, highlighting the matched-filter response and the location of the reconstructed onset.

Similar to the analysis of Campaign 2, the timing extraction method chosen was evaluated. The precision of the  $t_{50}$  measurement depends strongly on the signal-to-noise ratio present in the waveform, and therefore on the deposited energy. To quantify the energy dependence of this resolution, artificial “toy” waveforms are generated by combining an analytic signal model with noise traces harvested from externally triggered events that contain no physical



**Figure 8.4:** Representative distributions of the error in reconstruction of pulse onset as compared to the onset of the noiseless analytic waveform model for one detector (Ge-25). Distributions across detectors are all well approximated by gaussian distributions and comparable resolution. The onset finding resolution converges to the resolution of digitizer sampling.


pulses. To characterize the onset finding performance as a function of energy, toy waveforms are produced across a range of energies from [0.5 keV, 20 keV]. Each toy waveform has associated with it a true, noiseless- $t_{50}$ , to which the reconstructed  $t_{50}$  can be compared. This was evaluated on a sample of approximately 10,000 waveforms at each toy waveform amplitude. Example distributions at a few energies are shown in Figure 8.4, where the energies span the full range of the region of interest.

The distributions of reconstructed  $t_{50}$  are observed to be symmetric and approximately normally distributed around the expected value. Figure 8.5 shows, for another detector, the distributions of  $t_{50}$  onset reconstruction versus energy for a series of toy pulse amplitudes.

The relative error in the reconstruction of pulses is significantly improved from Campaign 2, and pulses are well reconstructed down to 0.5 keV.

## 8.2.4 Pulse Shape Discrimination

For the CEvNS events of interest in this analysis, only interactions that occur inside the bulk of the HPGe crystal constitute a viable signal. Energy deposited near the outside

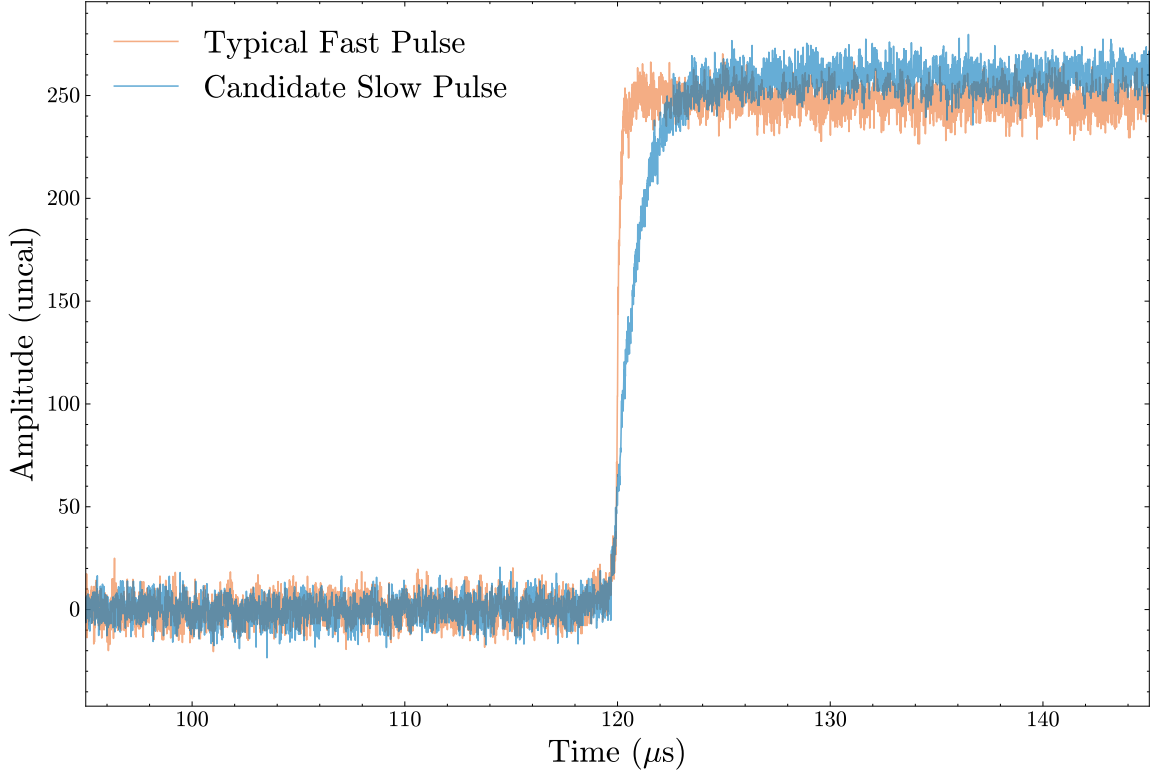


**Figure 8.5:** Energy vs onset reconstruction error for the analytic toy waveform study described above.

surface produce *surface events* which cannot be fully reconstructed. The degraded events can be reconstructed in the CEvNS energy region of interest, mimicking signal-like events. Discriminating between these two populations can reduce the steady-state background significantly.

The detectors used in this work are p-type point-contact (PPC) diodes. The n+ contact is formed by diffusing a thin layer of lithium that collects on the outer surfaces of the crystal. Inside this  $\sim 1$  mm thick lithium layer, the electric field used to drift charges is weaker compared to the crystal bulk. When ionizing radiation occurs on or near the surface, the charge carriers diffuse slowly out of the weak-field surface before reaching the bulk and then making their way to the point-contact. This causes two main pathologies for these surface events:

1. *Charge loss:* Not all charge will make it out of this weak field region. It is possible to measure a transition profile on the surface that corresponds to the fractional charge collection. Regions with sufficiently weak fields can contribute no charge collection at all, due to effects like recombination. For this particular analysis the exact transition



**Figure 8.6:** Two waveforms with different characteristic rise-times are shown to emphasize the difference between fast, signal-like pulses (orange) and slow, surfac events (blue).

profile is not relevant, beyond an estimate of the total active mass.

2. *Slow rise time:* After the slow diffusion out of the weak field regions on the surface, charge-carriers make their way to the point-contact. The resultant shape of the waveforms collected exhibits a long profiled rise-time.

By comparison, the events that occur in the bulk benefit from full charge collection and, after the typical drift-times in the bulk, prompt arrival of charge. Representative waveforms from bulk events (fast pulses) and surface events (slow pulses) are shown in Figure 8.6. The goal of the PSD in the analysis is to effectively differentiate between the slow pulses and fast pulses. Of equal importance is the need to be able to evaluate the signal acceptance to identify potential losses in signal.

The primary shape-based distinction between fast pulses and slow pulses is in the signal rise time. To exploit this difference, this analysis uses a PSD parameter based on the ratio


of two shaping filter outputs: a so-called **T/E** parameter. This ratio is constructed using two filters with distinct shaping times:

- The **T filter** is a fast, symmetric, triangular shaping filter. It is the same filter used in the onset-finding algorithm, and effectively serves as a matched filter for the rising edge of a fast signal. The peak of the T filter response can be interpreted as an energy estimator optimized for prompt charge collection. In this implementation, the filter rise time is  $1\ \mu\text{s}$ , corresponding to the response expected from a waveform with charge collection completed within that interval.
- The **E filter** is the trapezoidal filter optimized for energy resolution. It integrates the total collected charge over a longer shaping time, typically greater than  $8\ \mu\text{s}$ . Long shaping times are characteristic of energy-optimized filters because, in the presence of electronic noise, the uncertainty in the amplitude estimate decreases as  $1/\sqrt{N}$ , where  $N$  is the number of samples over which the signal is integrated. This results in improved resolution for signal-like events with prompt charge collection. As a by-product, the long shaping time also ensures that even delayed or slow-drifting charge from surface-like events is fully integrated, allowing the E filter to serve as a stable denominator for the T/E ratio.

The ratio T/E captures the relative promptness of the signal rise. For fast pulses, both filters yield comparable amplitudes, resulting in a T/E ratio close to unity. For slow pulses, the fast filter underestimates the energy due to incomplete charge arrival within its shaping window, leading to T/E values significantly less than one. This behavior is illustrated in Figure 8.7, where the filter responses to fast and slow pulses are shown.

To define PSD selection criteria and estimate the signal acceptance, it is necessary to characterize the expected distribution of the T/E parameter for a population of signal-like events.

In this context, signal-like events are defined as waveforms with the characteristics of interactions occurring within the detector bulk. The population of CEvNS events that must




**Figure 8.7:** An example slow pulse (left) and fast pulse (right), each displayed with the resulting T and E filters and their respective maximums. For the slow pulse, the T filter rises and falls significantly before the real amplitude of the charge collected, estimated by the E filter. For the signal-like, fast pulse the T and E filter give comparable estimates of the amplitude of charge collected. The ratio between these two identifies surface events.

survive any PSD is relatively small and blinded from the development of the parameter. As such, an alternative source of bulk-like events is required to characterize the T/E distribution.

For this purpose, intrinsic radioactivity of the HPGe itself was used to characterize the T/E distribution. Decays from isotopes such as  $^{71}\text{Ge}$  and  $^{65}\text{Zn}$  produce well-understood, uniformly distributed energy depositions within the detector bulk. These events exhibit waveform shapes consistent with prompt charge collection and thus serve as a reliable proxy for signal-like events.

The characterization procedure proceeded as follows:

1. A population of intrinsic background events known to originate within the detector bulk were selected. This is illustrated in Figure 8.8, which shows, for detector Ge-28, a histogram of the internally triggered calibration data energy distribution alongside the two dimensional view of energy versus the PSD parameter T/E. A tight band around the intrinsic 10.36 keV X-ray peak in the energy spectrum selects a very pure population of signal-like events. The waveforms from these events are recorded as a population of signal-like events.




**Figure 8.8:** Example of T/E distributions using internally triggered calibration data for Ge-28. The band near T/E of 1.0 is illustrative of signal-like, fast pulses originating in the bulk of the crystal.

2. An analytic equation was used to model the signal-like waveform shapes, parameterized by rise time. This equation is given by:


$$S(x) = \tanh\left(\frac{t - t_0}{\tau}\right), \quad (8.2)$$

where  $t_0$  is the signal onset and  $\tau$  parametrizes the rise. This equation is a simple re-parametrization of the sigmoid function used in Chapter 7, to match its use in the literature [76]. Using this model, the identified bulk event waveforms were fit to extract rise time distributions for each detector. An example of one of these fits is shown in Figure 8.9. These distributions, shown in Figure 8.10, sample the crystal volume uniformly and thus are expected to accurately model the distribution of rise-times of CEvNS events generated in the bulk.


3. For each detector, the rise-time distribution were sampled to create representative, signal-like toy waveforms. These templates were combined with real noise from externally triggered waveforms to compute the T/E parameter for each waveform across a range of pulse heights. An example of a distribution of the PSD parameter T/E for toy waveforms at a particular amplitude, weighted by the rise-time distribution



**Figure 8.9:** Example fit of a fast pulse extracted from the  $^{71}\text{Ge}$  x-ray to the analytic model described above.



