

A SEARCH FOR MUON NEUTRINO TO
ELECTRON NEUTRINO OSCILLATIONS
AT $\Delta m^2 > 0.1 \text{ eV}^2$

RYAN BENTON PATTERSON

A DISSERTATION PRESENTED TO THE FACULTY OF
PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF PHYSICS

NOVEMBER 2007

© Copyright by Ryan Benton Patterson, 2007. All Rights Reserved.

Abstract

The evidence is compelling that neutrinos undergo flavor change as they propagate. In recent years, experiments have observed this phenomenon of neutrino oscillations using disparate neutrino sources: the sun, fission reactors, accelerators, and secondary cosmic rays. The standard model of particle physics needs only simple extensions – neutrino masses and mixing – to accommodate all neutrino oscillation results to date, save one. The 3.8σ -significant $\bar{\nu}_e$ excess reported by the LSND collaboration is consistent with $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations with a mass-squared splitting of $\Delta m^2 \sim 1 \text{ eV}^2$. This signal, which has not been independently verified, is inconsistent with other oscillation evidence unless more daring standard model extensions (*e.g.*, sterile neutrinos) are considered.

The (Mini) Booster Neutrino Experiment (MiniBooNE) at the Fermi National Accelerator Laboratory is designed to search for $\nu_\mu \rightarrow \nu_e$ oscillations with sufficient sensitivity to confirm or refute the LSND signal. This dissertation presents the first $\nu_\mu \rightarrow \nu_e$ oscillation search performed at MiniBooNE. After a short introduction, we describe the experimental setup, which includes a $\sim 1 \text{ GeV}$ ν_μ source fed by an 8 GeV proton synchrotron; and an 800 ton mineral oil Cherenkov detector located 0.5 km downstream. We then present a detailed discussion of the analysis, which uses the full neutrino data set corresponding to 5.58×10^{20} protons-on-target and 1.7×10^6 neutrino interactions. No evidence for LSND-like $\nu_\mu \rightarrow \nu_e$ oscillations is seen either in the count of ν_e candidates [data: 380, expectation: $358 \pm 19(\text{stat.}) \pm 35(\text{syst.})$] or in the shape of the neutrino energy spectrum.

Acknowledgments

Foremost, I am truly grateful to my advisor, Peter Meyers, for having provided seven years of unerring guidance. It has been a privilege to work closely with him on all matters found in these pages, and I have doubtless learned more from Peter about mentoring, collaborating, and conducting research than I will ever actually realize.

As is typical for a particle physics experiment, the total human effort expended in support of this neutrino oscillation result is best measured in person-centuries. I cannot possibly acknowledge the innumerable individual contributions, so I say to all my MiniBooNE colleagues: Thank you for everything. It has been a pleasure.

In particular, I would like to thank Hirohisa Tanaka, who has been my go-to guy for most of my time at Princeton. Only through countless hours of discussion with Hiro has this thesis come to pass, and considerable credit for it goes to him. I also thank Ted Laird for his significant contributions to the analysis and for the many stimulating conversations we've had on all topics imaginable. I have thoroughly enjoyed working with Hiro and Ted, and I thank them for their valued friendship.

Despite his approximate job description of “all things to all people,” Steve Brice has always found time to engage in useful (to me) discussions, and for this and more I thank him. In addition, I have had many fruitful exchanges with Andrew Bazarko, Chris Green, Chris Polly, Mike Shaevitz, Ion Stancu, and Sam Zeller, and I thank them for their time and expertise. Also, if you anticipate spending four months working inside a 40-foot steel tank, you'll do well to bring Stan Chidzik, Hayes Lansford, and Bill Sands along.

Outside of research, I have had the good fortune of attending graduate school with two of my best friends from college, quite by chance. I've known Jimmy Cooley and Brian Patton for over a decade now, and I appreciate the wisdom, humor, skepticism, and all around camaraderie they bring. Ben Brantley, Tony Guerrera, and Craig Vieregge have always had an adventure of some sort waiting when I've had a little time off. And, living in central New Jersey was much more pleasant than one might expect thanks to Asad Aboobaker, Joe Haley, Kevin Hufferberger, Lewis Hyatt, Kelly Lee, Mike Leung, Glen Nixon, and Greg Riviello.

Lastly, I give my heartfelt thanks to my parents, Bob and Helen. They have offered their steadfast support of all my endeavors over the years and have never asked for a thing in return – except perhaps that I call home once in a while.

Contents

Abstract	iii
Acknowledgments	iv
List of Tables	x
List of Figures	xi
I Introduction	1
1 Neutrino oscillations and MiniBooNE	2
1.1 Extending the standard model	2
1.1.1 Neutrino masses	3
1.1.2 Neutrino mixing	4
1.2 Implications	4
1.3 Neutrino oscillations	7
1.3.1 General formalism	7
1.3.2 Quasi-two-neutrino oscillations	9
1.3.3 Matter effects	9
1.3.4 Experimental evidence	10
1.3.5 LSND	14
1.4 MiniBooNE	20
1.5 Layout of this document	21
2 Booster neutrino beam	22
2.1 Proton source	22
2.2 Target and horn	22
2.3 Decay region and flux	26
3 MiniBooNE detector	28
3.1 Overview	28
3.2 PMT support structure	28
3.3 Oil	29
3.4 PMTs	30
3.5 Electronics and DAQ	31
3.6 Calibration system	35

3.7	Trigger and readout	36
4	PMT charge and time	39
4.1	Determining t	39
4.2	Determining q	41
4.2.1	The shape of V_q	41
4.2.2	Calculating Q_{raw}	43
4.2.3	Establishing V_q^{ref}	44
4.2.4	Completing the picture	45
4.2.5	Verification	46
4.3	Obtaining calibration constants	46
4.3.1	The slope m	48
4.3.2	The timing offset Δt_{offset}	48
4.3.3	The extended V_q^{ref}	49
4.3.4	The charge scale factor g	50
4.3.5	The slew correction	50
II	Analysis	54
5	Analysis overview	55
5.1	The task	55
5.2	Subevents	55
5.3	Neutrino events	58
5.4	Blind analysis	60
5.5	Simulation chain	62
5.6	Systematic error propagation	63
5.7	Likelihood ratios and $M_{\gamma\gamma}$	63
6	Neutrino flux prediction	64
6.1	GEANT4 beamline description	64
6.1.1	Proton beam	64
6.1.2	Hadron production	65
6.1.3	Particle propagation	70
6.1.4	Production weighting and decays	71
6.2	Neutrino history	72
6.3	Resulting flux	72
6.4	Systematic uncertainties	72
7	Neutrino cross section model	75
7.1	Charged current quasi-elastic scattering	75
7.2	Resonant processes	79
7.3	Neutral current π^0 production	79

7.4	Final state interactions	80
7.5	Miscellany	81
7.6	Systematic uncertainties	82
8	Detector simulation	84
8.1	Geometry definition	84
8.2	Particle propagation	85
8.3	Optical photons	86
8.3.1	Production	86
8.3.2	Propagation	87
8.3.3	Detection	92
8.4	PMT response	93
8.5	Digitization	94
8.6	Beam-off activity	97
8.7	Tuning and systematic uncertainties	98
9	Event reconstruction	104
9.1	Introduction	104
9.2	The charge likelihood	105
9.2.1	Calculating predicted charge μ	106
9.2.2	Computational realities	114
9.2.3	The functions $P(\text{hit}; \mu)$ and $f_q(q; \mu)$	117
9.3	The time likelihood	117
9.3.1	Reducing the problem	119
9.3.2	Creating the primitive distributions	119
9.3.3	Using the primitive distributions	121
9.4	μ and e hypotheses	125
9.5	Two-track fits	125
9.6	Minimizing $F = -\log(\mathcal{L})$	127
9.6.1	e and μ hypotheses (one track)	127
9.6.2	π^0 and $\gamma\gamma$ hypotheses (two tracks)	128
9.7	Performance	131
10	ν_μ charged current quasi-elastic scattering	140
10.1	Selecting ν_μ CC QE events	140
10.2	Using the sample	143
10.2.1	Q^2 dependence of the CC QE cross section	143
10.2.2	E_ν dependent corrections	147
10.2.3	ν_e from μ decay	150
10.2.4	Closing	152

11 Neutral current π^0 events	154
11.1 π^0 as a $\nu_\mu \rightarrow \nu_e$ background	154
11.2 π^0 event selection	155
11.2.1 Momentum dependent correction	156
11.2.2 Coherent π^0 production (angular distribution)	157
11.2.3 Radiative Δ decay	161
11.2.4 Events from outside the tank	163
12 ν_e selection and signal fit	165
12.1 Selecting ν_e CC QE events	165
12.2 Signal χ^2	171
12.3 Forming the covariance matrix \mathbf{V}	171
12.3.1 The first eight component matrices	171
12.3.2 \mathbf{V}^{det}	173
12.3.3 \mathbf{V}^{qt}	174
12.3.4 \mathbf{V}^{MC}	175
12.4 Signal extraction	175
12.4.1 Best-fit oscillation parameters	175
12.4.2 Global scan confidence intervals	176
12.4.3 Raster scan confidence intervals	178
12.5 Oscillation sensitivity	179
12.6 Testing the fit	183
13 Verifying the predictions	189
13.1 Non-signal samples	189
13.1.1 The comparisons	190
13.2 Lifting blindness	212
14 Results	215
14.1 Summary	215
14.2 Additional discussion	218
14.3 Conclusions	218
III Appendices	224
A Data integrity	225
B Charge and time miscellany	226
B.1 The bleed-off time τ	226
B.2 Criteria for using a Q-ADCn in the Q_{raw} fit	226
B.3 Op-amp slew rate limitation	227
B.4 Additional QT board effects	231

C	Additional reconstruction details	234
C.1	Solid angle $\Omega(r)$ and acceptance $\epsilon(\eta)$	234
C.2	Transmission functions $T_{\text{sci}}(r)$ and $T_{\text{Ch}}(r)$	234
C.3	Scattering tables	235
C.4	Production profiles $\rho_{\text{sci}}(s)$, $\rho_{\text{Ch}}(s)$ and $g(\cos \theta; s)$	237
C.5	Glossary for Chapter 9	239
D	Low energy region	241
	References	243

List of Tables

3.1	Properties of the MiniBooNE mineral oil	29
3.2	Properties of the PMTs	31
3.3	MiniBooNE triggers	38
6.1	Beam profile parameters	64
6.2	Sanford-Wang parameters for π^+ , π^- , and K^0 production	68
6.3	Data used in K^+ production fits	70
6.4	Feynman scaling parameters for K^+ production	70
6.5	p -Be and p -Al inelastic cross section parameters	70
6.6	Neutrinos at the detector, by production channel	74
7.1	Axial vector masses used in nuance	82
8.1	Hadronic cross sections in the detector simulation	87
8.2	Parameters describing the index of refraction	92
10.1	Line cut coefficients for various detector realities	143
12.1	ν_e particle ID cut coefficients	167
12.2	Expected numbers of background and signal events	170
12.3	Summary of the component covariance matrices	172
12.4	Null hypothesis confidence level properties	177
13.1	Summary of comparisons between data and simulation	192
13.2	First round of data/simulation comparisons from the ν_e sample	212
13.3	Second round of data/simulation comparisons from the ν_e sample	214
14.1	Examining the χ^2 value	217
C.1	Parameters for the transmission curves in the reconstruction algorithm	235
D.1	Events in the low energy region	242

List of Figures

1.1	Solar neutrino flux, measurements and prediction	12
1.2	Evidence for $\bar{\nu}_e$ disappearance from KamLAND	13
1.3	Super-Kamiokande L/E spectrum	14
1.4	ν_μ disappearance in MINOS	15
1.5	Long-baseline summary (MINOS, K2K, Super-K)	16
1.6	H. Murayama's parameter space plot	17
1.7	LSND L/E spectrum	18
1.8	LSND allowed region with KARMEN and Bugey exclusion curves	19
2.1	MiniBooNE target and detector locations on the FNAL site	23
2.2	Experimental layout, from target to detector	24
2.3	Proton delivery to MiniBooNE	25
2.4	Drawing of the magnetic focusing horn	26
2.5	Neutrino flux at the detector	27
3.1	A cutaway drawing of the detector	29
3.2	A view of the main and veto detector regions during construction	30
3.3	Elevation view of the detector hall	31
3.4	PMT support structure	32
3.5	Preamplifier circuit	33
3.6	Front-end circuit	34
3.7	Relations between various signals in the electronics	35
3.8	A laser flask	36
4.1	The time quad and T_{raw}	40
4.2	Single-photoelectron charge response	41
4.3	Anode pulse	42
4.4	Integrated pulse, constructed	42
4.5	Integrated pulse, measured	43
4.6	Determining Q_{raw}	44
4.7	Establishing V_q^{ref}	45
4.8	Verification of charge determination procedure	47
4.9	Charge fit residuals	48
4.10	Determining the timing offsets	49