**Figure 8.10:** Bulk event rise-time distributions as measured using uniformly distributed intrinsic radiation events.



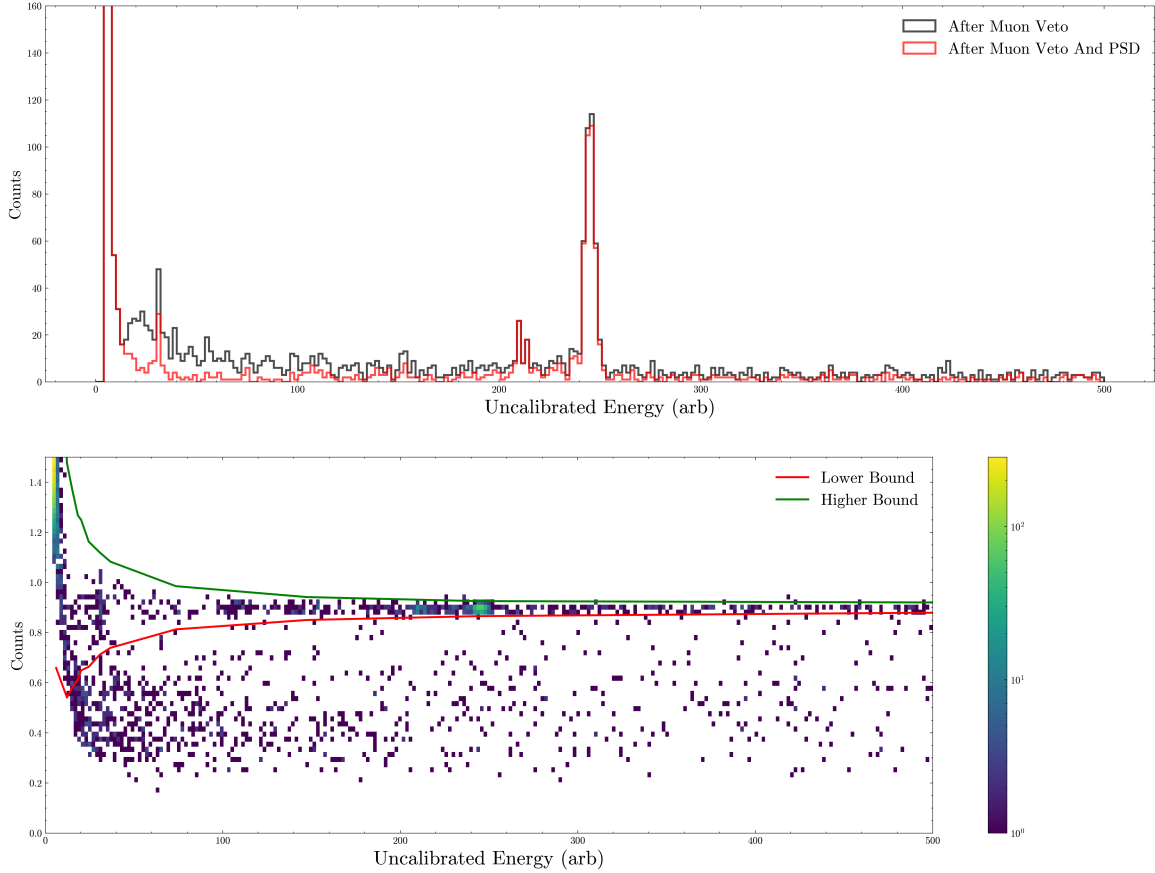
**Figure 8.11:** Distribution of T/E for samples of toy waveforms generated with signal-like rise times. Distributions are sampled across a series of rise-times and then weighted by the rise-time distributions measured. 99.7% credibility intervals define signal acceptance window as a function of energy.

is shown in Figure 8.11.

4. At a series of energies and for every detector, the T/E distributions are integrated to generate 99.7% credible intervals. These define the most probable regions where CEvNS events can lie in energy vs T/E. These curves, shown in Figure 8.12, define the acceptance regions for the PSD in the analysis.

This procedure yields energy-dependent T/E distributions representative of signal-like events for each detector. These distributions form the basis for constructing signal acceptance bands, shown in Figure 8.12, which define the PSD selection applied in the analysis.

For each energy evaluated in this procedure, the lower boundaries of the 99.7% T/E acceptance region is determined through a brute force integration method. In particular, events falling below the signal-acceptance region are slow-pulse candidates and can be safely



**Figure 8.12:** Application of 99.7% signal-accepting PSD parameter on internally triggered calibration data for Ge21.

rejected.

At the lowest energies, the ability to separate fast and slow pulses degrade and the populations merge together. The combination of the onset finding algorithm and the application of the PSD as defined in this section inform the 0.5 keV analysis threshold.

### 8.3 Data Quality and Preparation

Following the general outline of data-cleaning established in Campaign 2, a series of cuts must be considered to arrive at the final datasets used in the statistical analysis. Importantly, the resultant dead-time of all cuts must be determined. For this analysis, it must also be determined that all cuts applied have negligible impacts on signal acceptance. As a rule

of thumb, the backgrounds in this experiment are sufficiently low that there is no need for aggressive cuts that risk reducing the signal. As such all cuts are designed and ensured to retain all CEvNS events.

This section summarizes the cuts used in this analysis, and their impacts on the live-time. Where possible, the cuts that are unchanged from Campaign 2 are briefly summarized.

### **8.3.1 Timing Based Cuts**

#### **Pre-amplifier Reset**

Pre-amplifier resets and their subsequent baseline restores are still present in the data. The internally triggered logic signals indicating the presence of a reset are used to defined 200  $\mu$ s windows where data is ignored.

In general, detector rehabilitation of cryostat vacuums reduced the leakage currents in detectors so that the incurred dead-time from the regular resets reduced from Campaign 2.

#### **LN2 Fill**

LN2 Dewar refills occurred at regular daily intervals, as in Campaign 2. Events recorded during LN2 refills were cut from the analysis due to vibration induced noise. Similar to Campaign 2, the dead-time incurred by this removal is approximately 5%. There were slight differences in the individual detector dead-times as some detectors required longer exclusion windows after fills.

#### **Additional Noisy Periods**

On top of the expected fill periods, several other distinct periods were identified in the blinded analysis that indicated a high degree of microphonics-related noise not correlated to any fill. These events were clearly correlated in time and far above the expected backgrounds from typical rates. These time periods were identified in the data and are displayed in table 8.2. These periods were removed from the analysis and had small impacts on the overall live-time.

**Table 8.2:** Additional periods of microphonics-related noise identified during the blinded analysis. These intervals, not associated with cryogen fills and containing large, correlated background rates, were excluded.

Micro-phonics noise		
Date	Start(unix)	Dur. (s)
February 23	1740328256	735
April 3	1743688800	3600
April 3	1745174658	5400
April 20	11745146368	30000
April 21	1745240489	30000
April 26-27	1745691120	65880

### Cross-Talk

One of the major improvements in this analysis from Campaign 2 was the identification of the source of the “transient” noise after the data-taking period. Once revealed to be cross-talk between the energy channels of the HPGe detectors and the logic signals from a different detector’s pre-amplifier signaling a reset, the coincidence window was extended to all detectors. In a 200  $\mu$ s window around any detector reset, all channels from all detectors were ignored.

### 8.3.2 Waveform Based Cuts

#### Saturated Events

Similar to Campaign 2, some events recorded could not be reconstructed accurately due to a saturation of the configured dynamic range of the digitizer. There were both positive and negative saturation, and were often related to a pre-amplifier resets that did not have an accompanying TTL signal. In any case these events were removed from the dataset.

#### Stray Microphonics

Apart from the dedicated time-based cuts, one detector (Ge-25) was particularly sensitive to a long-tailing and persistent population of microphonic events. A conservative waveform

based cut was developed to further reduce this population.

The parameter definition was related to linear regressions of different components of the pre-trigger baseline. A short baseline window was defined as a linear regression of the first 5,000 samples and the standard baseline used the first 14,000 samples. The difference between the slopes of the short and standard window were used to identify pathological, noisy waveforms that exhibited bad baselines. The particular parameter definition was:

$$\left| \frac{m_s - m_l}{m_l} \right|,$$

where  $m_s$  and  $m_l$  are the short and long baseline slopes, respectively. Events were rejected if they had a slope parameter value less than 0.01 and larger than 0.25. These values were picked to be conservative and were ensured through handscanning of all background events to not remove signal-like events.

As a result, it does not aggressively reduce all microphonics events in Ge-25, but assists in reducing the background further.


### **Negative Pulses**

A subset of waveforms in this analysis exhibited anomalous negative-going pulses, the origin of which remains undetermined. These events, shown in Figure 8.13, distort the baseline and bias the energy reconstruction if left unremoved. Since they are un-physical artifacts rather than genuine energy deposits, they can be safely rejected without impacting the signal acceptance.

The trapezoidal filter minimum parameter, described in Chapter 7, was used again to identify and remove such events. Specifically, events with a trapezoidal-filtered waveform minimum exceeding the maximum were rejected using the parameter definition

$$\left| \frac{\max(\text{Trap-Filtered Waveform})}{\min(\text{Trap-Filtered Waveform})} \right|,$$

with a cut value of 0.5 applied in the analysis.



**Figure 8.13:** Example of a digitized waveform in an HPGe detector containing a so-called “negative pulse.” These anomalous events lead to faulty baseline reconstruction and biased energy estimates, and are therefore excluded from the analysis.

## PSD

One of the major developments in this analysis was the application of PSD. The cut is defined in more detail in Section 8.2.


### 8.3.3 Muon Veto

The essential features of the physics muon veto system are largely the same from Campaign 2. In the analysis, efforts were made to improve the live-time of the experiment by optimizing the veto conditions.

Specifically, energy thresholds were applied to the plastic scintillator signals inside the muon-veto as a part of the veto logic. HPGe detector signals were only rejected if, beyond the time coincidence conditions described in Campaign 2, the energy in scintillator was above a channel-dependent threshold. This is graphically illustrated in Figure 8.14.

### 8.3.4 Dead-Time

The cuts used in this analysis are summarized in Table 8.3. In general, the key features of this analysis are the lowered energy threshold and the improved background reduction



**Figure 8.14:** An example distribution of energy deposited one of the scintillator panels that comprises the muon veto. The large, low energy population to the left of the dashed red line is largely comprised of gamma rays and not likely to be correlated to a background in the HPGe detectors. Imposing energy thresholds on the coincidence reduces dead-time. Original figure courtesy of Emma van Nieuwenhuizen.

techniques. It is worth mentioning that particular effort was made to improve the total live-time of the experiment as well. From Campaign 2 to Campaign 3, approximately 5% of the live-time was recovered.

**Table 8.3:** Summary of cuts in Campaign-3 and the resultant dead-time. Future data-taking campaigns can address the cross-talk driven dead-time and lower the muon-veto dead-time with more aggressive coincidence windows. The LN2 fill dead-time is driven by long, daily fills. The length of Dewar refills is limited by operational safety constraints.

Cut	Type	Dead-Time	relative to Campaign-2
LN2 Fill	Timing	$\sim 5\%$	similar, adapted window
Muon-Veto	Timing	$\sim 8\%$	similar
TRP Inhibit	Timing	$\sim 1\%$	includes cross talk replaces trap min cut
Noisy Period/Microphonics	Timing	$\ll 1\%$	same
Missed Resets	Waveform	$\sim 1\%$	same
Negative Event	Waveform	$\ll 1\%$	new
Saturated Event	Waveform	$\ll 1\%$	same
Event Energy	Waveform	NA	lower threshold
Event Time	Waveform	NA	44 mus instead of 40 mus
Event PSD	Waveform	NA	new

## 8.4 Backgrounds

This section reviews the backgrounds present in the data collected, after all cuts. It is necessary to consider both steady-state backgrounds and beam-related backgrounds, and both require modeling.

The steady-state backgrounds are evaluated using off-beam data and calibration data.

### 8.4.1 Off-Beam Data After All Cuts

Before all cuts described in Section 8.3, there are approximately 380 million off-beam triggers and 380 million on-beam triggers<sup>1</sup>. After all data cleaning cuts, LN fill period removals, and

<sup>1</sup> As a reminder, because of the external triggering scheme all events are recorded with 100% trigger efficiency. A threshold is applied in the analysis after the fact

the muon-veto cut there are roughly 330 million triggers remaining corresponding to the live-time of 88%.

From there, events are only considered with energies in the range  $[0.5 - 20]$  keV and reconstructed pulse onsets between samples  $[14,500 - 20,000]$ , which corresponds to the time range  $[-4 - 40]$   $\mu s$  centered around the externally triggered SNS timing signal indicating protons on target. Due to cable delays for various timing signals, the companion on-beam triggers expect neutrino arrival to begin at sample 14,980, or pulse onset of -160 ns.

Finally, PSD is applied to arrive at the total off-beam dataset. The events left in the energy and time region of interest (ROI) after all cuts from the original total off-beam triggers corresponds to a nearly one-to-one exposure with the companion on-beam dataset. Figure 8.15 shows the energy spectra, integrated over the 44  $\mu s$  analysis window, for all detectors.


It must be emphasized that these histograms contain all the events considered in the off-beam dataset collected over almost four months of beam data. The key feature consistent across all detectors is the remarkably low backgrounds. The central challenge of the statistical analysis is modeling the background accurately with so few counts. In this experiment there is no detailed model or simulation of the backgrounds. The backgrounds are instead directly measured, and simple models are used to capture the main features present.

### **Key Features of Background Spectrum**


The background model can be motivated by a few heuristic features in the recorded data. The energy and timing of the steady-state background are expected to be uncorrelated and so they are discussed here separately.

#### **Time**


The first thing to consider is the distribution of steady state background events in time. The expected reconstructed  $t_{50}$  distributions are flat for each detector. In general the steady state background events are expected to follow a Poisson distribution but the rate of backgrounds in this experiment is so low that the distribution of events in a particular




(a) Ge-21



(b) Ge-25



(c) Ge-26



(d) Ge-28


**Figure 8.15:** Full Campaign 3, externally triggered background energy spectra for HPGe detectors. The integral of restricted energy range is reported to emphasize the low-energy backgrounds outside of the range of the intrinsic X-rays. Ge-25 has a low energy excess background stemming from particular sensitivity to microphonics events.

trigger window is flat.


The reconstructed events for each detector in time are shown in Figure 8.16. The variation in total background index is observed, but all backgrounds in detectors are well modeled by a flat distribution.

### Energy


In the energy spectra there are three key features of the steady state backgrounds that are modeled in this analysis. These features are typical of low energy, low background germanium based experiments. They have been extensively described in, for example, a thorough background modeling of another CEvNS experiment [77]. They are summarized,




(a) Ge21



(b) Ge25



(c) Ge26



(d) Ge28

**Figure 8.16:** Full Campaign 3, externally triggered background pulse onsets for HPGe detectors. The onset of steady-state backgrounds is well-modeled by a flat distribution for all detectors.

broadly, as:

- **Continuum:** In the continuum between roughly 1 keV and 20 keV the distribution is expected to rise slowly towards low energies due to a mixture of down scattered gammas and muon induced neutron scatters. This slow, monotonic increase is well modeled by an exponential decay with a small decay constant.
- **Low energy rise:** At the lowest energies, below 1 keV, there is a separate, faster rise of background events seen in all detectors. These events are a mixture of surface radiation events not filtered by the PSD and non-radiation events (such as stray microphonics) which survive all data-cleaning cuts. Each detector sees a different amplitude in this low energy rise, but all distributions are well modeled by an exponential decay.

- **Peaks:** At known positions in the energy spectra, there are a series of intrinsic sources of radiation which are seen as peaks in the energy distributions. Most prominent is the  $^{71}\text{Ge}$  x-ray at 10.36 keV seen in all externally triggered spectra. The K and L shells of all x-ray peaks are well modeled by gaussian distributions.

With this understanding, the backgrounds in energy are well described by the following composite model:

$$B(E) = N_{low}e^{-\lambda_{low}E} + N_{high}e^{-\lambda_{high}E} + \sum N_{peak}^i e^{-(E-\mu_i)^2/2\sigma_i^2} \quad (8.3)$$

where:

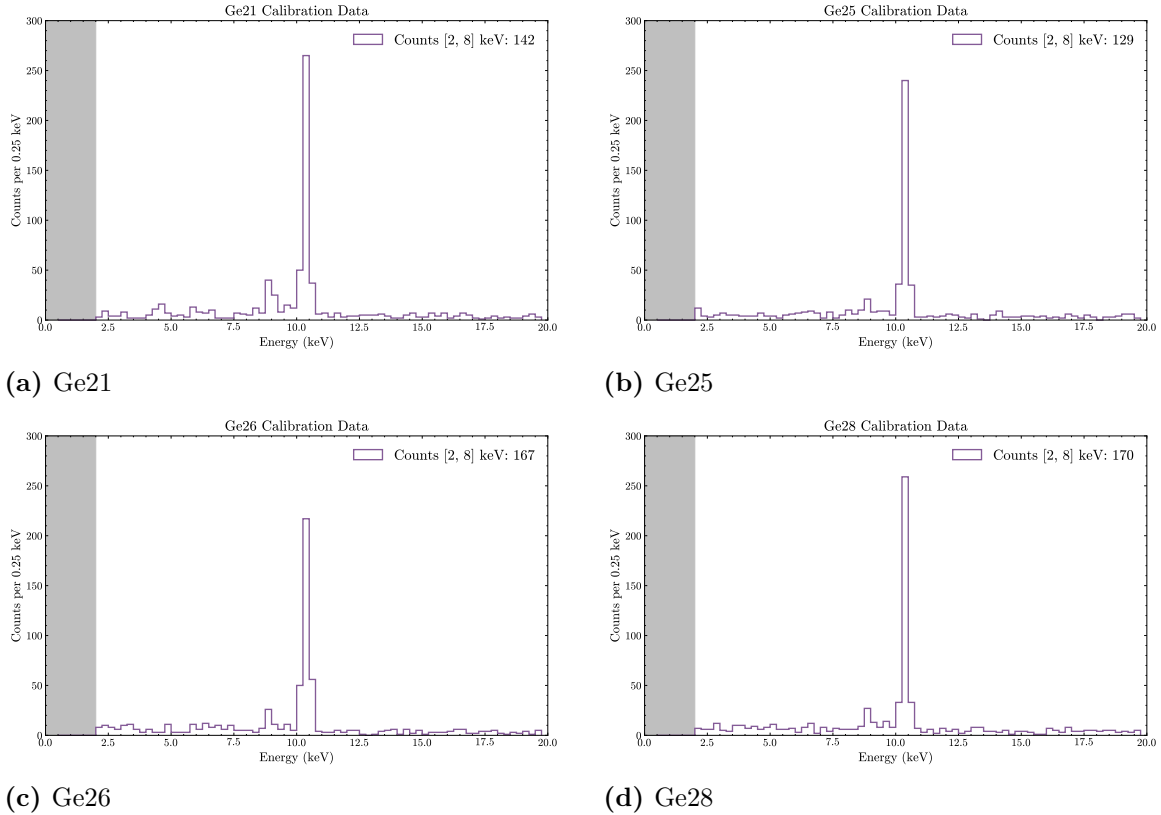
- $N_{low}$  and  $N_{high}$  represent the amplitude of the exponential decays
- $\lambda_{low}$  and  $\lambda_{high}$  describe the shape of the exponential decays
- $N_{peak}^i$  represents the amplitudes of the low energy x-rays that are present in the background data
- $\mu_i$  and  $\sigma_i$  are the peak positions and resolutions for each x-ray

This function,  $B$ , is used to model the steady-state background data is part of the composite signal model in the on-beam data.

#### 8.4.2 High Statistics Internally Triggered Dataset

On its own, the low statistics in the off-beam dataset has limited power to constrain the background model parameters. The higher-statistics calibration data are used along-side the off-beam data to better constrain these parameters.

The full datasets of internally triggered calibration data are shown in Figure 8.17. These datasets are were collected approximately once a week throughout full Campaign 3 data-taking period, and are assumed to be representative of the steady-state background distributions seen in the off-beam dataset.



**Figure 8.17:** Combined, internally triggered calibration data spectra for Ge detectors. With the higher-statistics internally triggered data the K shells of  $^{71}\text{Ge}$  and  $^{65}\text{Zn}$  are both visible. The region below 2 keV is excluded in the analysis of the internally triggered data because of the trigger efficiency which is not present in the beam data. Also reported are the integrals of the continua from 2 to 8 keV in each detector highlighting the variation in the steady-state background.

However, the trigger efficiency for the **internally triggered** data at low energies is unknown. It cannot be used to reliably describe the low energy features. This is represented graphically as a gray band below 2 keV. The calibration data below 2 keV are not considered for background modeling.

There are several different ways to make use of this dataset to help constrain the free model parameters that describe the steady-state background in the statistical analysis. The method that is chosen in this analysis is a simultaneous fit of both the off-beam data and the calibration data sharing model parameters that are visible in both datasets. Table 8.4 highlights which model parameters are shared between the two datasets and which model

parameters are fixed during this analysis. Model parameters are fixed if the exact values are known from literature, such as the peak positions and branching ratios of K and L shell X-rays. Model parameters are shared if they are present above 2 keV. Features below 2 keV, such as the low energy rise, cannot be shared between the two datasets and are constrained exclusively by the off-beam data.

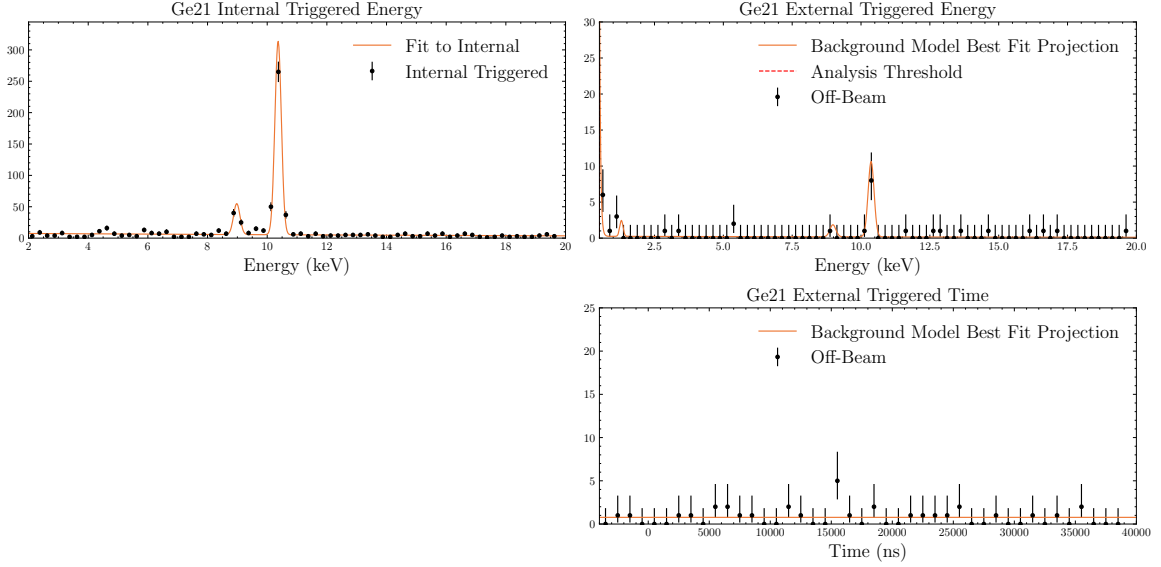
**Table 8.4:** Model parameters and whether they are shared or fixed in the Campaign-3 statistical analysis.

Parameter	Shared	Fixed
$N_{\text{low}}$	×	×
$\lambda_{\text{low}}$	×	×
$N_{\text{high}}$	✓	×
$\lambda_{\text{high}}$	✓	×
${}^{71}\text{Ge}_K \text{ N}$	✓	×
${}^{71}\text{Ge}_K \mu$	✓	✓
${}^{71}\text{Ge}_K \sigma$	✓	×
${}^{71}\text{Ge}_L \text{ N}$	×	✓
${}^{71}\text{Ge}_L \mu$	×	✓
${}^{71}\text{Ge}_L \sigma$	×	×
${}^{65}\text{Zn}_K \text{ N}$	✓	×
${}^{65}\text{Zn}_K \mu$	✓	✓
${}^{65}\text{Zn}_K \sigma$	✓	×

### 8.4.3 Simultaneous Fits to Background Data

The results of the steady-state background analysis are shown here to highlight the quality of the model and parameter estimation relevant for the final statistical analysis.

Figures 8.18 through 8.21 show the results of the simultaneous fits to the background model. The fits to the off-beam data are 2D dimensional fits with a constant background, while the calibration data is a one-dimensional fit in energy only. The fits are in excellent agreement with the Poisson statistics of the counts distributed in the spectra and the simultaneous fits greatly improve the constraints on the parameter estimation. Further



**Figure 8.18:** Simultaneous fit of calibration data and off-beam data for Ge21.

details on the results of the statistical analysis are described in Section 8.6.