4.11	The extended V_q^{ref}	50
4.12	Time-dependent variation in the charge response	51
4.13	The slew tables	52
5.1	Identifying subevents	56
5.2	Veto multiplicities	57
5.3	Michel electron spectrum	58
5.4	Rejecting cosmic activity; beam timing	59
5.5	A candidate ν_μ CC QE event	60
5.6	Cherenkov ring profiles, data and simulation	61
5.7	Summary of important neutrino processes	62
6.1	Momentum and angle distributions for relevant secondaries	66
6.2	Parametrized and measured π^+ cross sections	67
6.3	Parametrized and measured K^+ cross sections	69
6.4	Parametrized and measured p -Be and p -Al inelastic cross sections	71
6.5	Energy spectra of neutrinos at the detector, by production channel	73
7.1	Charged current neutrino cross sections versus energy	76
7.2	Measured and predicted ν_μ CC QE cross sections	78
7.3	Measured and predicted CC π^+ production cross sections	80
7.4	Measured and predicted resonant π^0 production cross sections	81
8.1	A view of the detector enclosure in the simulation	85
8.2	Rates of optical photon processes	88
8.3	Measurements of light scattering in mineral oil	90
8.4	Fluorescence spectra	91
8.5	R5912 PMT quantum efficiency in oil	93
8.6	PMT timing distribution in laser flask events	94
8.7	PMT time response	95
8.8	PMT charge response	96
8.9	Charge nonlinearities	97
8.10	Progression of data/simulation agreement in Michel energy	101
8.11	Progression of data/simulation agreement in the late time distribution	102
8.12	Reduction of parameter uncertainties, as seen in four distributions	103
9.1	Reconstruction geometry: point source	107
9.2	Reconstruction geometry: extended source	108
9.3	Scintillation production profiles for 300 MeV muons and electrons	109
9.4	Reconstruction geometry: Cherenkov light	110
9.5	Cherenkov profiles for 300 MeV muons and electrons	111
9.6	Reconstruction geometry: indirect isotropic light	112
9.7	Reconstruction geometry: indirect Cherenkov light	113
9.8	An example of $J(s)$	116

9.9	Examples of $f(q; \mu)$	118
9.10	Building the primitive distributions (Part 1)	120
9.11	Building the primitive distributions (Part 2)	121
9.12	Building the primitive distributions (Part 3)	122
9.13	Verifying the Cherenkov primitive distributions	123
9.14	Verifying the scintillation primitive distributions	124
9.15	A graphical view of the one- and two-track event parameters	126
9.16	An example trap in the two-track likelihood surface	129
9.17	Reconstructed radial event position plotted against truth	133
9.18	μ , e , and π^0 reconstructed kinetic energy plotted against truth	134
9.19	Radius, kinetic energy, and direction resolutions	135
9.20	Fitted invariant mass $M_{\gamma\gamma}$ for NC π^0 events (Part 1)	136
9.21	Fitted invariant mass $M_{\gamma\gamma}$ for NC π^0 events (Part 2)	137
9.22	Reconstructed γ energies in NC π^0 events	138
9.23	Energy and direction resolutions for γ 's in NC π^0 events	139
10.1	Range <i>vs.</i> energy for true ν_μ CC QE events (simulation)	142
10.2	Range <i>vs.</i> energy for events passing the initial ν_μ CC QE selection	142
10.3	Establishing the line cut coefficients	143
10.4	ν_μ CC QE final selection	144
10.5	Distribution of Q^2 for ν_μ CC QE candidate events	145
10.6	Insensitivity of the E_ν^{rec} distribution to the low- Q^2 shape	146
10.7	Mapping from E_ν^{rec} to E_ν^{true} for Monte Carlo ν_μ CC QE events	147
10.8	E_ν^{rec} in the ν_μ CC QE sample (data and uncorrected Monte Carlo)	148
10.9	Inferred E_ν^{true} distributions for data and simulation and their ratio	149
10.10	$f(E_\nu^{\text{true}})$ for ten π^+ cross section variants	151
10.11	Example $\nu_\mu \rightarrow \nu_e$ E_ν^{rec} spectrum with(out) the ν_μ CC QE constraint	151
10.12	Cartoon of the production of ν_e from μ^+ decay in the decay pipe	152
10.13	Mapping from E_ν^{rec} to E_ν^{true} for Monte Carlo ν_μ CC QE events	153
10.14	E_ν^{rec} spectrum for ν_e from μ^+ with(out) the ν_μ CC QE constraint	153
11.1	Cartoon of partially contained π^0 events	155
11.2	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$ distribution for events passing the basic π^0 cuts	156
11.3	2D distribution of $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ and $M_{\gamma\gamma}$ after the π^0 pre-selection	157
11.4	$M_{\gamma\gamma}$ and $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ after the π^0 pre-selection (data and simulation)	158
11.5	$M_{\gamma\gamma}$ and $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ after the π^0 pre-selection and one blindness cut	159
11.6	Reconstructed $M_{\gamma\gamma}$ distributions, data and simulation	160
11.7	The NC π^0 weighting function	161
11.8	π^0 momentum distributions and ν_e PID efficiency for π^0 events	162
11.9	Determination of the coherent π^0 fraction	163
11.10	Dirt events: R_{twb} distribution for a low-energy, high-radius sample	164
12.1	PID cut values versus E_e ; cut efficiencies versus E_ν^{rec}	166

12.2	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$ distribution for events passing the ν_e CC QE pre-selection . . .	167
12.3	$M_{\gamma\gamma}$ and $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ distributions after ν_e pre-selection and μ removal . .	168
12.4	$M_{\gamma\gamma}$ and $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ distributions after all other ν_e cuts	169
12.5	Predicted E_ν^{rec} histograms for ν_e backgrounds and example ν_e signal	170
12.6	$\Delta\chi^2(\alpha_{\text{null}})$ distribution from simulated experiments	177
12.7	90% C.L. $\Delta\chi^2$ cutoffs throughout parameter space	179
12.8	90% C.L. sensitivity contour, all errors	180
12.9	90% C.L. sensitivity contours with various errors removed (Part 1)	181
12.10	90% C.L. sensitivity contours with various errors removed (Part 2)	182
12.11	90% C.L. sensitivity contour for the counting analysis	183
12.12	Demonstrating that $\Delta\chi^2(\alpha_{\text{true}})$ is non-negative	184
12.13	Examining many simulated experiments, no systematic errors	185
12.14	Examining many simulated experiments, all errors	186
12.15	“Fully faked” simulated experiments	187
12.16	Sufficiency of number of detector model variants	188
13.1	Distribution of cumulative χ^2 probabilities	193
13.2	ν_μ CC QE plots (Part 1)	195
13.3	ν_μ CC QE plots (Part 2)	196
13.4	ν_μ CC QE plots (Part 3)	197
13.5	High energy ν_e plots (Part 1)	199
13.6	High energy ν_e plots (Part 2)	200
13.7	Sideband A plots (Part 1)	202
13.8	Sideband A plots (Part 2)	203
13.9	Sideband B plots (Part 1)	205
13.10	Sideband B plots (Part 2)	206
13.11	Sideband C plots (Part 1)	208
13.12	Sideband C plots (Part 2)	209
13.13	Sideband C plots (Part 3)	210
13.14	Sideband C plots (Part 4)	211
13.15	The discrepancies δ_i for the E histogram	213
13.16	Oscillation sensitivity for two different E_ν^{rec} thresholds	214
14.1	E_ν^{rec} spectra for data, prediction, and example signal	216
14.2	E_ν^{rec} spectra, including data and the best-fit prediction	216
14.3	90% upper limit on $\sin^2 2\theta$ as a function of Δm^2	217
14.4	A collection of exclusion curves from various searches	219
14.5	Upper limit and global scan contours together	220
14.6	Simple global scan contours with approximate unified approach contours . .	221
14.7	Combined LSND-KARMEN allowed region and our results	222
14.8	E_ν^{rec} spectrum below 475 MeV	223
B.1	V_q for a simple pulse	227

B.2	Bleed-off time for the laser sync. channel	228
B.3	The V_q op-amp	229
B.4	Slew rate limitation for a sample pulse	229
B.5	Measuring the op-amp slew rate	230
B.6	The rising edge of V_q at 1 PE	230
B.7	Seeing the op-amp slew rate in data	231
B.8	Charge proportionality breakdown	232
B.9	Hard limit on charge	233
C.1	Transmission curves for Cherenkov and scintillation light	236
C.2	The scintillation scattering table	237
C.3	Part of the Cherenkov scattering table	238
D.1	Distribution of reconstructed neutrino energies for ν_e candidate events . . .	241
D.2	Cross sections for γ interactions on carbon as a function of energy	243
D.3	Reconstructed energy E and direction U_z in two energy regions	244

A SEARCH FOR MUON NEUTRINO TO
ELECTRON NEUTRINO OSCILLATIONS
AT $\Delta m^2 > 0.1 \text{ eV}^2$

Part I

Introduction

Chapter 1

Neutrino oscillations and MiniBooNE

*I have done a terrible thing. I have postulated
a particle than cannot be detected.*

Wolfgang Pauli

In 1953, twenty-three years after Pauli suggested the existence of a weakly interacting light neutral particle to explain the continuous nature of the beta decay spectrum [1], neutrinos emanating from a nuclear reactor were detected [2] [3], and experimental neutrino physics was born. Today, much of the experimental effort relates in some way to neutrino mass, and in this chapter we review the phenomenology of massive neutrinos. In particular, we introduce neutrino oscillations and the motivation for the work described in this thesis.

1.1 Extending the standard model

In the standard electroweak model of Glashow, Weinberg, and Salam [4], neutrinos are massless spin- $\frac{1}{2}$ members of left-handed lepton doublets:

$$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L \quad \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L \quad \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L \quad e_R \quad \mu_R \quad \tau_R . \quad (1.1)$$

The left- and right-handed fields shown are the chiral projections¹

$$\psi_{L,R} = \frac{1 \mp \gamma^5}{2} \psi \quad (1.2)$$

¹We adopt the “Dirac” convention [5] for the Dirac matrices, wherein:

$$\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

of the four-component Dirac spinors ψ which are solutions to the Dirac equation [6]. Since chirality is conserved for massless particles, right-handed neutrinos (and left-handed antineutrinos) do not participate in the weak interaction due to the V–A nature of the weak current ($J^\mu \sim \gamma^\mu(1 - \gamma^5)$). Since they cannot interact via the strong or electromagnetic forces either, the $(\nu_x)_R$ fields are absent from the standard model.

1.1.1 Neutrino masses

Detailed treatments of neutrino masses may be found in Refs. [6], [7], and [8]. This discussion follows [6].

The Dirac Lagrangian²

$$\mathcal{L} = \bar{\psi} (i\gamma^\mu \partial_\mu - m) \psi \quad (1.3)$$

includes a kinetic energy term and a mass term $\mathcal{L}_M = -m\bar{\psi}\psi$ which we can write in terms of the chiral states as

$$-\mathcal{L}_M = m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L) . \quad (1.4)$$

(Note that $\bar{\psi}_L\psi_L = \bar{\psi}_R\psi_R = 0$.) If we assume in the neutrino case ($\psi = \nu$) that a right-handed field ψ_R (and $\bar{\psi}_R = \psi_R^\dagger\gamma^0$) exists, this mass term can be present in the theory if $m = m_\nu \neq 0$. However, the $\mathcal{O}(\text{eV})$ upper limits on neutrino mass [9] imply that the Yukawa coupling λ_ν which leads to m upon electroweak symmetry breaking [10] must be anomalously and inexplicably small, $\lambda_\nu < \mathcal{O}(10^{-11})$, compared to the Yukawa couplings for the other fermions ($\sim 10^{-6}$ to ~ 1).

Since the neutrino is neutral, the above mass term is not the most general. We can combine ψ with its charge conjugate field ψ^c to form additional Lorentz invariant bilinears which, for charged particles, would violate charge conservation. The most general possible Lagrangian for neutrinos includes these ‘‘Majorana’’ mass terms:

$$-2\mathcal{L} = \begin{pmatrix} \bar{\psi}_L & \bar{\psi}_L^c \end{pmatrix} \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix} \begin{pmatrix} \psi_R^c \\ \psi_R \end{pmatrix} + \text{h.c.} \quad (1.5)$$

Setting $m_L = m_R = 0$ recovers the ‘‘Dirac’’ case of Eq. (1.4), with $m = m_D$.