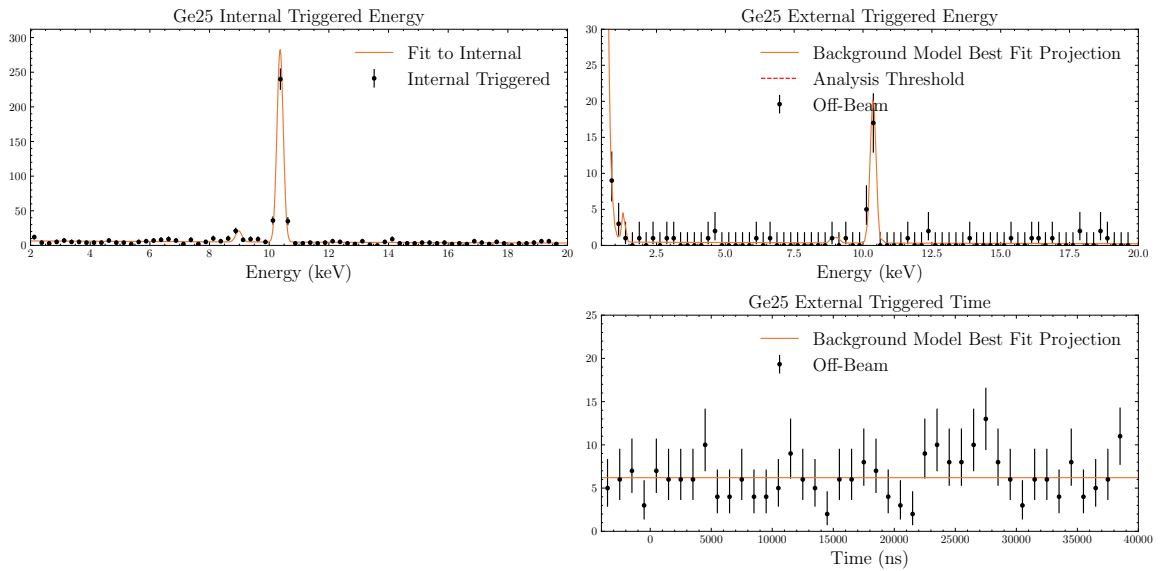
#### 8.4.4 Beam Related Backgrounds

The beam-related neutron backgrounds were estimated for Campaign 2 with a dedicated MCNP simulation. The estimated total counts across all detectors in Campaign 2 was approximately 2% of the total Standard Model CEvNS expectation. For Campaign 3, with the increased exposure, this background was expected to grow slightly. Constrained by this simulation, in the energy region of interest for this analysis,  $2.2 \pm 0.6$  counts are expected. This accounts for less than 2% of the number of CEvNS counts.

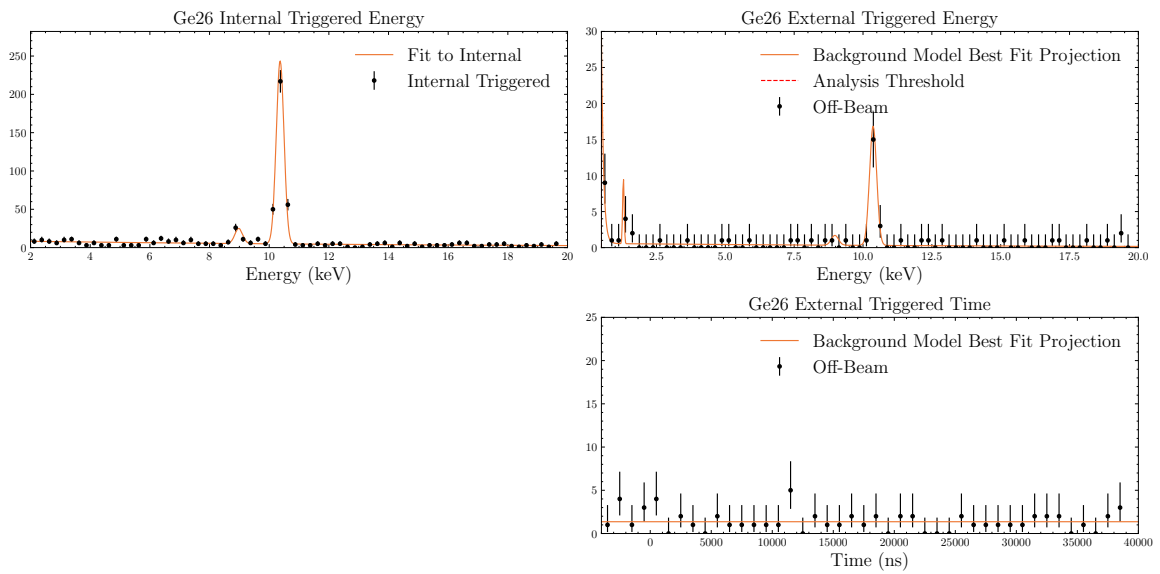
In summary, the total beam-related backgrounds are considered negligible for this analysis were 2% is still significantly lower than the systematic or statistical uncertainties considered. Future analyses will need to contend with this background – and its own uncertainty – in the pursuit of a more precise CEvNS measurement.

#### 8.4.5 Background Model

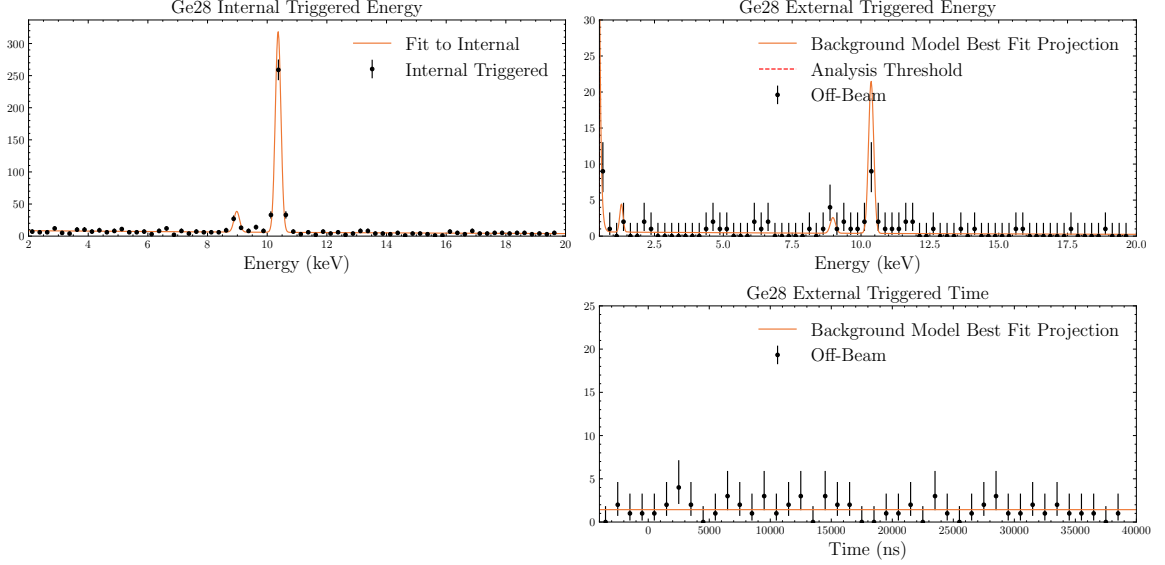
Having characterized the steady state backgrounds and put a limit on the beam related backgrounds, the final 2D background model is summarized as:



**Figure 8.19:** Simultaneous fit of calibration data and off-beam data for Ge25.



**Figure 8.20:** Simultaneous fit of calibration data and off-beam data for Ge26.



**Figure 8.21:** Simultaneous fit of calibration data and off-beam data for Ge28.

$$B(E, t) = N_{low}e^{-\lambda_{low}E} + N_{high}e^{-\lambda_{high}E} + \sum N_{peak}^i e^{-(E-\mu_i)^2/2\sigma_i^2} \quad (8.4)$$


This model is taken to be the probability distribution function in energy and time of all backgrounds present in this analysis. It is used in the statistical analysis in the simultaneous fits with the on-beam data.

## 8.5 Signal Expectation

### 8.5.1 Overview

The spirit of this section remains very similar to the signal expectation work of Campaign 2. A few details are worth an explicit mention that have changed or improved:

- The shape of the proton-on-target traces from the SNS were re-evaluated after the **PPU!** (**PPU!**) to ensure a faithful neutrino flux input.
- The new and improved onset finding algorithm introduces a new energy-dependent smearing of pulse onset. This is handled more rigorously in this analysis as an energy-dependent smearing, rather than an approximate constant.



**Figure 8.22:** Average beam trace over Campaign-3 after application of the beam stability cut. A FWHM of  $\sim 350$  ns is estimated.

- The energy regions that were previously excluded from the statistical analysis are retained in the signal PDF, due to better modeling of the backgrounds.


## 8.5.2 Improved Smearing

### POT Trace

The average proton-on-target trace was re-evaluated over the course of the Campaign 3 dataset to serve as a faithful starting point in the generation of two-dimensional neutrino fluxes. This two dimensional flux is integrated over to generate recoil spectra and so it is important the time dimension is considered carefully.

The collaboration records a limited sample of traces of the SNS proton bunches. This sample, limited to periods of stable beam operations, constitutes an accurate representation of the shape of proton bunches that generate the neutrinos for this measurement.

These POT were collated and averaged, weighted by beam-power, to produce a representative proton arrival distribution. This proton distribution is seen in Figure 8.22. The deviations from this average were taken to be small in comparison to the further smearing of predicted recoils described subsequently.



**Figure 8.23:** Differential recoil spectra, integrated over recoil energies with and without convolution over POT trace.


The convolution with the two-dimensional recoil spectrum is illustrated in a one-dimensional projection in time in Figure 8.23. As a point of clarification, the Campaign 2 analysis also used a POT convolved neutrino spectrum, but it was not specifically re-evaluated during that analysis. Provided the pulses are well-behaved, the changes from SNS run to SNS run are relatively small.

### Energy and Time

The energy smearing procedure is largely the same from Campaign 2, though the specific energy-dependent energy resolutions for each detector changed subsequent the detector rehabilitation and recommissioning.

The energy-dependent onset finding algorithm has an improved resolution at all energies. Compared to Campaign 2 recoil times were integrated over with an energy dependent smearing function, as opposed to an approximately constant resolution.

The energy dependent resolutions, taken together with the drift time distributions, encode a transformation of every point in a two dimensional recoil spectrum. An example of such a transformation is shown in Figure 8.24. After integrating over both recoil energy and




**Figure 8.24:** An example of smearing functions for a single point  $(T, t)$ . The energy is smeared by the energy resolution function and the neutrino scattering time is smeared by the drift time and recoil energy dependent drift time.

interaction time, all detector effects have been included.

### 8.5.3 Signal PDF

The resultant two-dimensional recoil spectrum, after all detector effects, is expected to model accurately the distributions of neutrino scatters observed in the data.

An example of projections in energy and time are seen in Figure 8.25 for one detector. The recoil distributions are decomposed by flavor, though in the first-pass statistical analysis only the combined recoil spectrum is considered. It is possible to see the different amplitudes of each flavor, which has its origins in the different energies of the neutrino spectra, and fixed relative weights are used to generate the final signal PDF. The relative amplitudes are the integrated ratios of the integrated distributions and are  $1 : 0.864 : 0.663$  for  $\nu_\mu : \nu_e : \bar{\nu}_\mu$ . These flavors of neutrinos – the dominant flavors produced at the SNS through the pion decay at rest chain – are produced in equal quantities. The fact that they are not observed in equal quantities is due to the kinematics of the decays resulting in different energy distributions. The different energies for each flavor give rise to different differential rates.



**Figure 8.25:** The final signal PDF used in the statistical analysis for expected signal in observable energy and pulse onset as measured from the start of neutrino arrival.

#### 8.5.4 Expected Counts

The total number of CEvNS events predicted depends on an exposure calculation related to the live-time of the dataset, after all cuts. This live-time is detector dependent and was considered for each detector separately.

Table 8.5 shows the final exposure before and after applying the beam power cut are presented. This exposure calculation depends in part on the active mass of the HPGe detectors. The same active masses as Campaign 2 are assumed.

The resultant Standard Model expected counts are summarized in Table 8.6.