Diagonalizing Eq. (1.5) yields mass eigenvalues

$$m_{1,2} = \frac{1}{2} \left[(m_L + m_R) \pm \sqrt{(m_L - m_R)^2 + 4m_D^2} \right] . \quad (1.6)$$

(Positive masses can always be obtained through a redefinition of the fields.) If m_D is similar to other fermion masses ($m_D \sim m_e$) and $m_R \gg m_D, m_L$, the eigenvalues become

$$\begin{aligned} m_1 &\approx \frac{m_D^2}{m_R} \\ m_2 &\approx m_R \end{aligned}$$

²We use natural units $\hbar = c = 1$ throughout this chapter.

where one eigenstate (call it ν) has mass $m_1 \ll (m_D \sim m_e)$ and another (N) has mass $m_2 \gg m_e$. Many standard model extensions place m_R at the GUT scale³ (motivating the condition $m_R \gg m_D, m_L$) and offer a candidate heavy neutrino field N , either sterile or coupled through a right-handed current mediated by a GUT-scale boson [11] [12]. This formalism, known as the “see-saw” mechanism in the literature, provides us a light neutrino while avoiding the fine tuning needed with a pure Dirac mass term.

1.1.2 Neutrino mixing

For N_ν neutrino generations, there will be N_ν mass eigenstates ν_i . Nothing requires that these eigenstates correspond to the flavor (or weak) eigenstates ν_α which appear in the weak current. In general, the mass (ν_i) and flavor (ν_α) eigenstates will be connected by an $N_\nu \times N_\nu$ unitary matrix \mathbf{U} :

$$|\nu_\alpha\rangle = \sum_{i=1}^{N_\nu} U_{\alpha i}^* |\nu_i\rangle . \quad (1.7)$$

\mathbf{U} (often called the Maki-Nakagawa-Sakata (MNS) matrix [13]) is analogous to the perhaps more familiar CKM matrix governing quark mixing.

As we shall see, the evidence is overwhelming that neutrinos have mass and that the weak eigenstates are admixtures of the mass eigenstates. We review some of the possible experimental signatures of neutrino mass and mixing before discussing the phenomenon relevant to the present work: neutrino oscillations.

1.2 Implications

β -decay spectrum

Nuclear β -decay

$${}^A_Z X_i \rightarrow {}^A_{Z+1} X_f + e^- + \bar{\nu}_e \quad (1.8)$$

produces electrons with maximum energy

$$E_e^{\max} = E_0 - m_\nu , \quad (1.9)$$

where $E_0 = M_i - M_f$ is the mass difference between the initial and final nuclei and where we have neglected the nuclear recoil energy which can be no larger than $\frac{E_0^2 - m_e^2}{2M_i}$ (or about 0.003 eV for tritium β -decay). With massless neutrinos, the energy spectrum $K(E_e)$ of the outgoing electron approaches the endpoint $E_e^{\max} = E_0$ linearly. For $m_\nu \neq 0$, the endpoint shifts by m_ν and the linearity of $K(E_e)$ fails near $(E_e^{\max} - E_e) \sim m_\nu$, with the spectrum

³ 10^{14} GeV or so. To get $m_1 \sim 0.05$ eV (the neutrino mass lower limit implied by atmospheric neutrino oscillations, §1.3.4) using $m_R = 10^{14}$ GeV, one needs $m_D \sim 70$ GeV which, while stretching the relation $m_D \sim m_e$, is a perfectly reasonable electroweak value.

taking the form [14]

$$K(E_e) \sim \left[(E_0 - E_e) \sqrt{(E_0 - E_e)^2 - m_\nu^2} \right]^{\frac{1}{2}} . \quad (1.10)$$

By measuring the energies of near-endpoint electrons, one can look for this m_ν -induced modification to $K(E_e)$. If found, the observed spectrum would provide a measure of m_ν . Tritium is a good choice for these so-called “kinematic” (also, “direct”) mass searches, as (a) its relatively low E_0 (18.6 keV) keeps the fraction of near-endpoint decays high⁴, (b) its half-life is short enough that useful rates can be obtained without too much line-broadening within the source, and (c) final state electron and nuclear energy levels (which confuse the shape of Eq. (1.10)) can be calculated.

β -decay experiments have thus far not observed a neutrino mass signal, and the Particle Data Group global limit is $m_\nu < 2$ eV [9]. Although a $\bar{\nu}_e$ is emitted in the decay, we cannot treat this as a limit on $m_{\bar{\nu}_e}$ or (invoking CPT invariance) m_{ν_e} . These quantities are not well-defined in light of neutrino mixing, and the mass actually measured in these experiments is given by [6]

$$m_\nu^2 = \sum_i |U_{ei}|^2 m_i^2 , \quad (1.11)$$

assuming the experimental resolution δ_E is much larger than any mass difference ($\delta_E \gg |m_i - m_j|$). In the case that $|U_{ei}|^2 \ll 1$ for all i except one, m_ν can be approximately identified as the “electron neutrino mass”.

Neutrinoless double- β decay

Many nuclei can undergo the second-order “double β -decay” ($\beta\beta$) transition

$${}^A_Z\text{X}_i \rightarrow {}^A_{Z+2}\text{X}_f + e^- + e^- + \bar{\nu}_e + \bar{\nu}_e \quad (1.12)$$

(or similar β^+ and electron capture transitions). In practice, this reaction is relevant only for the few nuclei (⁷⁶Ge, ¹³⁶Xe, others) that can undergo $\beta\beta$ -decay even though they are stable to β -decay. For these even-even nuclei, the nominal β^\pm -decay daughter has a higher mass due the nucleon pairing force [16]. Thus, β -decay is kinematically forbidden and cannot overpower the otherwise rarer $\beta\beta$ -decay.

If neutrinos are massive Majorana particles, they can mediate a related neutrinoless double beta decay reaction ($0\nu\beta\beta$)

$${}^A_Z\text{X}_i \rightarrow {}^A_{Z+2}\text{X}_f + e^- + e^- . \quad (1.13)$$

Majorana neutrinos ($\nu = \bar{\nu}$) are required since the $\bar{\nu}$ propagator leaving one weak vertex is absorbed as a ν at the other. Massive neutrinos are needed so that the chirality may be different at the two vertices as required by the V-A current (although models involving right-handed currents often require nonvanishing neutrino mass for $0\nu\beta\beta$ anyway [17]).

⁴although only $\sim 10^{-13}$ of decays put E_e within 1 eV of the endpoint [15]

The search for $0\nu\beta\beta$ is active, with upcoming experiments reaching for $m=0.1$ eV sensitivity [18]. A controversial positive signal was reported by Klapdor-Kleingrothaus *et al.* [19]; it has not yet been independently verified. If a signal is confirmed, the measured $0\nu\beta\beta$ rate can be related to an effective Majorana mass $m_{\beta\beta}$, with

$$m_{\beta\beta}^2 = \left| \sum_i U_{ei}^2 m_i \right|^2 . \quad (1.14)$$

In contrast to m_ν in Eq. (1.11), complex CP phases in the MNS matrix can lead to cancellations in this sum, reducing $m_{\beta\beta}^2$ and the $0\nu\beta\beta$ rate. Additionally, large uncertainties in necessary nuclear matrix elements limit the precision to which $m_{\beta\beta}$ can be determined.

Cosmology

In the hot big bang model, light neutrinos in the early universe were in thermal equilibrium until around $kT = 1$ MeV. Since then, neutrinos have been frozen out and should be present today with a number density related to that observed for the cosmic microwave background photons [12]

$$n_\nu = \frac{3}{11} n_\gamma = 113 \text{ cm}^{-3} \text{ (per generation)}. \quad (1.15)$$

The contribution Ω_ν of these neutrinos to the cosmological density parameter Ω is

$$\begin{aligned} \Omega_\nu &= \frac{1}{\rho_c} n_\nu \sum_{i=1}^{N_\nu} m_i \\ &= \frac{1}{(93.5 \text{ eV}) h^2} \sum_{i=1}^{N_\nu} m_i , \end{aligned} \quad (1.16)$$

where $\rho_c = 1.88 \times 10^{-29} h^2 \text{ g/cm}^3$ is the critical density and h is Hubble's constant parametrized as $H_0 = (100 h) \text{ Mpc}^{-1} \text{ km s}^{-1}$. The observation that $\Omega \approx 1$ [20] (and, therefore, $\Omega_\nu < 1$) implies an upper limit on the sum of light neutrino masses (taking $h \approx 0.5$)

$$\sum_{i=1}^{N_\nu} m_i \lesssim 20 \text{ eV} \quad (1.17)$$

Much heavier neutrinos, $m > 1$ MeV, become non-relativistic before decoupling, and their number density is given not by Eq. (1.15) but rather by the Boltzmann factor, which includes a mass-dependent suppression $n_\nu \sim e^{-m/T}$ [12]. For $m \gtrsim 2$ GeV, the resulting density parameter Ω_ν again becomes compatible with observations. Intermediate mass neutrinos ($20 \text{ eV} < m < 2 \text{ GeV}$) are allowed only if they are unstable with lifetimes shorter than the age of the universe [8].

In addition to contributing to the matter density in the universe, neutrinos can influence the formation of large scale structure (*e.g.*, superclusters). Neutrinos streaming out of regions of relatively high density into regions of lower density tend to erase small scale

density perturbations. This *collisionless damping* has observable effects on the matter distribution [6]. The mass limits inferred from the measured matter power spectrum are (assuming three neutrino generations) [21]

$$\sum_{i=1}^3 m_i < 0.7 - 2.2 \text{ eV} . \quad (1.18)$$

Onward

Despite the variety of experimental possibilities (including many not mentioned here⁵), only upper limits on neutrino masses have been set thus far *except* via the interferometric phenomenon of neutrino oscillations.

1.3 Neutrino oscillations

A neutrino created through a charged current weak interaction is, by definition, in a flavor eigenstate. The particular flavor eigenstate $|\nu_\alpha\rangle$ ($\alpha = e, \mu, \tau$) can be determined experimentally by observing the charged lepton α^\pm involved. (An e^+ gets produced alongside a ν_e , and so forth.) Neutrino mixing implies that $|\nu_\alpha\rangle$ is not (in general) a mass eigenstate and, therefore, is not necessarily a stationary state of the free-space Hamiltonian. Consequently, a neutrino of definite flavor need not keep that flavor as it propagates.

1.3.1 General formalism

A mass eigenstate $|\nu_i\rangle$ with mass m_i and momentum p evolves in time according to

$$|\nu_i(t)\rangle = e^{-iE_i t} |\nu_i\rangle , \quad (1.19)$$

where $E_i = \sqrt{p^2 + m_i^2}$. A flavor eigenstate $|\nu\rangle = |\nu_\alpha\rangle$ produced with momentum p is a superposition of these, and its time evolution is obtained by combining Eq. (1.19) and Eq. (1.7):

$$\begin{aligned} |\nu(t)\rangle &= \sum_i U_{\alpha i}^* e^{-iE_i t} |\nu_i\rangle \\ &= \sum_i \sum_\beta U_{\beta i} U_{\alpha i}^* e^{-iE_i t} |\nu_\beta\rangle , \end{aligned} \quad (1.20)$$

where in the second line we have used the unitarity of \mathbf{U} to invert Eq. (1.7):

$$|\nu_i\rangle = \sum_\beta U_{\beta i} |\nu_\beta\rangle . \quad (1.21)$$

⁵neutrino magnetic moments [22], neutrino decay [23], supernovae observations [24]

We can immediately write down the probability $P(\nu_\alpha \rightarrow \nu_\beta)(t)$ that the neutrino $|\nu\rangle$, if detected at time t , will have its flavor measured to be β :

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta)(t) &= |\langle \nu_\beta | \nu(t) \rangle|^2 \\ &= \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-iE_i t} \right|^2. \end{aligned} \quad (1.22)$$

A highly relativistic neutrino covers a distance $L \approx t$ in time t (recalling our unit convention $c = 1$) and has $E_i \approx p + \frac{m_i^2}{2p}$. Thus,

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-ipL} e^{-i\frac{m_i^2}{2p}L} \right|^2 \\ &= \sum_i \sum_j U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j} e^{\frac{-iL}{2p}(m_i^2 - m_j^2)}. \end{aligned} \quad (1.23)$$

Regrouping terms in the sum and using the unitarity of \mathbf{U} yields [9]

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= \delta_{\alpha\beta} \\ &\quad - 4 \sum_{i>j} \Re(U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*) \sin^2[\kappa \Delta m_{ij}^2 (L/E)] \\ &\quad + 2 \sum_{i>j} \Im(U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*) \sin[2\kappa \Delta m_{ij}^2 (L/E)], \end{aligned} \quad (1.24)$$

where $\kappa = 1.27 \text{ GeV/km/eV}^2$, $\Delta m_{ij}^2 = m_i^2 - m_j^2$, L is the source-to-detector distance (the “baseline”), and E is the energy of the neutrino.⁶ Some comments:

- The sinusoidal nature of $P(\nu_\alpha \rightarrow \nu_\beta)(L)$ motivates the appellation “neutrino oscillations”.
- Experiments typically search for oscillations by looking for a deficit of a particular flavor (“disappearance” experiments) or an excess of a particular flavor, often one that is nearly absent in the neutrino flux to begin with (“appearance” experiments).
- If CPT holds, the conjugate oscillation process has a probability $P(\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha)$ given by Eq. (1.24) with the modification $\mathbf{U} \mapsto \mathbf{U}^*$.
- If \mathbf{U} is real, neutrinos and antineutrinos have identical oscillation probabilities. Additionally, the last term in Eq. (1.24) vanishes, and the signs of the Δm_{ij}^2 splittings become irrelevant for vacuum (*i.e.* non-MSW, §1.3.3) oscillations.⁷

⁶The energy is well-defined only because of the relativistic approximation made in deriving Eq. (1.24). For a more detailed treatment which addresses this and other nuances, see Ref. [25].