## 8.6 Statistical Analysis

The statistical analysis is largely the same as that of Campaign 2 with a few key differences:

- With enough exposure to expect a stand-alone measurement in each detector it was of particular interest to perform individual detector measurements. Given the variable steady-state backgrounds in each detector, separate fits maximized the sensitivity of

**Table 8.5:** Total exposure before and after cuts in the unit of GWhkg. Exposure is similar between all detectors and variations are mainly attributed to excluded windows due to LN2 fills.

detector	all	with beam stability	with beam stability +dead time
	GWhkg	GWhkg	GWhkg
Ge21	7.208	6.862	5.905
Ge25	7.208	6.862	5.867
Ge26	7.189	6.843	5.929
Ge28	7.224	6.877	5.950
sum	28.829	27.444	23.650

**Table 8.6:** Standard model signal expectation for Campaign-3. Each detector alone expected a number of signal counts comparable to the total expected counts of Campaign 2.

detector	CEvNS counts [0.5,20.0]keV <sub>ee</sub>
Ge21	31.03
Ge25	30.84
Ge26	31.15
Ge28	31.26
total	124.28

the experiment.

- The internally triggered calibration data was simultaneously fit with the off-beam and on-beam datasets.

As a result, each detector is considered as its own sub-experiment with three datasets given by:  $\{x_{cal}^i, x_{off}^i, x_{on}^i\}$ . Each experiment is described by a joint likelihood  $\mathcal{L}_{\text{joint}} = \mathcal{L}_{\text{cal}}\mathcal{L}_{\text{off}}\mathcal{L}_{\text{on}}$  which is the product of the three likelihoods that describe calibration data, off-beam data, and on-beam data.

For each sub-experiment it is possible to determine an individual maximum likelihood estimate that gives a best-fit value of  $N_{CEvNS}$ . There are several nuisance parameters related to the amplitudes and shapes of components of the background model.

### 8.6.1 Results

The on-beam data above 0.5 keVee was unblinded on the 12th of September 2025. Best fits carried out for each detector separately are shown in Figure 8.26 to Figure 8.29.

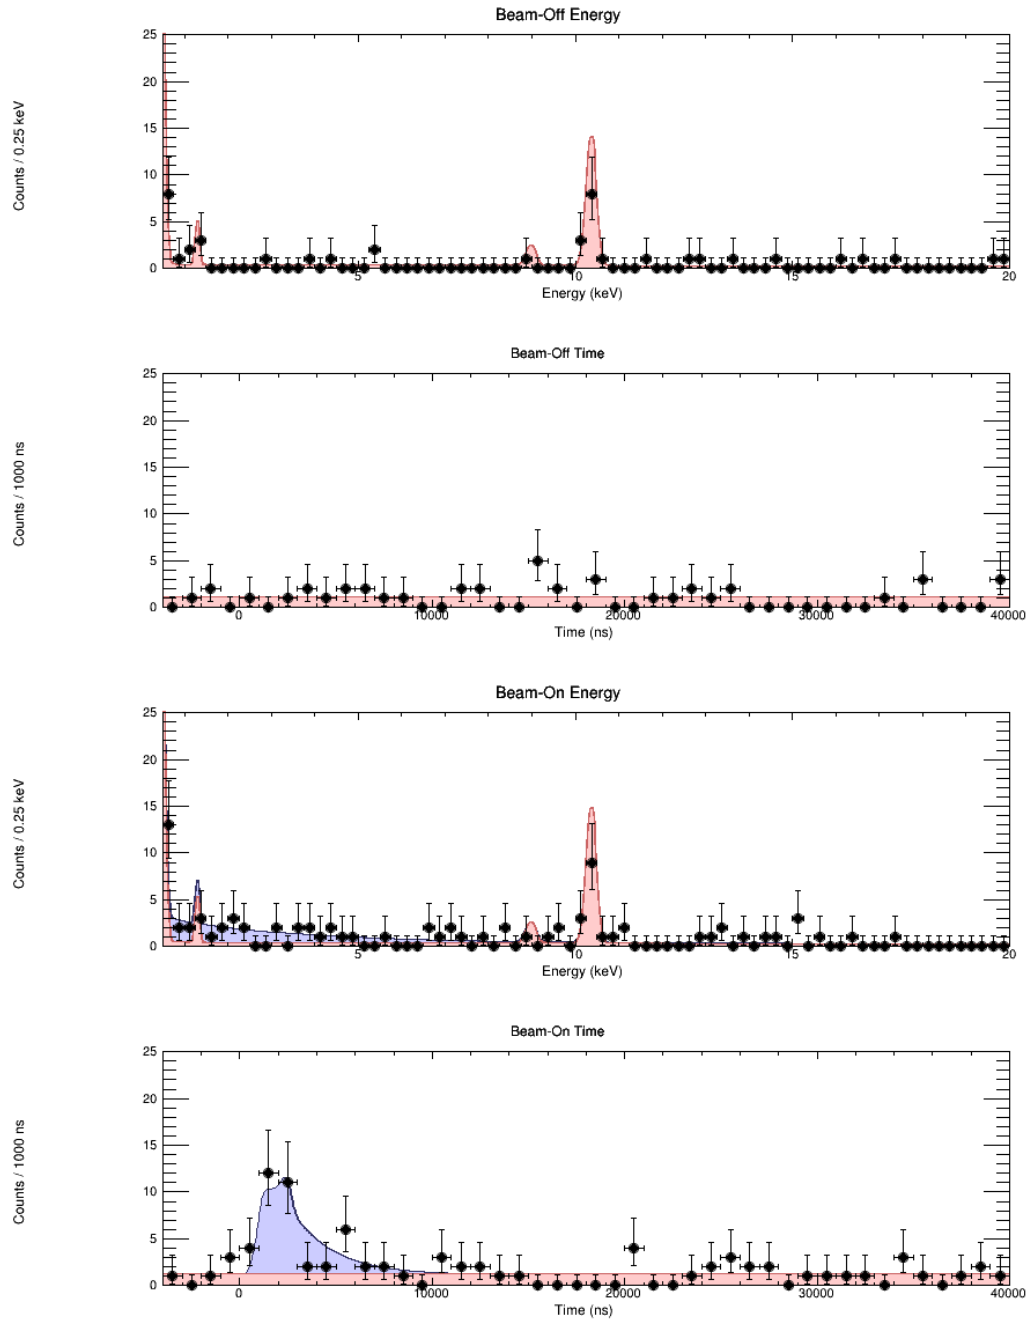
Every detector on its own shows a clear presence of signal in coincidence with the beam. Profiled log likelihoods for the number of signal counts in each detector were performed. The ratio in comparison to the Standard Model is evaluated as well.

## 8.7 Discussion

The statistical precision of the each detector's result is evaluated through profiled log-likelihoods of each observed number of counts. This is shown in Figure 8.30.

A simple way to combine the sensitivity of all the detectors is to take the ratio of the observed counts to the expected number of counts and add the profiled likelihoods in this dimension. This gives a total sensitivity of the experiment and a comparison to the Standard Model as a percentage. This is shown in Figure 8.31.

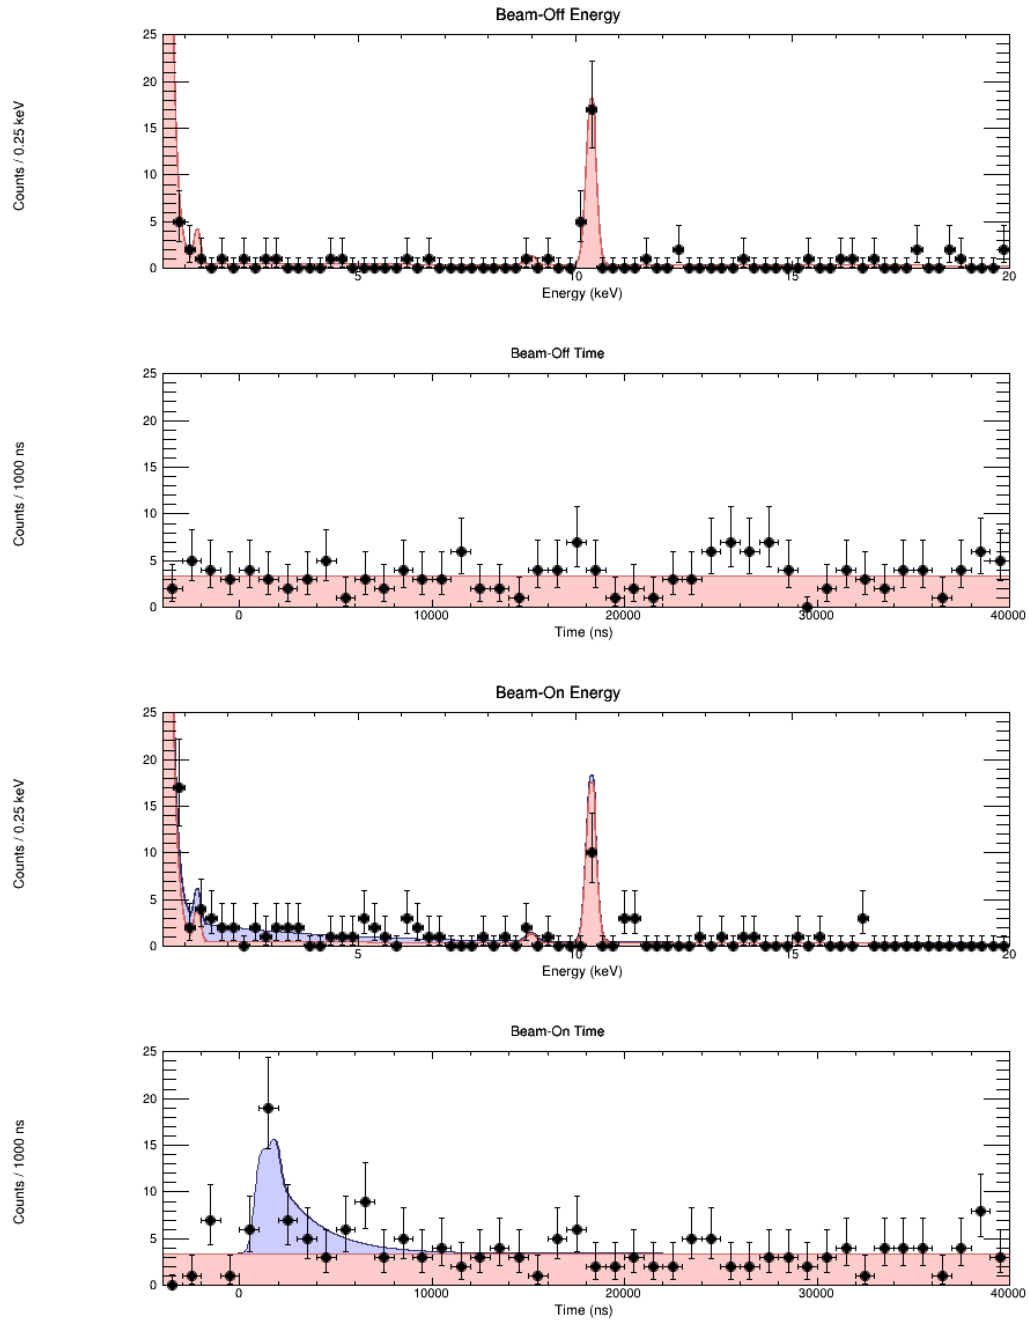
Finally, it may be of particular interest to see the spectra combined for visual effect, despite their individual underlying fits. This is done by summing the individually fit models