⁷One often sees Δm_{ij}^2 in the literature where $|\Delta m_{ij}^2|$ is intended. For better or for worse, we, too, use this shorthand when no ambiguity results.

- Neutrino oscillations depend only on mass-squared differences Δm_{ij}^2 . The values of the individual neutrino masses cannot be determined with oscillations.
- The experimental parameter L/E controls which mass-squared splittings can be observed. Terms in Eq. (1.24) with $\kappa\Delta m_{ij}^2(L/E)\ll 1$ will be experimentally negligible. E or L/E spectral distortion may be observable if $\kappa\Delta m_{ij}^2(L/E)\sim 1$. Oscillations may also be visible when $\kappa\Delta m_{ij}^2(L/E)\gg 1$, but resolving Δm_{ij}^2 will not be possible.

1.3.2 Quasi-two-neutrino oscillations

If either:

1. the produced flavor eigenstate $|\nu_\alpha\rangle$ couples significantly to only two mass eigenstates $|\nu_m\rangle$ and $|\nu_n\rangle$, or
2. one family of splittings ΔM^2 is much larger than all the other splittings and the experiment is sensitive to ΔM^2 (*i.e.*, $\kappa\Delta M^2(L/E) \gtrsim 1$)

then Eq. (1.24) can be greatly simplified to

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta \sin^2[\kappa\Delta m^2(L/E)] . \quad (1.25)$$

The splitting Δm^2 represents either Δm_{mn}^2 (Case 1) or ΔM^2 (Case 2), and the mixing parameter θ can be written in terms of the elements of \mathbf{U} [9]. Eq. (1.25) has exactly the form one obtains by considering two neutrino mass eigenstates separated in squared mass by Δm^2 and related to two flavor eigenstates via the mixing matrix

$$\mathbf{U} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} . \quad (1.26)$$

Many oscillation analyses (including this one) use this “quasi-two-neutrino” oscillation expression. Results are often presented by specifying allowed/excluded regions in the two-dimensional space $(\sin^2 2\theta, \Delta m^2)$.

1.3.3 Matter effects

Neutrinos traveling through matter undergo forward scattering which can be treated with an additional potential term in the Hamiltonian. Since electron neutrinos experience charged current scattering off the electrons in bulk matter while muon and tau neutrinos do not, the ν_e potential has an extra contribution

$$V = \sqrt{2}G_F N_e \quad (1.27)$$

where G_F is Fermi’s constant and N_e is the electron density of the propagation medium [6]. In the context of neutrino oscillations, this extra potential can be cast as an effective ν_e

mass which, upon diagonalization of the mass matrix, results in modified mass eigenvalues. For $\nu_e \leftrightarrow \nu_{\mu,\tau}$ oscillations in matter, the splitting Δm^2 of Eq. (1.25) gets replaced by

$$\Delta m_m^2 = \Delta m^2 f_{\text{MSW}} \quad (1.28)$$

and the mixing parameter $\sin^2 2\theta$ becomes

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta}{f_{\text{MSW}}}, \quad (1.29)$$

where the energy-dependent scale factor f_{MSW} is given by

$$f_{\text{MSW}} = \sqrt{\left(\frac{2VE}{\Delta m^2} - \cos 2\theta\right)^2 + \sin^2 2\theta}. \quad (1.30)$$

For an electron density of

$$N_e = \frac{\Delta m^2 \cos 2\theta}{2\sqrt{2}G_F E} \quad (1.31)$$

flavor oscillations become maximal ($\sin^2 2\theta_m = 1$) even for small vacuum mixing angles ($\sin^2 2\theta \ll 1$), and the oscillation wavelength is scaled up from its vacuum value by $1/\sin 2\theta$.

This phenomenon, often termed the Mikheyev-Smirnov-Wolfenstein (MSW) effect [26], is negligible for MiniBooNE (the experiment described in this thesis). Typical MSW-related quantities in MiniBooNE are $E=0.8$ GeV, $\Delta m^2=1$ eV², $\sin^2 2\theta=0.004$, and $N_e=10^{24}$ cm⁻³, yielding a scale factor of $(1 - f_{\text{MSW}}) = 3 \times 10^{-5}$. Within the entire LSND allowed region (§1.3.5), the effect is never larger than $(1 - f_{\text{MSW}}) = \mathcal{O}(10^{-4})$. We do not mention the MSW effect outside of this chapter.

1.3.4 Experimental evidence

Solar and reactor neutrinos

The fusion reactions that power the sun, such as



produce electron neutrinos. In 1964, John Bahcall put forth a prediction for the flux of these solar neutrinos reaching Earth [27] and Raymond Davis described [28] and performed [29] an experiment to measure them. The experiment used 390 000 liters of perchloroethylene whose chlorine served as a target for the inverse β -decay process⁸



⁸This reaction has a threshold of 814 keV and cannot detect the pp neutrinos of Eq. (1.32). It is sensitive primarily to ${}^8\text{B}$ neutrinos [6].

The produced ^{37}Ar was periodically collected and counted. The solar ν_e flux inferred from the ^{37}Ar production rate was found to be significantly lower than that predicted. After another quarter-century of data collection and solar model refinements, the discrepancy remained:

$$\begin{aligned} \text{observed: } & 2.56 \pm 0.16 \text{ SNU}^9 \quad (\text{Homestake chlorine experiment (1996) [30]}) \\ \text{predicted: } & 7.6_{-1.1}^{+1.3} \text{ SNU} \quad (\text{Bahcall } et \text{ al. (2001) [31]}) . \end{aligned}$$

This so-called ‘‘solar neutrino problem’’, wherein the observed rate of solar neutrinos falls short of expectation, was confirmed in similar gallium-based experiments (GALLEX/GNO and SAGE) which were sensitive to lower energy neutrinos, including those from the pp reaction above [32] [33] [34].

Cherenkov detectors also recorded solar ν_e deficits. Kamiokande [35] and later Super-Kamiokande (Super-K) [36], using 2 kiloton (Kamiokande) and 22 kiloton (Super-K) fiducial volumes of water surrounded by grids of photomultiplier tubes, observed fewer solar neutrino elastic scattering events

$$\nu + e^- \rightarrow \nu + e^- \tag{1.34}$$

than expected. (This reaction occurs for ν_e as well as ν_μ and ν_τ , although the cross section differs for ν_e since a charged current channel is available.) The more recent Sudbury Neutrino Observatory (SNO) performed a similar measurement using heavy water [37], which offers two additional detection channels: (1) charged current dissociation of deuterium

$$\nu_e + d \rightarrow e^- + p + p \tag{1.35}$$

which sees only the ν_e flux, and (2) neutral current dissociation

$$\nu + d \rightarrow \nu + p + n \tag{1.36}$$

which is flavor-agnostic. SNO could thus measure the ν_e flux alongside the *total* neutrino flux. As shown in Figure 1.1, the inferred total flux is in good agreement with the standard solar model. However, only one-third of the flux is detected as ν_e , with the rest appearing as ν_μ or ν_τ , flavors which are not produced in the sun. SNO’s landmark result solidified the oscillation explanation for the solar neutrino problem.

Commercial fission reactors produce electron antineutrinos with energies comparable to those of solar neutrinos. Reactor-based experiments at baselines ranging from 10 m to 1 km saw no deficits in $\bar{\nu}_e$ event rates. In 2003, the much longer baseline ($L \sim 200$ km) KamLAND experiment reported a $\bar{\nu}_e$ flux that was lower than expectation by the factor $0.611 \pm 0.085(\text{stat.}) \pm 0.041(\text{syst.})$ [39]. The KamLAND signal (Figure 1.2), along with the previous null reactor results and the amassed solar neutrino observations, forms a consistent picture in which electron (anti)neutrinos undergo quasi-two-neutrino oscillations with a mass-squared splitting of $\Delta m_\odot^2 = 7.9_{-0.5}^{+0.6} \times 10^{-5} \text{ eV}^2$ and a mixing angle of $\tan^2 \theta_\odot =$

⁹ ‘‘solar neutrino unit’’ = 10^{-36} argon atoms produced per chlorine atom per second

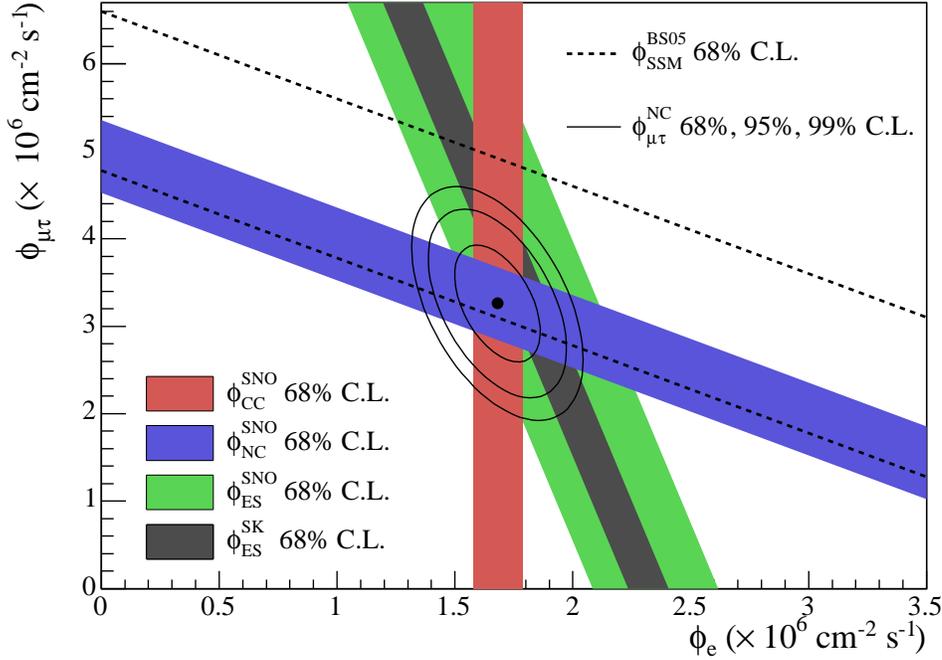


Figure 1.1: Solar neutrino flux: measurements (SNO and Super-K) and prediction. The three reactions (elastic scattering (ES, gray and green), charged current dissociation (CC, red), neutral current dissociation (NC, blue)) have different sensitivities to the ν_e flux and the ν_μ/ν_τ flux. SNO’s inferred fluxes ϕ_e and $\phi_{\mu\tau}$ are shown here along with the ES-based measurement from Super-K. The standard solar model 68% C.L. prediction is bounded by the dashed lines. From Ref. [38].

$0.40^{+0.10}_{-0.07}$ [40]. The vastly different solar and reactor neutrino baselines are reconciled by important MSW effects within the sun.

Atmospheric and accelerator neutrinos

Primary cosmic rays (mostly protons, with some heavier nuclei) interact in the upper atmosphere producing secondary hadrons (π^\pm and K^\pm) whose decay chains include neutrinos. For example,

$$\pi^+ \rightarrow \mu^+ \nu_\mu \quad (1.37)$$

followed by

$$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \quad (1.38)$$

results in two muon (anti)neutrinos and one electron neutrino. The initial indications that these neutrinos oscillate came from the Kamiokande [41] [42] and IMB [43] groups, which observed anomalously low ν_μ rates relative both to absolute predictions and to the observed ν_e rates. (The flavor ratio ν_μ/ν_e can be predicted more precisely than the absolute rates

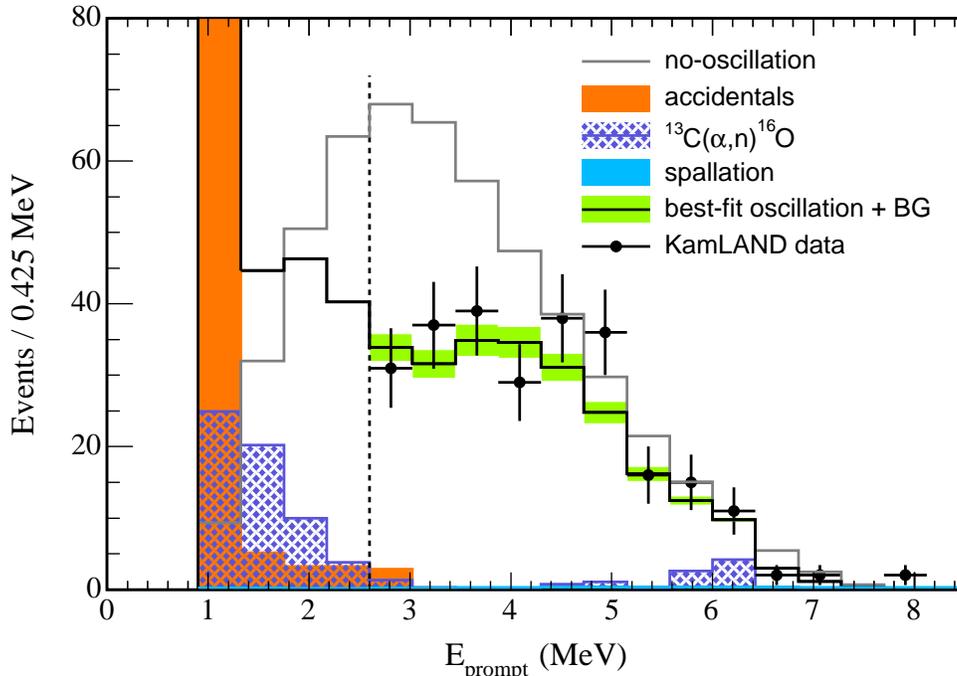


Figure 1.2: Evidence for $\bar{\nu}_e$ disappearance from KamLAND [40]. The energy-dependent deficit in the spectrum is clear (data and best-fit compared with “no-oscillation” curve).

since the ratio is governed primarily by the decay modes of the secondaries.) Super-K subsequently presented evidence strongly favoring the hypothesis that atmospheric muon neutrinos oscillate into tau neutrinos [44] [45]. The L/E -dependent ν_μ flux suppression seen in Super-K is shown in Figure 1.3.

A particle accelerator can be used to create a neutrino source not unlike the natural atmospheric one. By sending a beam of protons through a fixed target, one can generate mesons whose decay chains produce neutrinos, as in Eqs. (1.37) and (1.38). The MINOS experiment used such a beam to perform a search for ν_μ disappearance in the Δm^2 region preferred by the atmospheric data. Figure 1.4 shows the measured MINOS ν_μ spectrum which includes a clear deficit relative to a no-oscillation prediction. The K2K collaboration, using a similar beam aimed at the Super-K detector, also reported evidence of ν_μ disappearance. The MINOS and K2K observations are consistent with the Super-K atmospheric $\nu_\mu \rightarrow \nu_\tau$ oscillation parameters. Figure 1.5 shows this agreement. Taken together, these results imply (at $\sim 90\%$ C.L.): $1.9 \times 10^{-3} \text{ eV}^2 < \Delta m_{\text{atm}}^2 < 3.0 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{\text{atm}} > 0.9$. No evidence for ν_e appearance or disappearance has been seen at Δm_{atm}^2 , with the CHOOZ experiment setting the best limit: $\sin^2 2\theta < 0.1$ (90% C.L.) [47].

We show in Figure 1.6 a graph from H. Murayama [49] which displays the allowed and excluded parameters from (almost) all oscillation experiments to date. (Notably, the MINOS and K2K results have not yet been incorporated.) The LSND signal in the graph is our next topic.

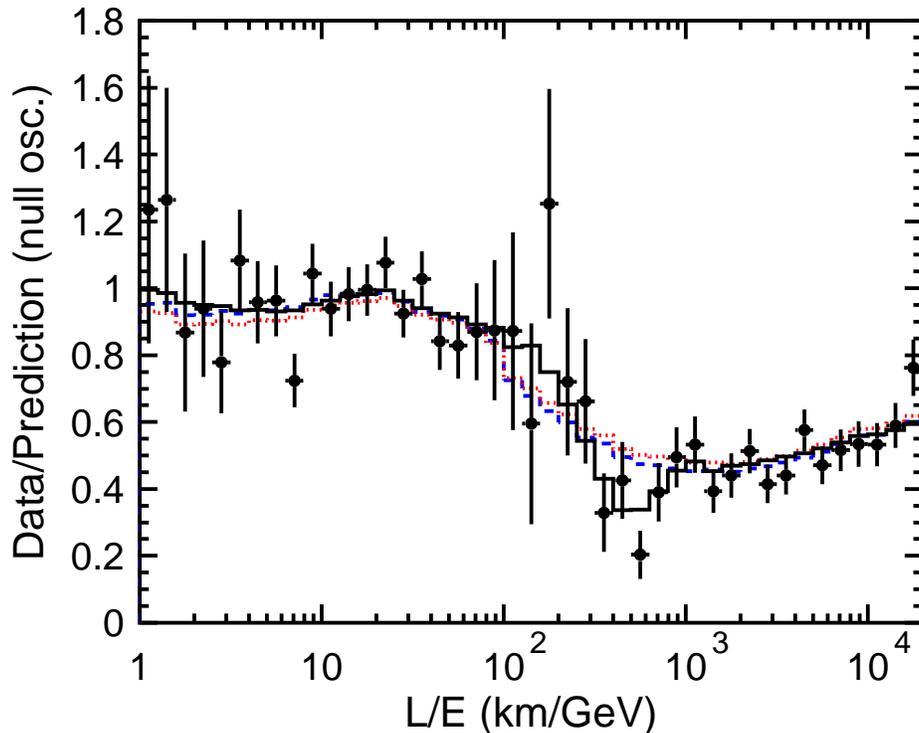


Figure 1.3: Super-Kamiokande L/E spectrum from atmospheric neutrinos. The data (solid points) and best-fit $\nu_\mu \rightarrow \nu_\tau$ oscillations (solid) are shown as ratios to the no-oscillation Monte Carlo prediction. The best-fit expectations for neutrino decay (dashed) and neutrino decoherence (dotted) are also shown. From Ref. [46].

1.3.5 LSND

The Liquid Scintillator Neutrino Detector (LSND) at Los Alamos National Laboratory looked for the reaction

$$\bar{\nu}_e + p \rightarrow e^+ + n \quad (1.39)$$

using ~ 200 tons of mineral oil doped with b-PBD scintillator and viewed by 1,220 photomultiplier tubes. This reaction was identified by the prompt light from the positron and the delayed 2.2 MeV γ released in the neutron capture $n + p \rightarrow d + \gamma$. The detector was situated 30 m from a stopped- μ^+ $\bar{\nu}_\mu$ source. Appearance of $\bar{\nu}_e$ events in excess of the small expected backgrounds would suggest $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ flavor change.

Indeed, a 3.8σ -significant excess of $87.9 \pm 22.4 \pm 6.0$ events was found [50]. Figure 1.7 shows the observed L/E spectrum. The excess is consistent with $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations with $\Delta m^2 > 0.03$ eV². Figure 1.8 shows the LSND allowed oscillation parameters along with two exclusion curves discussed below. The best-fit oscillation scenario gives a spectrum-averaged oscillation probability of $(0.264 \pm 0.067 \pm 0.045)\%$. Earlier analyses using a subset of the final data sample were consistent with this one [51] [52].

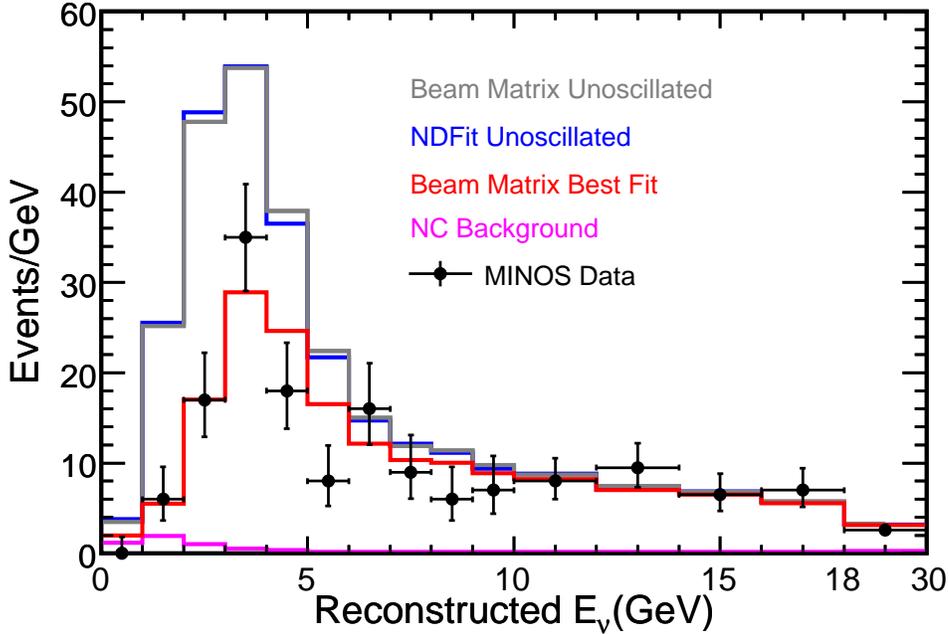


Figure 1.4: Observation of ν_μ disappearance in MINOS (data and best-fit compared with unoscillated spectrum) [48].

In 1998, LSND presented a $\nu_\mu \rightarrow \nu_e$ oscillation analysis using the ν_μ flux from π^+ decay-in-flight. The energies ($60 \text{ MeV} < E_\nu < 200 \text{ MeV}$), background sources, and event signatures all differed from the $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ case above. Thus, while less sensitive than the μ^+ decay-at-rest results, this analysis served as a useful cross check for the oscillation interpretation of the $\bar{\nu}_e$ excess. A total of 40 beam-on events were observed, in excess of the expected 21.9 ± 2.1 events. The inferred oscillation probability of $(0.26 \pm 0.1 \pm 0.05)\%$ is consistent with the decay-at-rest value.

Bugey

A few year prior, an experiment near the Bugey nuclear reactor in France collected large samples of $\bar{\nu}_e$ events in three detectors at baselines of 15 m, 40 m, and 95 m [54]. Each detector comprised ~ 100 optically independent modules containing liquid scintillator doped with ${}^6\text{Li}$. As in LSND, the signature reaction was $\bar{\nu}_e$ -induced inverse β -decay, Eq. (1.39). Photomultiplier tubes viewing the scintillator recorded the prompt light from the positrons and the delayed light from the neutron captures on ${}^6\text{Li}$. The three-detector configuration reduced flux uncertainties in the oscillation search to 2%.

The experiment saw no evidence for $\bar{\nu}_e$ disappearance. Since reactor antineutrinos have $E_\nu \sim 3 \text{ MeV}$, the Bugey baselines of $L \sim 40 \text{ m}$ imply sensitivity to $\bar{\nu}_e$ oscillations at $\Delta m^2 \sim (E/L)/\kappa \sim 0.06 \text{ eV}^2$, which is within the LSND allowed region. As shown in Figure 1.8, the null Bugey result excludes LSND allowed parameters with $\sin^2 2\theta > 0.04$.

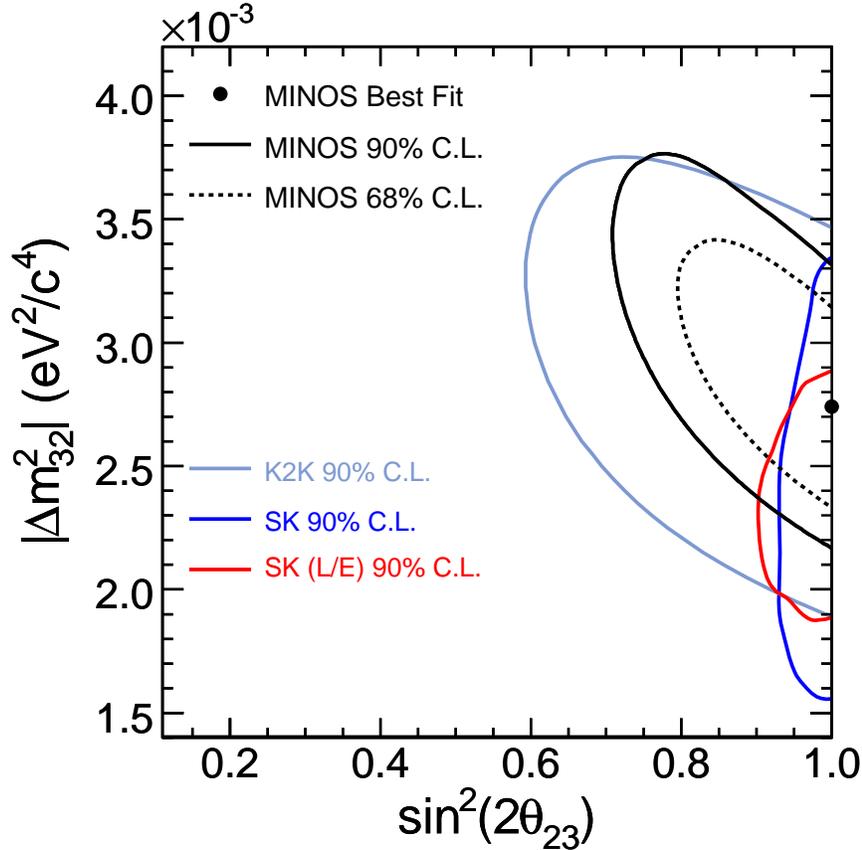


Figure 1.5: Preferred oscillation parameters from MINOS, K2K, and Super-K [48].