**Figure 8.26:** Ge21 unblinded data and best fits to signal and background are shown.

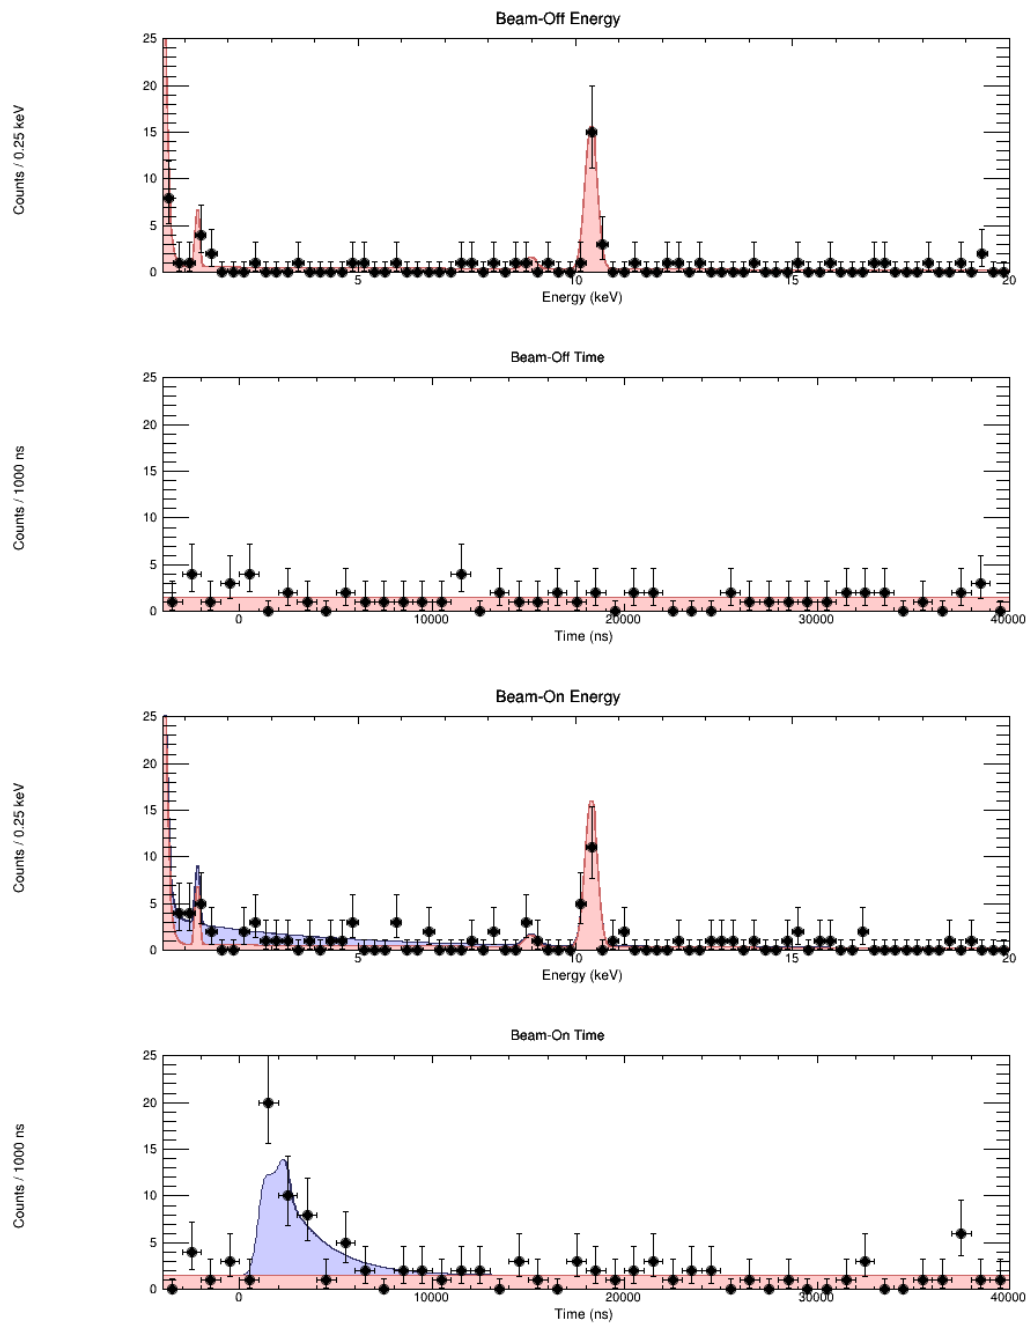
and shown agains the combined observed counts in Figure 8.32.

All signal counts are tabulated in Table 8.7, including counts from a naive counting experiment, the best fit, and the Standard Model expectation.

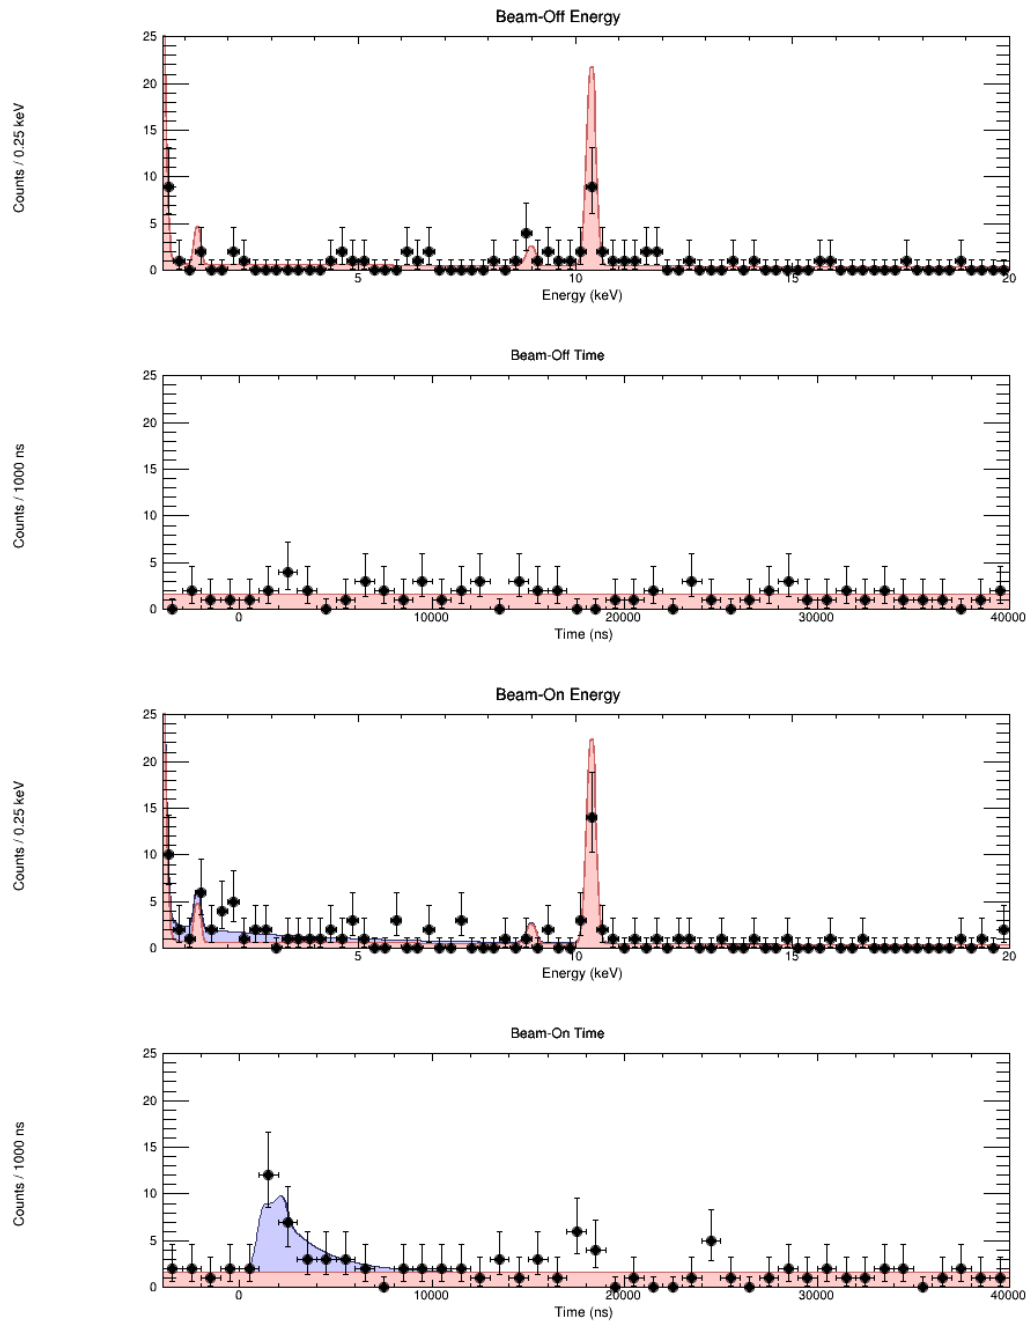


**Figure 8.27:** Ge25 unblinded data and best fits to signal and background are shown.


Campaign 3 announces the most precise measurement of CEvNS to-date. The measurement is in excellent agreement with the Standard Model.




**Figure 8.28:** Ge26 unblinded data and best fits to signal and background are shown.



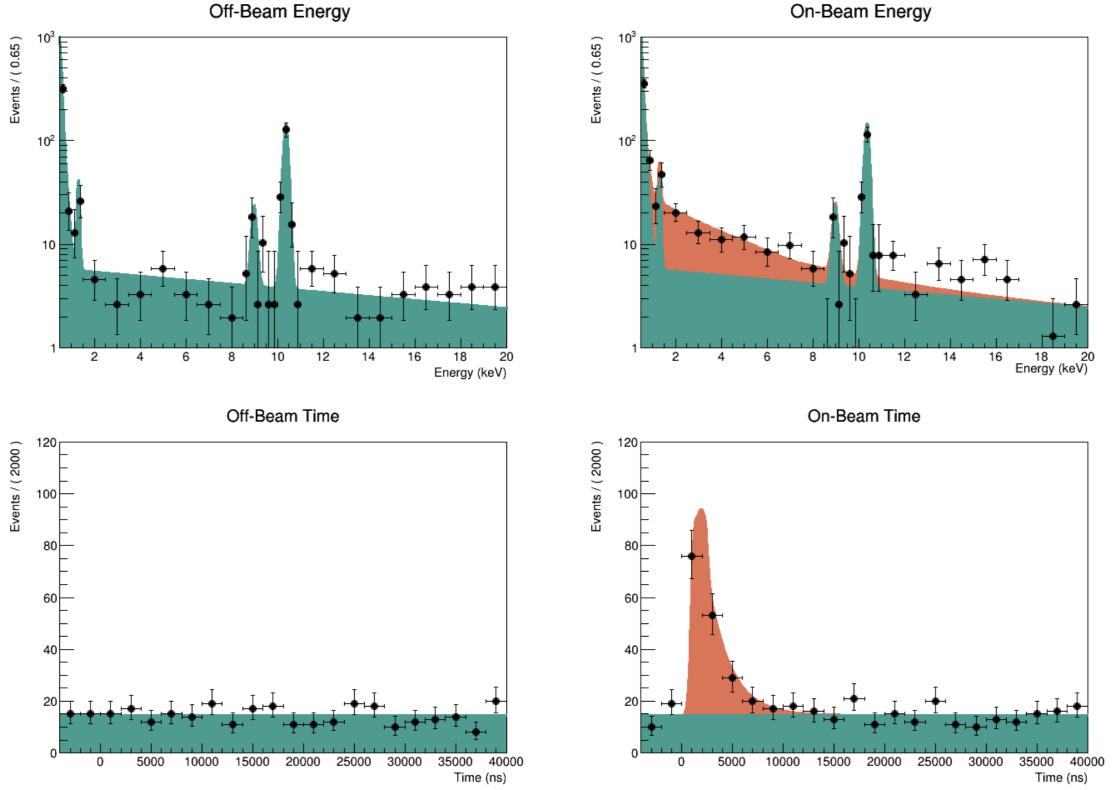
**Figure 8.29:** Ge28 unblinded data and best fits to signal and background are shown.



**Figure 8.30:** Profiled log-likelihoods for the number of counts in each detector. 68% confidence intervals are quoted under the assumption of an asymptotically  $\chi^2$  distributed likelihood ratio test statistic.



**Figure 8.31:** The PLL of each detector is expressed as a ratio of the observed number of counts to the predicted number of counts. The shaded region around the Standard Model prediction encompasses the 10% systematic uncertainty.



**Figure 8.32:** Combined unblinded datasets across all four detectors in Campaign 3. The best fits of each detector are added together to produce a combined background model and combined signal model.

**Table 8.7:** Unblinding results: the counting analysis was carried out for the ROI optimized for the best S/B, while the likelihood covers the full range. Uncertainties for the combined fit results are estimated by combining the single results.