KARMEN

The Karlsruhe Rutherford Medium Energy Neutrino (KARMEN) experiment searched for oscillations using a stopped- μ^+ source of $\bar{\nu}_\mu$ similar to that employed by LSND [53]. The KARMEN detector was a 56-ton segmented liquid scintillator calorimeter located 18 m from the neutrino source. The $\bar{\nu}_e$ detection principle was similar – prompt positron light followed by neutron capture – and the capture efficiency was enhanced by the placement of Gd-doped linings between the detector modules. KARMEN also used beam timing to reduce backgrounds from π^+ -produced neutrinos. ($\tau_{\pi^+} = 26$ ns, $\tau_\mu = 2.2$ μ s.) The experiment recorded 15 candidate $\bar{\nu}_e$ events, consistent with the 15.8 ± 0.5 events expected from non-oscillation sources. Figure 1.8 shows the resulting exclusion curve.

Implications of the LSND signal

While the Bugey and KARMEN results limit the range of allowed parameters, the LSND excess is still consistent with $0.1 \text{ eV}^2 < \Delta m_{\text{LSND}}^2 < 10 \text{ eV}^2$ and $0.001 < \sin^2 2\theta_{\text{LSND}} < 0.04$.

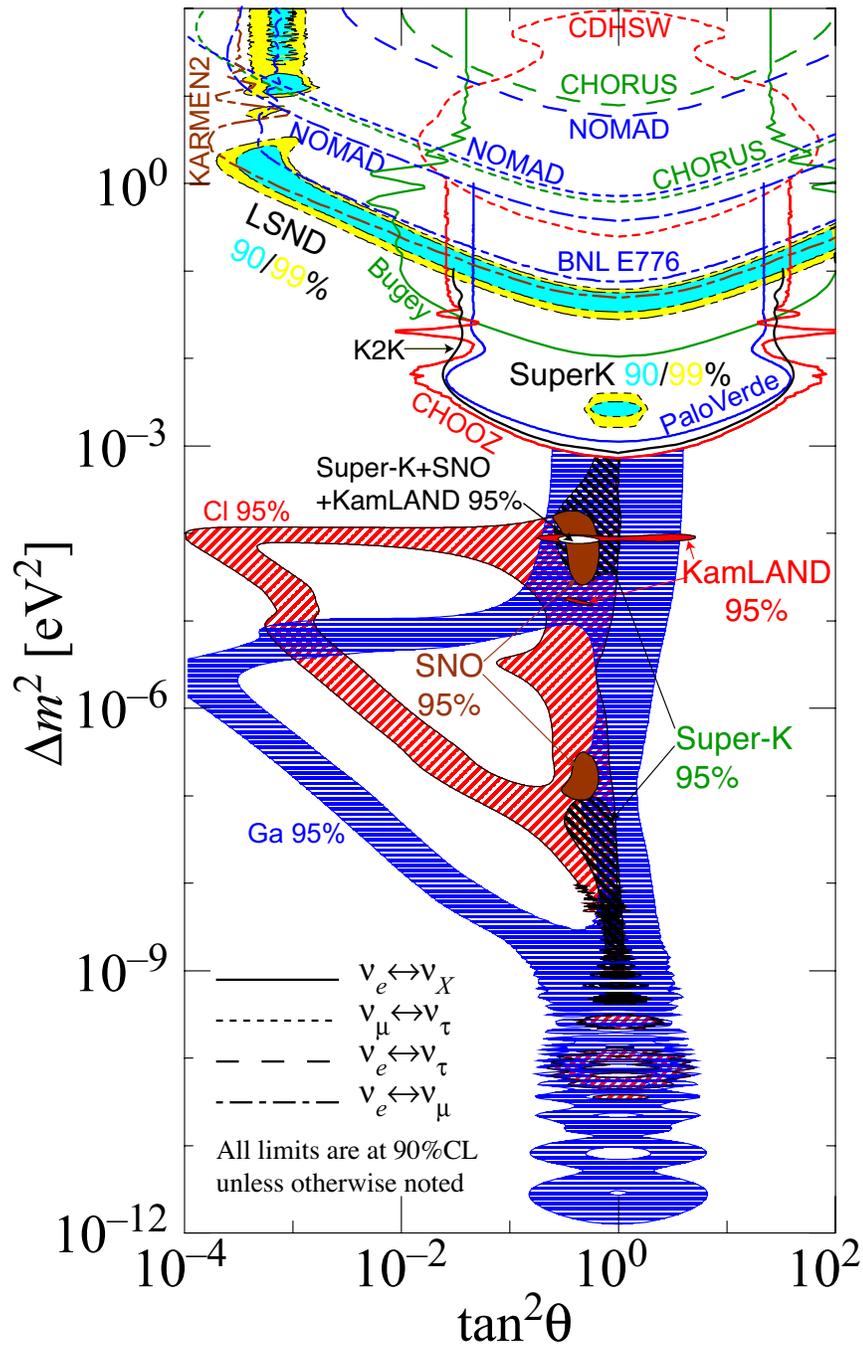


Figure 1.6: A compendium of neutrino oscillation results through 2005. From [49]. Also, [9]. Note that $\tan^2\theta$ is plotted rather than $\sin^2 2\theta$ to reveal Δm^2 sign information obtained in MSW-influenced experiments.

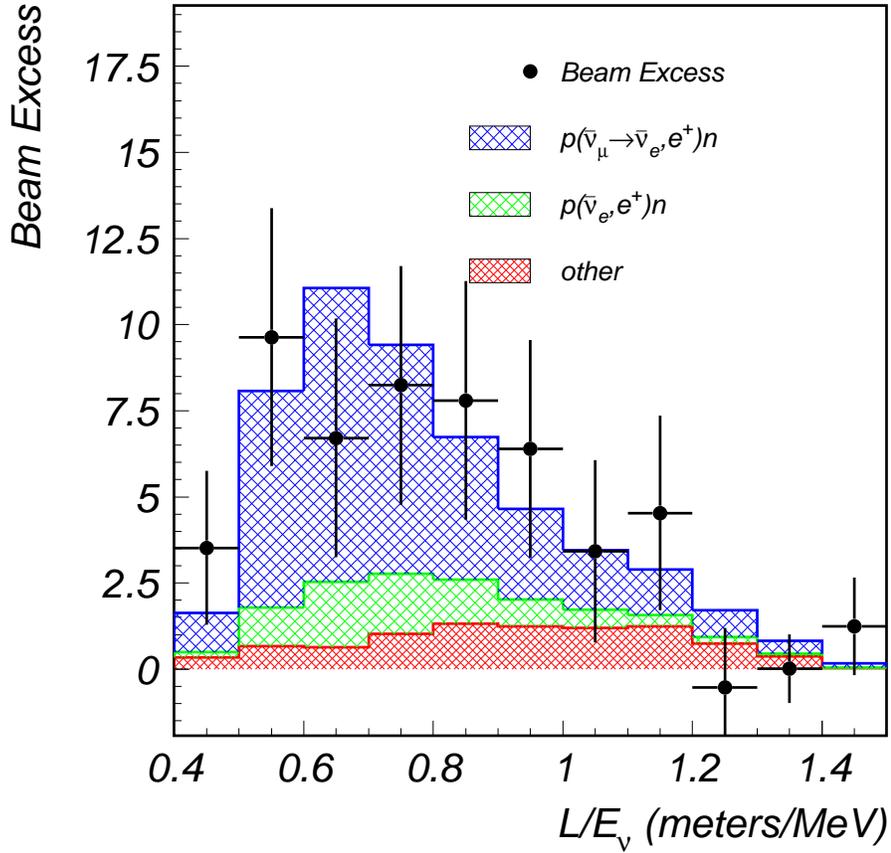


Figure 1.7: $\bar{\nu}_e$ candidate events in LSND. Beam-off backgrounds have been subtracted. The expected beam-related backgrounds are indicated by the bottom two histograms (red and green). An excess in the data (points) is seen and is well described by a $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation hypothesis (blue). Note that this plot was produced with a different set of cuts from those used in the primary analysis. From Ref. [50].

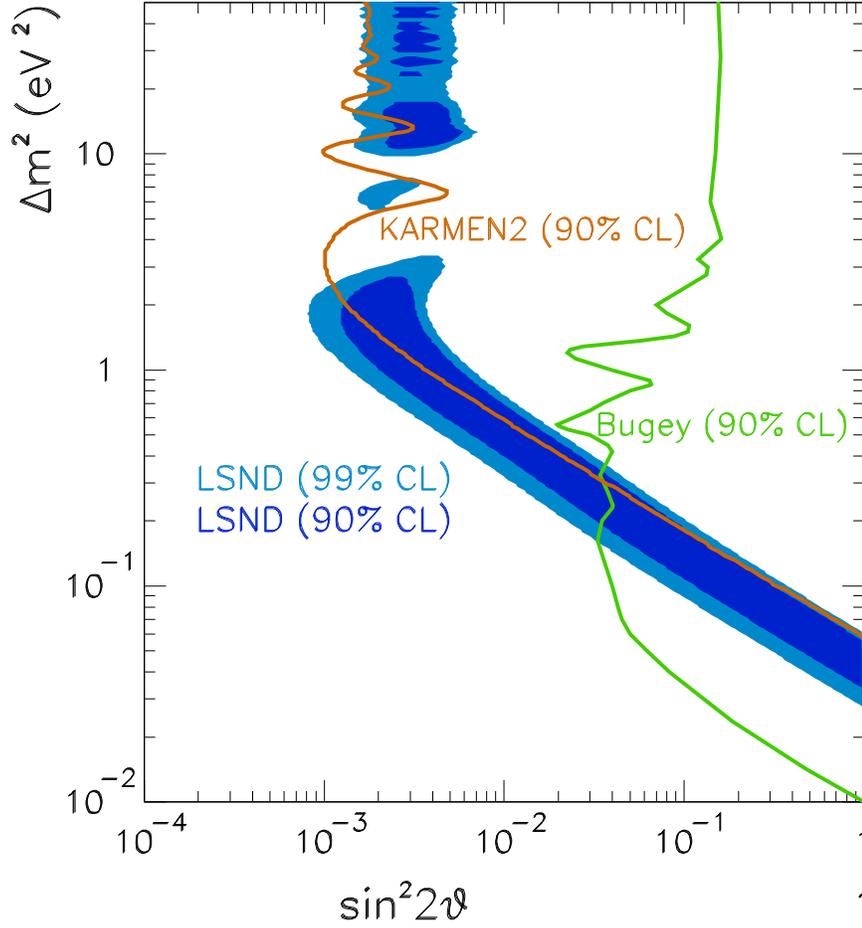


Figure 1.8: The LSND allowed region (90% and 99% C.L.) along with exclusion curves from KARMEN [53] and Bugey [54].

With three neutrino generations, the spectrum of mass-squared splittings must satisfy

$$\Delta m_{12}^2 + \Delta m_{23}^2 + \Delta m_{31}^2 = 0 . \quad (1.40)$$

The gap between $\Delta m_{\odot}^2 \approx 8 \times 10^{-5} \text{ eV}^2$ and $\Delta m_{\text{atm}}^2 \approx 3 \times 10^{-3} \text{ eV}^2$ implies that the third splitting must be approximately degenerate with the atmospheric value,

$$\Delta m_{\text{third}}^2 \approx \Delta m_{\text{atm}}^2 \approx 3 \times 10^{-3} \text{ eV}^2 , \quad (1.41)$$

yet $\Delta m_{\text{LSND}}^2 \gg \Delta m_{\text{atm}}^2$. That is, within a three neutrino framework, the LSND result is inconsistent with the numerous experiments of §1.3.4 which have measured Δm_{\odot}^2 and Δm_{atm}^2 . This inconsistency can also be demonstrated with a three-neutrino global analysis, as in Ref. [55].

A possible reconciliation is to introduce a fourth neutrino separated in mass-squared

from the others by $\sim \Delta m_{\text{LSND}}^2$. From measurements of the invisible partial width of Z^0 decays at LEP, we know that the number of light neutrinos that couple to the Z^0 is 2.9841 ± 0.0083 , consistent with $N_\nu = 3$ [56]. This fourth neutrino, then, must not take part in the weak interaction and is thus termed a *sterile* neutrino. Models that introduce one sterile neutrino ν_s leave some tension in the short-baseline data, but it has been shown that a second sterile neutrino can remedy this [57]. SNO and Super-K, by measuring the active neutrino flux with neutral current reactions and by looking for sterile-like MSW effects, provide limits on the amount of mixing that ν_s can have with the active sector [45] [58], but many models (especially those with multiple sterile states) are still allowed.