Detector	counting [0.5, 8.5]keV <sub>ee</sub> [0, 8] $\mu$ s	likelihood [0.5, 20]keV <sub>ee</sub> [-4, 40] $\mu$ s	SM CEvNS [0.5, 20]keV <sub>ee</sub> full range
Ge21	27.0 $\pm$ 5.9 (4.6 $\sigma$ )	31.95+6.54-5.83 (5.5 $\sigma$ )	31.0 $\pm$ 3.2
Ge25	32.0 $\pm$ 8.0 (4.0 $\sigma$ )	31.88+6.81-6.06 (5.3 $\sigma$ )	30.8 $\pm$ 3.2
Ge26	31.0 $\pm$ 6.7 (4.7 $\sigma$ )	36.24+7.00-6.28 (5.8 $\sigma$ )	31.1 $\pm$ 3.2
Ge28	15.0 $\pm$ 5.7 (2.6 $\sigma$ )	23.78+5.62-4.92 (4.8 $\sigma$ )	31.3 $\pm$ 3.2
all	105.0 $\pm$ 13.2 (7.9 $\sigma$ )	123.87+13.1-11.5 ( $\sim$ 10 $\sigma$ )	124.3 $\pm$ 12.8

# Chapter 9

## Non-Standard Interactions


This thesis has, so-far, presented a dual picture within the framework of EFT. On the one hand, theoretical developments in the prediction of low-energy weak interactions were shown and this was connected to tests of new physics beyond the Standard Model. On the other hand, an experimental result was presented, providing the world's most precise measurement of CEvNS to date. These two contributions were concerned primarily with the subject of precision and how both sides of this study are necessary to enable discovery.

As an example of how these pieces can come together in practice, it is useful to consider the example of non-standard interactions (NSIs) in weak neutral currents and what this experimental result can tell us about new physics. The study of NSIs fits nicely in the overall EFT framework of this thesis as they typify a key theme that has been emphasized: the parametrization of both Standard Model physics and beyond Standard Model physics in one language that allows low-energy experiments to probe structure at significantly higher energies.

The example of NSIs as a use-case for a CEvNS measurement is particularly illustrative as it captures a broad set of potential new physics into a parametrized deviation from the standard expectation. Beyond this example there are many applications of the data presented in this work, ranging from electromagnetic properties of the neutrino [78] to studies of the properties of the germanium nucleus [79]. This chapter will focus on what this data can say about NSIs, building on existing tests with COHERENT data [80].

### 9.1 Formalism

The study of NSIs is a broadband search for new physics. In particular, it is a low-energy parametrization of new physics beyond the Standard Model that does not necessarily depend



**Figure 9.1:** Two-dimensional, profiled allowed regions for  $(\varepsilon_d^{\mu\mu}, \varepsilon_u^{\mu\mu})$  plane, with all other parameters fixed to zero. 1, 2, and 3  $\sigma$  allowed contours are shown. Double bands are characteristic of the degeneracy in parameters for NSIs.

on the specification of the nature of that physics. It was first considered in the literature in the context of neutrino oscillation phenomena [81], and since then has been studied extensively. Here, a primarily phenomenological perspective is taken to make use of the tools of parametrization of a particular type of deviation from the Standard Model without much emphasis placed on the source of that deviation. Indeed, if such deviations solidified themselves with the coming programs of increasingly precise CEvNS experiments more dedicated hypothesis testing would then become a more pressing endeavor.

The effective four-fermion Lagrangian that parametrizes the NSIs modification of the dominant vector component of the CEvNS differential rate is given by [82]:

$$\mathcal{L}_{\text{NSI}}^{\text{NC}} = -2\sqrt{2}G_F \sum_{\alpha,\beta=e,\mu,\tau} (\bar{\nu}_{\alpha L}\gamma^\rho\nu_{\beta L}) \sum_{f=u,d} \varepsilon_{\alpha\beta}^{fV} (\bar{f}\gamma_\rho f), \quad (9.1)$$

where  $G_F$  is the Fermi constant and  $\varepsilon_{\alpha\beta}^{fV}$  parametrizes the strength of the non-standard

deviation of the vector coupling. These coupling parameters depend on the neutrino flavor as well as on the type of quark.

The usual weak charge  $Q_W^2$  is modified by the presence of the non-standard couplings and now depends on the particular neutrino flavor, which is labeled here as  $Q_\alpha^2$ . This can be expressed as

$$Q_\alpha^2 = \left[ \left( g_V^p + 2\varepsilon_{\alpha\alpha}^{uV} + \varepsilon_{\alpha\alpha}^{dV} \right) ZF_Z(|\vec{q}|^2) + \left( g_V^n + \varepsilon_{\alpha\alpha}^{uV} + 2\varepsilon_{\alpha\alpha}^{dV} \right) NF_N(|\vec{q}|^2) \right]^2 + \sum_{\beta \neq \alpha} \left| \left( 2\varepsilon_{\alpha\beta}^{uV} + \varepsilon_{\alpha\beta}^{dV} \right) ZF_Z(|\vec{q}|^2) + \left( \varepsilon_{\alpha\beta}^{uV} + 2\varepsilon_{\alpha\beta}^{dV} \right) NF_N(|\vec{q}|^2) \right|^2. \quad (9.2)$$

Here,  $Z$  and  $N$  denote the number of protons and neutrons in the nucleus. Each nucleon has an associated weak charge described by the valence quark content. There are several ways that the nucleon charges can be modified through diagonal and off-diagonal couplings.


Evident from Equation 9.1, the total number of parameters that can contribute to NSIs is large, and even for limited subsets of parameters degeneracies exist. For simplicity, here, only a few sets of diagonal parameters are considered.

## 9.2 Statistical Analysis

The statistical treatment followed in this analysis is largely similar to the statistical analysis presented in the CEvNS searches. The non-standard parameters  $\varepsilon$  were constrained by performing through a hybrid maximum-likelihood estimate with gaussian pull-terms to capture the systematic uncertainties on the predicted number of events. The negative log-likelihood (NLL) minimized in the fit is given by For a data set containing  $N_{\text{obs}}$  events with recoil energies  $\{E_i\}$ , the extended likelihood function is

$$\mathcal{L}(\boldsymbol{\varepsilon}, \boldsymbol{\theta}) = e^{-N^{\text{pred}}(\boldsymbol{\varepsilon}, \boldsymbol{\theta})} \prod_{i=1}^{N_{\text{obs}}} \lambda(E_i, t_i | \boldsymbol{\varepsilon}, \boldsymbol{\theta}) \times \prod_j \exp \left[ -\frac{1}{2} \left( \frac{\eta_j}{\sigma_j} \right)^2 \right], \quad (9.3)$$

where  $\lambda(E_i, t_i | \boldsymbol{\varepsilon}, \boldsymbol{\theta})$  is the differential event rate evaluated at the measured energy and time  $(E_i, t_i)$ . Here,  $\boldsymbol{\varepsilon}$  denotes the NSI parameters of interest and  $\boldsymbol{\theta}$  represents nuisance parameters



**Figure 9.2:** Two-dimensional, profiled allowed regions for  $(\epsilon_u^{ee}, \epsilon_u^{\mu\mu})$  plane, with all other parameters fixed to zero.

describing detector response and normalization systematics. In particular the set of model parameters  $\epsilon$  modifies the expected signal PDFs by modifying the ratios of neutrino flavors present in the total recoil spectrum but does not distort the shape of the individual flavor recoil spectra.

Systematic uncertainties in the prediction are captured in the parameter estimation through a Gaussian constraint on a global normalization parameter. This is dominated by neutrino flux, though includes smaller effects such as those from the active mass uncertainty or nuclear form factor. This is written as  $\eta_{\text{norm}}$  given by

$$N^{\text{pred}}(\epsilon, \theta) = N_0(\epsilon) (1 + \eta_{\text{norm}}), \quad \eta_{\text{norm}} \sim \mathcal{N}(0, 0.103^2), \quad (9.4)$$

where  $N_0(\epsilon)$  is the central-value prediction and the 10.3% uncertainty reflects the total systematic uncertainty on the prediction.

In practice, numerically, the negative log-likelihood is minimized with respect to the model parameters and nuisance parameters.

### 9.3 Results

Performing a subset of searches for parameters which contribute to NSIs reveals agreement with the Standard Model predictions at the  $2\sigma$  level. Two-dimensional planes of parameters were considered where constraints are the strongest:  $(\varepsilon_d^{\mu\mu}, \varepsilon_u^{\mu\mu})$  and  $(\varepsilon_u^{ee}, \varepsilon_u^{\mu\mu})$ . The constraints presented here strengthen the existing constraints on these parameters [82], and give slight hints of a non-standard preference.

The future programs precision programs at the SNS will be able to probe further the appearance of any subtle deviations from the Standard Model.

# Chapter 10

## Conclusion

Throughout this work the framework of low-energy, effective theories has been used to communicate predictions and measurements of weak interaction. The success of this framework has been to generate tools for simple and precise predictions, which parametrize the known and unknown physics relevant in the laboratory. The search for this unknown physics depends on dedicated campaigns of measurements capable of pushing these predictions to their limits. To date, in low-energy neutrino interaction, the precision of prediction are still far-ahead of the companion experimental precision. This thesis highlights contributions that reflects a dual effort to close the gap between the precision of theory and experiment in study of low energy weak interaction.

### **Theoretical Contributions**

The theoretical contributions described in this work focus on the rare, doubly-weak process of  $2\nu\beta\beta$  decay. The growing international platform for the search of  $0\nu\beta\beta$  has generated increasingly precise measurements of the Standard Model process. As the precision of these measurements improve it becomes increasingly necessary to maintain precise predictions so has to parametrize possible new physics, and at the same time to not confuse it for higher order corrections typically excluded from existing predictions.

The contributions to this effort, described in Chapter 3, summarize two specific examples that illustrate the interplay between higher order corrections and potential new physics. They are:

1. NLO corrections in  $2\nu\beta\beta$  decay arising from induced currents from the non-relativistic expansion of nuclear currents.
2. Exotic currents at the quark level from a hypothetical tensor Lorentz structure, and how this new interaction can mimic Standard Model physics.

As next-generation double beta decay experiments such as LEGEND continue to collect more data it will be important to use careful predictive frameworks to use the  $2\nu\beta\beta$  differential rates as a probe of the Standard Model.

### **Experimental Contributions**

The experimental contributions in this work are centered around the HPGe detector-system within the COHERENT collaboration. Two separate measurement campaigns are reported, which include both the first measurement by the collaboration on a germanium target, as well as the world's most precise CEvNS measurement to-date. This measurement has a broad range of applications in the search for physics beyond the Standard Model. As the most precise measurement of CEvNS to-date it is of interest to model builders in their attempts to parametrize physics that may distort the distribution or rates of measured events.

The COHERENT germanium detector system, Ge-Mini, shows great promise in continuing to make increasingly precise measurements of the CEvNS differential rate. It will continue taking data increasing its statistical precision, following the analysis routines outlined in this work.

However, the statistical uncertainties associated with the measurement reported in this thesis have already reached the point where they are smaller than the systematic uncertainties on the prediction. This relatively large systematic uncertainty is due to the 10% uncertainty in the neutrino flux production at the SNS. An important next step for the collaboration will be the first results from the heavy-water detector system, which will serve to normalize the neutrino flux at the SNS. The combination of these detector systems will allow for future analyses to classify deviations from predictions as being related to new physics or to underlying systematic uncertainties. As other neutrino detector systems come online at the SNS, combined analyses can further increase the sensitivity of CEvNS as a probe for new physics.