An alternate scenario is that the mass spectrum for neutrinos differs from that for antineutrinos. This requires a violation of *CPT* in the lepton sector; Refs. [59] and [60] put forth such models.

1.4 MiniBooNE

The LSND results still lack independent confirmation or refutation, and this is the *raison d'être* of the Booster Neutrino Experiment (MiniBooNE) at the Fermi National Accelerator Laboratory (FNAL) in Batavia, Illinois.¹⁰

MiniBooNE is a $\nu_\mu \rightarrow \nu_e$ appearance search that utilizes the FNAL Booster neutrino beam, a broadband source of ~ 1 GeV muon neutrinos or (if so configured) antineutrinos.¹¹ A 0.8 kton mineral oil Cherenkov detector sits 0.5 km downstream of the neutrino source. An excess of ν_e candidate events in the detector would be an indication of $\nu_\mu \rightarrow \nu_e$ flavor change. While MiniBooNE operates at neutrino energies $10 - 20\times$ higher than those in LSND, the neutrino baseline is larger by a corresponding amount ($16\times$), giving MiniBooNE and LSND sensitivity to comparable Δm^2 ranges.

The higher energies also imply a different set of event signatures. Charged current ν_μ interactions are not kinematically forbidden, and these account for the majority of events. Charged current ν_e events are sifted out by recognizing the Cherenkov patterns characteristic of e -induced electromagnetic showers. The neutrino beam has $\sim 0.6\%$ ν_e contamination which forms an irreducible background in the ν_e search. Additionally, neutral pions created in ν_μ interactions may decay to an electromagnetic final state ($\pi^0 \rightarrow \gamma\gamma$) which can mimic a ν_e event.

This thesis presents the first $\nu_\mu \rightarrow \nu_e$ oscillation search performed at MiniBooNE, using 1.7 million neutrino interactions collected from February 4, 2003, to October 30, 2005. An LSND-like $\nu_\mu \rightarrow \nu_e$ oscillation signal would appear as an excess of $\mathcal{O}(100)$ ν_e events after selection, alongside a comparable number of intrinsic ν_e and misidentified ν_μ interactions.

¹⁰“MiniBooNE” refers to the single-detector experiment currently running. A confirmation of the LSND signal would motivate the construction of a second, longer baseline detector and the birth of “BooNE”.

¹¹This beamline was constructed for MiniBooNE but it is proving to be more generally useful.

1.5 Layout of this document

The remainder of Part I outlines MiniBooNE's experimental setup. Part II discusses the $\nu_\mu \rightarrow \nu_e$ oscillation analysis, beginning with an overview (Chapter 5) and continuing with three chapters describing the Monte Carlo simulation chain. Chapter 9 discusses the event reconstruction algorithm, which provides event kinematics and particle identification information. Chapters 10 and 11 describe the selection and use of two tuning samples (ν_μ charged current quasi-elastic events and neutral current π^0 events). Chapter 12 presents the ν_e selection and signal extraction procedure. Chapter 13 discusses how we verify the readiness of the analysis, and Chapter 14 concludes the main text with the results of the oscillation search. Appendices follow.

Chapter 2

Booster neutrino beam

The neutrino source in MiniBooNE begins with 8 GeV kinetic energy protons from the FNAL Booster proton synchrotron. Protons are directed onto a 71.1 cm long, 0.48 cm radius beryllium target that sits inside a magnetic focusing horn. Positive hadrons produced in the p -Be collisions are focused by the horn into a region where their decays produce the desired neutrinos. The decay region terminates with a steel and concrete absorber. The neutrinos continue through the absorber and ~ 500 m of undisturbed earth before reaching the MiniBooNE detector. Figure 2.1 shows the location of the target hall and detector on the FNAL site, and Figure 2.2 gives the layout of the experiment, from target to detector.

2.1 Proton source

Protons from the Booster arrive at the beryllium target in $1.6 \mu\text{s}$ spills (a.k.a. batches, pulses) with $\sim 4 \times 10^{12}$ protons in each spill.^{1,2} The Booster cycle frequency (and, thus, the instantaneous spill rate) is 15 Hz, although spills to MiniBooNE are limited to a time-averaged rate of 5 Hz due to target heating and horn structural limits.

To prevent unacceptable levels of radioactivation within the Booster tunnel, local and total beam losses are monitored and controlled. These Booster losses were the typical proton rate limitation during the neutrino run, although a series of Booster improvements kept the average rate on a steady upward climb (Figure 2.3). The analysis in this thesis uses the complete neutrino data set, which corresponds to 5.58×10^{20} protons-on-target (POT).³

2.2 Target and horn

The 71.1 cm target, comprising seven cylindrical beryllium slugs, provides 1.9 inelastic interaction lengths of material for the incoming protons. The slugs sit along the central axis

¹The spills have 53 MHz fine structure which is ignored by the oscillation analysis but can be seen in the neutrino data.

²Beamline toroids provide the protons-per-pulse measurements with $<5\%$ uncertainty, negligible in this analysis.

³*cf.* the MiniBooNE goal: 1×10^{21} POT [62]

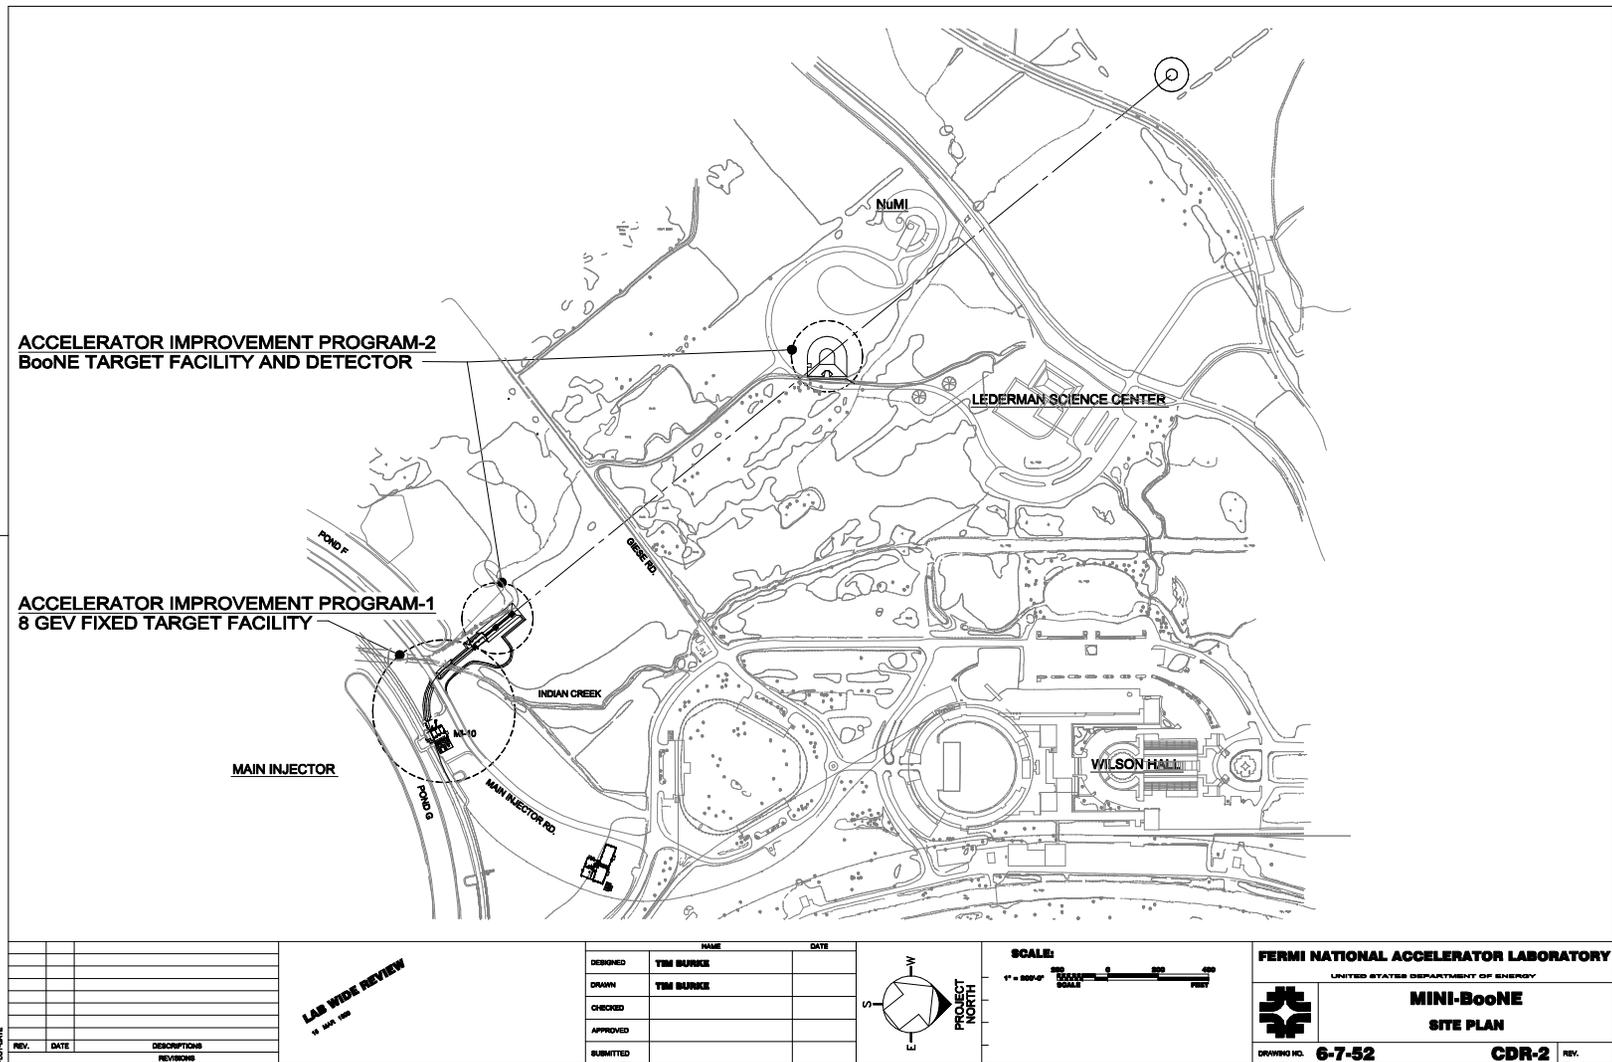


Figure 2.1: FNAL site plan, showing the locations of the 8 GeV target hall and the MiniBooNE detector (smaller dashed circles). The target-to-detector distance is 541 m. From Ref. [61].

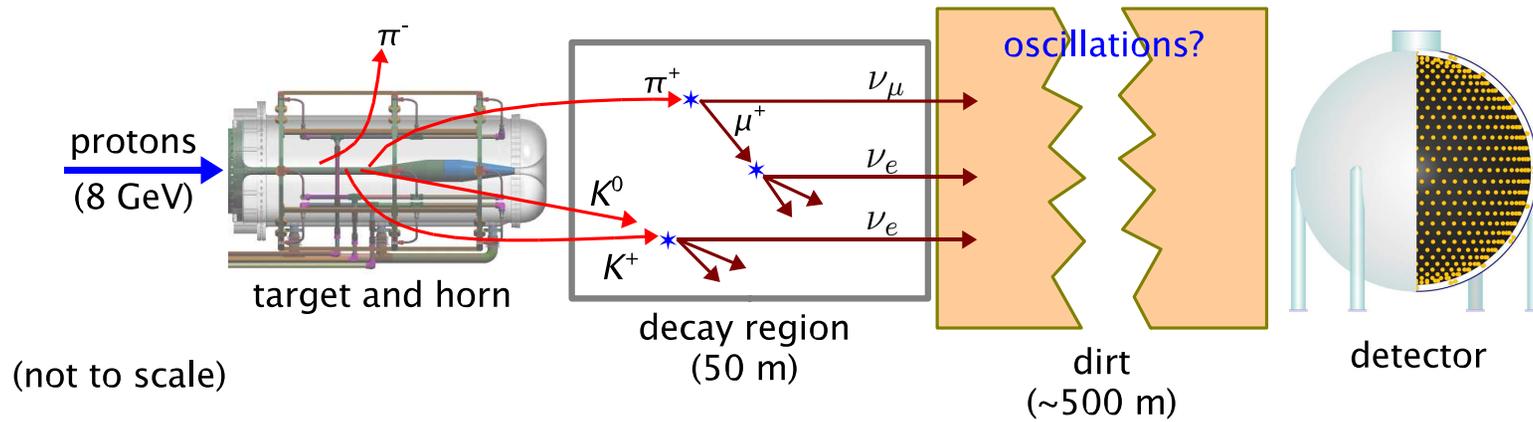


Figure 2.2: A cartoon of the experimental setup (not to scale). The $\pi^+ \rightarrow \mu^+ \nu_\mu$ decay chain is the primary source of neutrinos. The ν_e -producing decays indicated produce only 0.6% of the total flux at the detector, but this small ν_e contamination leads to the largest background in the $\nu_\mu \rightarrow \nu_e$ oscillation search.