# Bibliography

- [1] D. Aristizabal Sierra, *Neutrinos: Theory Lectures*, 2023.
- [2] C. S. Wu et al., *Experimental Test of Parity Conservation in Beta Decay*, **105**, 1413–1415 (1957).
- [3] R. P. Feynman and M. Gell-Mann, *Theory of the Fermi Interaction*, **109**, 193–198 (1958).
- [4] A. Sirlin, *Radiative Corrections in the  $\mathit{SU}(2)_L$   $\mathit{ifm-mode} \times \mathit{else} \mathit{timestimes} \mathit{fi}$   $\mathit{U}(1)$  Theory: A Simple Renormalization Framework*, **22**, 971–981 (1980).
- [5] M. Hoferichter, J. Menéndez, and A. Schwenk, *Coherent Elastic Neutrino-Nucleus Scattering: EFT Analysis and Nuclear Responses*, **102**, 074018 (2020).
- [6] D. Djukanovic et al., *Electromagnetic Form Factors of the Nucleon from  $\mathit{N}_{f=2+1}$  Lattice QCD*, **109**, 094510 (2024).
- [7] G. F. Chew and G. C. Wick, *The Impulse Approximation*, **85**, 636–642 (1952).
- [8] T. Tomoda, *Double Beta Decay*, **54**, 53 (1991).
- [9] M. E. Rose and R. K. Osborn, *Nuclear Matrix Elements in Beta Decay*, **93**, 1326–1336 (1954).
- [10] N. V. Dessel, V. Pandey, H. Ray, and N. Jachowicz, *Cross Sections for Coherent Elastic and Inelastic Neutrino-Nucleus Scattering*, **9**, 207 (2023).
- [11] J. Menéndez, D. Gazit, and A. Schwenk, *Chiral Two-Body Currents in Nuclei: Gamow-Teller Transitions and Neutrinoless Double-Beta Decay*, **107**, 062501 (2011).
- [12] W. Pauli, *Pauli Letter Collection*.
- [13] V. Cirigliano, S. Gardner, and B. R. Holstein, *Beta Decays and Non-Standard Interactions in the LHC Era*, *Fundamental Symmetries in the Era of the LHC* **71**, 93–118 (2013).
- [14] M. Goeppert-Mayer, *Double Beta-Disintegration*, **48**, 512–516 (1935).
- [15] S. R. Elliott, A. A. Hahn, and M. K. Moe, *Direct Evidence for Two-Neutrino Double-Beta Decay in  $^{82}\mathit{Se}$* , **59**, 2020–2023 (1987).

- [16] B. Longfellow et al., *Improved Tensor Current Limit from  $^8\mathrm{B}$   $\beta$  Decay Including New Recoil-Order Calculations*, **132**, 142502 (2024).
- [17] A. Branca, *Recent Progresses on BSM and Dark Matter Searches with CUORE*, 054 (2023).
- [18] GERDA Collaboration et al., *Final Results of GERDA on the Two-Neutrino Double- $\beta$  Decay Half-Life of  $^{76}\mathrm{Ge}$* , **131**, 142501 (2023).
- [19] M. Doi, T. Kotani, and E. Takasugi, *Double Beta Decay and Majorana Neutrino*, **83**, 1–175 (1985).
- [20] W. C. Haxton and G. J. Stephenson, *Double Beta Decay*, **12**, 409–479 (1984).
- [21] F. Šimkovic, R. Dvornický, D. Štefánik, and A. Faessler, *Improved Description of the  $^2\nu\beta\beta$ -Decay and a Possibility to Determine the Effective Axial-Vector Coupling Constant*, **97**, 034315 (2018).
- [22] S. el Morabit et al.,  *$2\nu\beta\beta$  Spectrum in Chiral Effective Field Theory*, **2025**, 82 (2025).
- [23] T. D. Lee and C. N. Yang, *Question of Parity Conservation in Weak Interactions*, **104**, 254–258 (1956).
- [24] I. Brivio and M. Trott, *The Standard Model as an Effective Field Theory*, **793**, 1–98 (2019).
- [25] A. Falkowski, M. González-Alonso, and O. Naviliat-Cuncic, *Comprehensive Analysis of Beta Decays within and beyond the Standard Model*, **2021**, 126 (2021).
- [26] Y. Aoki et al., *FLAG Review 2021*, **82**, 1–296 (2022).
- [27] P. Galison, *How Experiments End* (University of Chicago Press, 1987).
- [28] D. Z. Freedman, *Coherent Effects of a Weak Neutral Current*, **9**, 1389–1392 (1974).
- [29] A. Drukier and L. Stodolsky, *Principles and Applications of a Neutral-Current Detector for Neutrino Physics and Astronomy*, **30**, 2295–2309 (1984).
- [30] B. Cabrera, L. M. Krauss, and F. Wilczek, *Bolometric Detection of Neutrinos*, **55**, 25–28 (1985).
- [31] O. Tomalak, P. Machado, V. Pandey, and R. Plestid, *Flavor-Dependent Radiative Corrections in Coherent Elastic Neutrino-Nucleus Scattering*, **2021**, 97 (2021).
- [32] *[1901.04190] CYGNO: A CYGNUs Collaboration 1  $M^3$  Module with Optical Readout for Directional Dark Matter Search*, <https://arxiv.org/abs/1901.04190>.

- [33] R. H. Helm, *Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei*, **104**, 1466–1475 (1956).
- [34] S. R. Klein and J. Nystrand, *Exclusive Vector Meson Production in Relativistic Heavy Ion Collisions*, **60**, 014903 (1999).
- [35] G. Co', M. Anguiano, and A. M. Lallena, *Tensor Force and Deformation in Even-Even Nuclei*, **104**, 014313 (2021).
- [36] T. Ohlsson, *Status of Non-Standard Neutrino Interactions*, **76**, 044201 (2013).
- [37] C. Giunti, *General COHERENT Constraints on Neutrino Nonstandard Interactions*, **101**, 035039 (2020).
- [38] M. A. Corona et al., *Momentum Dependent Flavor Radiative Corrections to the Coherent Elastic Neutrino-Nucleus Scattering for the Neutrino Charge-Radius Determination*, **2024**, 271 (2024).
- [39] J. F. Beacom, W. M. Farr, and P. Vogel, *Detection of Supernova Neutrinos by Neutrino-Proton Elastic Scattering*, **66**, 033001 (2002).
- [40] J. R. Vanhoy et al.,  *$^{54}\text{Fe}$  Neutron Elastic and Inelastic Scattering Differential Cross Sections from 2–6 MeV*, **972**, 107–120 (2018).
- [41] D. Akimov et al., *Observation of Coherent Elastic Neutrino-Nucleus Scattering*, **357**, 1123–1126 (2017).
- [42] COHERENT Collaboration et al., *Measurement of the Coherent Elastic Neutrino-Nucleus Scattering Cross Section on CsI by COHERENT*, **129**, 081801 (2022).
- [43] COHERENT Collaboration et al., *First Measurement of Coherent Elastic Neutrino-Nucleus Scattering on Argon*, **126**, 012002 (2021).
- [44] *Results from a Prototype TES Detector for the Ricochet Experiment - ScienceDirect*, <https://www.sciencedirect.com/science/article/pii/S0168900223007568?via%3Dihub>.
- [45] T. N. collaboration et al., *Exploring Coherent Elastic Neutrino-Nucleus Scattering of Reactor Neutrinos with the NUCLEUS Experiment*, 2022.
- [46] The COHERENT Collaboration et al., *Evidence of Coherent Elastic Neutrino-Nucleus Scattering with COHERENT's Germanium Array*, **134**, 231801 (2025).
- [47] N. Ackermann et al., *Direct Observation of Coherent Elastic Antineutrino–Nucleus Scattering*, **643**, 1229–1233 (2025).

- [48] PICO Collaboration et al., *Absorption of Fermionic Dark Matter in the PICO-60  $\{C\}_3\{F\}_8$  Bubble Chamber*, **135**, 011001 (2025).
- [49] P. Dhakal et al., *Performance of Oak Ridge National Laboratory Spallation Neutron Source Proton Power Upgrade Cavities and Cryomodule Production*, **27**, 102001 (2024).
- [50] COHERENT Collaboration et al., *Simulating the Neutrino Flux from the Spallation Neutron Source for the COHERENT Experiment*, **106**, 032003 (2022).
- [51] R. Rapp, *Characterizing Neutrino and Neutron Fluxes from the Spallation Neutron Source for the COHERENT Experiment*, Thesis, 2022.
- [52] .collaborationEtal., *A D2O Detector for Flux Normalization of a Pion Decay-at-Rest Neutrino Source*, **16**, P08048 (2021).
- [53] P. Luke, F. Goulding, N. Madden, and R. Pehl, *Low Capacitance Large Volume Shaped-Field Germanium Detector*, **36**, 926–930 (1989).
- [54] *Large-Mass Ultralow Noise Germanium Detectors: Performance and Applications in Neutrino and Astroparticle Physics - IOPscience*, <https://iopscience.iop.org/article/10.1088/1475-7516/2007/09/009>.
- [55] M. Agostini et al., *Characterization of Inverted Coaxial  $^{76}\text{Ge}$  Detectors in GERDA for Future Double- $\beta$  Decay Experiments*, **81**, 505 (2021).
- [56] *WELL™ Germanium Well Detector*, <https://www.mirion.com/products/technologies/spectroscopy-scientific-analysis/gamma-spectroscopy/detectors/hpge-detectors-accessories/well-germanium-well-detector>.
- [57] D. Cintas et al., *Measurement of the Sodium and Iodine Scintillation Quenching Factors across Multiple NaI(Tl) Detectors to Identify Systematics*, **110**, 014613 (2024).
- [58] A. Bonhomme et al., *Direct Measurement of the Ionization Quenching Factor of Nuclear Recoils in Germanium in the keV Energy Range*, **82**, 815 (2022).
- [59] E. Gatti and P. F. Manfredi, *Processing the Signals from Solid-State Detectors in Elementary-Particle Physics*, **9**, 1–146 (1986).
- [60] V. Radeka, *Low-Noise Techniques in Detectors*, **38**, 217–277 (1988).
- [61] *Mirion, Radiation Measurement & Detection Devices*, <https://www.mirion.com/>.
- [62] *MPOD High Voltage Module | W-IE-NE-R Power Electronics*.
- [63] *Struck Innovative Systeme Digitizer Overview Page*, <https://www.struck.de/digov-erview.html>.

- [64] *16-Bit Digitizer Family*, <https://www.struck.de/sis3316-2.html>.
- [65] K. Kodama et al., *Observation of Tau Neutrino Interactions*, **504**, 218–224 (2001).
- [66] R. Davis, *A Review of the Homestake Solar Neutrino Experiment*, **32**, 13–32 (1994).
- [67] D. Radford, *Radforddc/Icpc\_siggen*, 2025.
- [68] Majorana Collaboration et al., *Charge Trapping Correction and Energy Performance of the Majorana Demonstrator*, **107**, 045503 (2023).
- [69] *MCNP® Website*, <https://mcnp.lanl.gov/>.
- [70] *Home - Eljen Technology*, <https://eljentechnology.com/>.
- [71] J. Lindhard, M. Scharff, and H. E. Schiøtt, *RANGE CONCEPTS AND HEAVY ION RANGES (NOTES ON ATOMIC COLLISIONS, II)*, **Vol: 33: No. 14** (1962).
- [72] L. Li, *A Measurement of The Response of A High Purity Germanium Detector to Low-Energy Nuclear Recoils*, 2022.
- [73] S. Navas et al., *Review of Particle Physics*, **110**, 030001 (2024).
- [74] H. Dembinski et al., *Scikit-Hep/Iminuit*, Zenodo, 2025.
- [75] G. Turin, *An Introduction to Matched Filters*, **6**, 10.1109/TIT.1960.1057571 (1960).
- [76] H. Bonet et al., *Pulse Shape Discrimination for the CONUS Experiment in the keV and Sub-keV Regime*, **84**, 139 (2024).
- [77] H. Bonet et al., *Full Background Decomposition of the CONUS Experiment*, **83**, 195 (2023).
- [78] C. Giunti, K. Kouzakov, Y.-F. Li, and A. Studenikin, *Neutrino Electromagnetic Properties*, **75**, 1–33 (2025).
- [79] M. Atzori Corona et al., *Joint Analysis of Reactor and Accelerator  $CE\nu$   $\nu$   $NS$  Data on Germanium: Implications for the Standard Model and Nuclear Physics*, **869**, 139856 (2025).
- [80] J. Liao and D. Marfatia, *COHERENT Constraints on Nonstandard Neutrino Interactions*, **775**, 54–57 (2017).
- [81] L. Wolfenstein, *Neutrino Oscillations in Matter*, **17**, 2369–2374 (1978).

- [82] V. M. Lozano, G. S. Garcia, and A. Terrones, *Neutrino Nonstandard Interactions: Confronting COHERENT and LHC Data*, **112**, 055017 (2025).
- [83] C. R. Harris et al., *Array Programming with NumPy*, *Nature* **585**, 357–362 (2020).
- [84] P. Virtanen et al., *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nat. Methods* **17**, 261–272 (2020).
- [85] J. D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9**, 90–95 (2007).