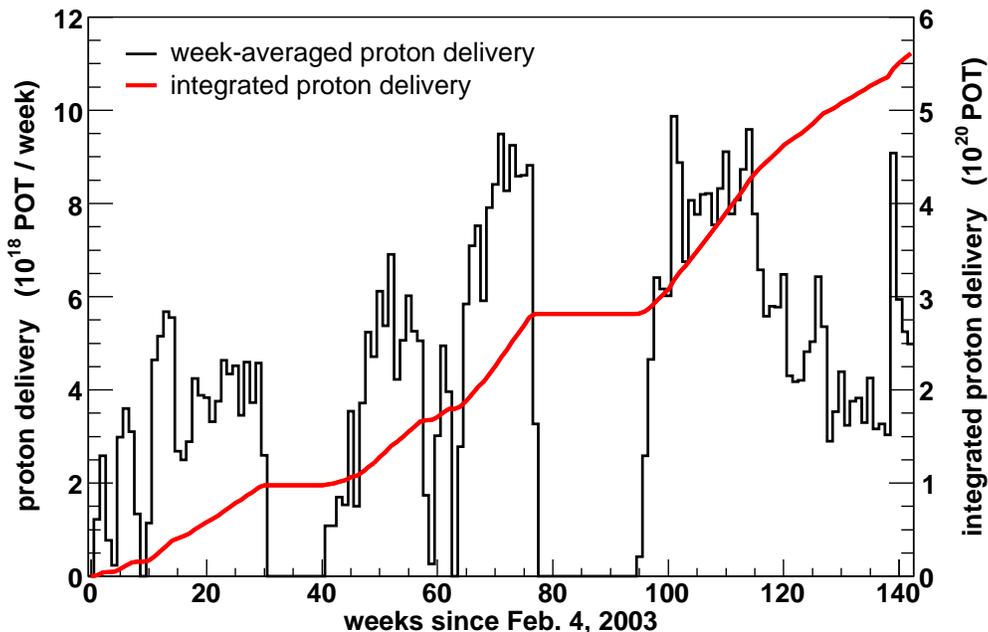


Figure 2.3: Protons to the MiniBooNE target. The black histogram shows the week-averaged proton rate. The drop near week 115 corresponds to the turn-on of the NuMI neutrino beam. The two large gaps are from accelerator shutdowns. The red curve shows the integrated proton delivery. Events that do not pass beam- or detector-related data integrity checks (2.5%, Appendix A) are not included.

of the horn and are cooled by a closed circulating air system which removes the ~ 600 W of beam-induced heating. Upstream, multiwires and beam profile monitors report the location and direction of the proton beam as it approaches the target. Beam spills with targeting information that is out of tolerance (Appendix A) are excluded from the analysis.

The magnetic focusing horn provides a toroidal magnetic field (peak: 1.5 T) that steers positive secondary hadrons toward the decay region.⁴ The horn is driven by a 143- μ s, 174-kA current pulse which peaks in time with the arrival of protons at the target. Current enters the horn along the cylindrical outer conductor (radius 30 cm, length 185 cm). The conductor folds inward at the downstream end to become the inner conductor, whose radius varies from 2.2 cm to 6.5 cm. The bulk of the horn, including the conductors and support hardware, is made of aluminum (alloy 6061 T6).

A closed water system cools the horn. Nozzles attached around the outer conductor spray water onto the inner conductor. Because the current pulses induce large vibrations in the horn, the nozzles are coupled to the outer conductor with bellows to provide mechanical isolation. A truss system surrounding the horn supports the water cooling system.

⁴MiniBooNE is now running with the horn polarity reversed, focusing negative secondaries to produce a primarily $\bar{\nu}_\mu$ beam.

Figure 2.4 shows a drawing of the horn and parts of the cooling system.

The target and horn sit within an iron shielding pile that controls the radiation level in the target hall. 2 m of shielding separates the target region from the room above. Additionally, a collimator just downstream of the horn absorbs secondary particles destined to miss the decay pipe. The collimator aperture grows from a radius of 30 cm (upstream) to 35.5 cm (downstream) and has a length of 214 cm, beginning 259 cm from the upstream end of the target.

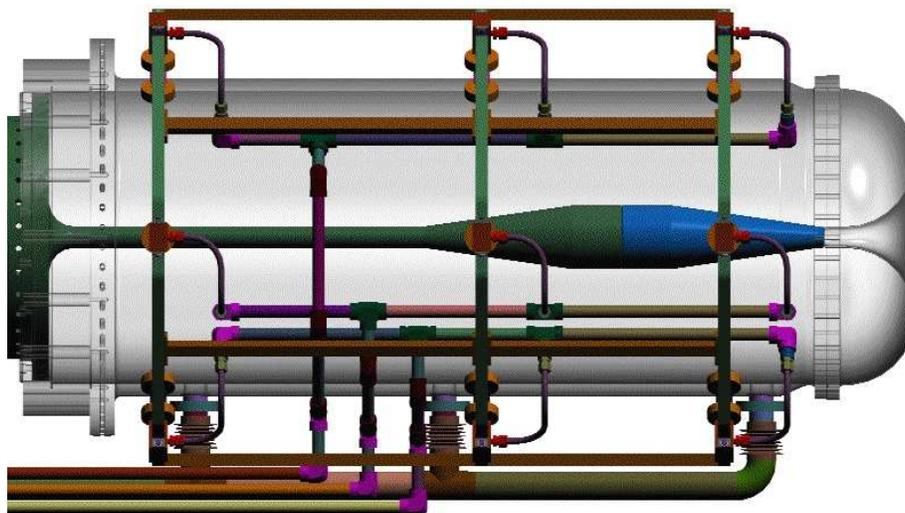


Figure 2.4: Side view of the horn. The outer conductor has been rendered partially transparent to reveal the inner conductor and parts of the water cooling system. From Ref. [61].

2.3 Decay region and flux

Secondary mesons that pass the collimator enter a ~ 45 m long, 90 cm radius decay pipe. The vast majority of secondaries are π^+ which decay to $\mu^+\nu_\mu$ to provide most of the neutrino flux reaching the detector. The decays of μ^+ , K^+ , and K_L^0 produce a small ν_e flux that amounts to a significant background in the oscillation search. Details of the flux composition are saved for Chapter 6. For now, we show in Figure 2.5 the predicted energy spectra of the neutrinos reaching the MiniBooNE detector. The relative integrated fluxes of $\nu_\mu : \bar{\nu}_\mu : \nu_e : \bar{\nu}_e$ are $1 : 6 \times 10^{-2} : 6 \times 10^{-3} : 6 \times 10^{-4}$.

The decay pipe terminates with a steel and concrete absorber located 50 m from the upstream end of the target. A second absorber can be inserted at 25 m, but this option was not exercised for the neutrino run. A muon spectrometer situated 7° off-axis views a

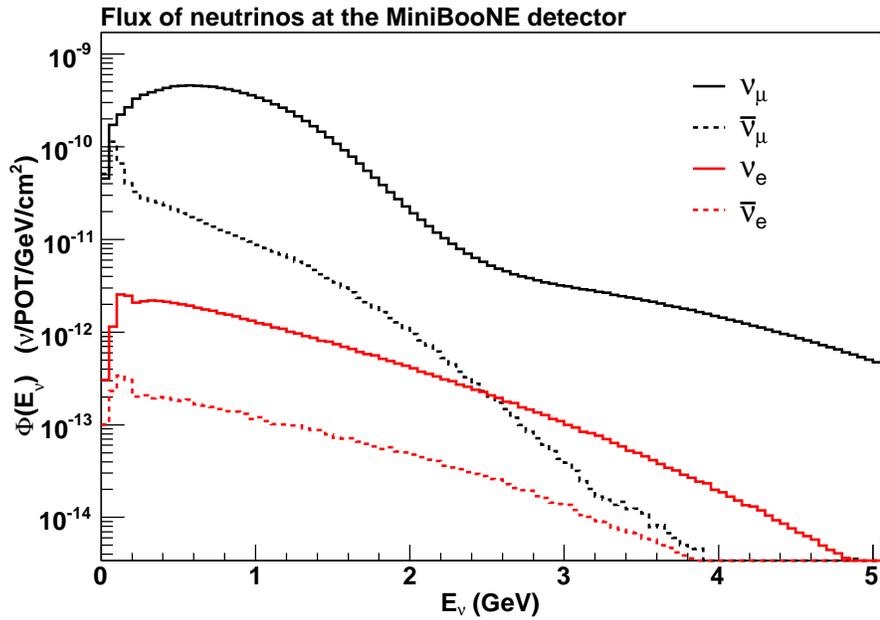


Figure 2.5: Neutrino fluxes at the MiniBooNE detector.

small portion of the decay region and can, in principle, provide a measure of K^+ production via the muons it sees. However, the analysis of this data was not complete in time for the oscillation search.

The neutrinos produced within the decay pipe continue through the absorber and head toward the MiniBooNE detector, the center of which lies 541 m downstream of the target and 2 m above the beam axis. We describe it next.

Chapter 3

MiniBooNE detector

3.1 Overview

The MiniBooNE detector consists of a 12.2 m diameter spherical carbon steel tank filled with 800 tons of mineral oil. Supported within the tank is an opaque spherical shell of diameter 11.5 m (the “optical barrier”) that divides the oil volume into a thin outer veto region and a large inner main region. The veto region is instrumented with 240 8-inch photomultiplier tubes (PMTs, tubes) mounted in pairs on the tank wall, and the main region is viewed by 1,280 8-inch PMTs lining the inner surface of the optical barrier. Charged particles created in neutrino interactions in the oil produce Cherenkov and scintillation light which is detected by the PMTs. The spatial and temporal distributions of the detected light provide all of MiniBooNE’s event identification information.

Figure 3.1 shows a cartoon of the primary detector components. Figure 3.2 shows the veto and main regions during detector construction. The veto region is painted white to maximize light collection from particles entering or leaving the detector, while the main region is painted black to minimize reflections which would confuse the Cherenkov light patterns.

The tank sits in a 15.2 m diameter cylindrical vault immediately below ground. A concrete floor separates this vault from the counting room above it. The readout electronics, data acquisition computers, HVAC system, and oil handling equipment sit in the upper room. The entire detector enclosure is shielded by a few meters of overburden. Figure 3.3 shows this layout.

3.2 PMT support structure

The optical barrier forms part of the PMT support structure (PSS), shown in Figure 3.4. The inner PMTs are attached with wire frames to the $\frac{1}{16}$ -inch thick optical barrier panels. Each panel supports two PMTs and is roughly 1 m wide by 0.5 m tall. The panels are held against a series of latitudinal pipes (“lat pipes”) by strips of aluminum sheet overlapping the panels and U-bolted to the lat pipes. The pipes themselves are supported by ~ 30 cm struts attached to bosses welded to the steel tank. The individual strut lengths differ to

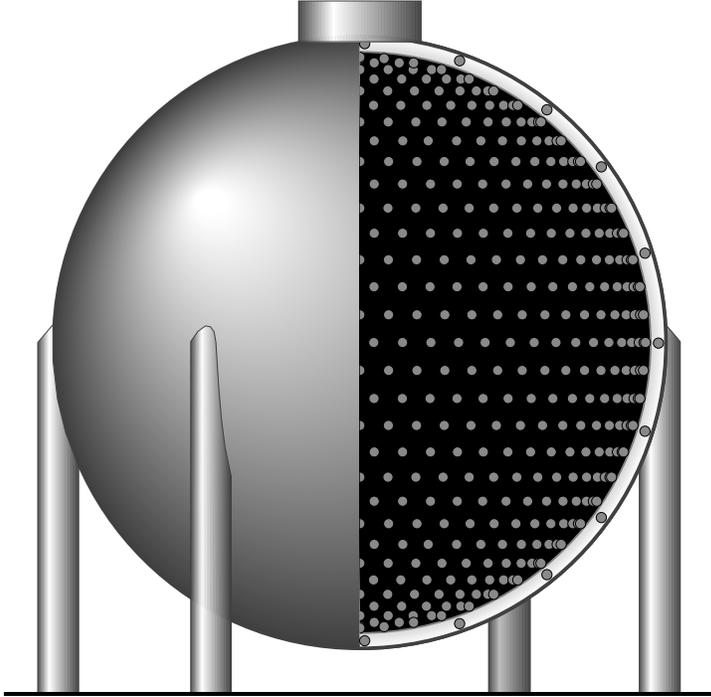


Figure 3.1: A cutaway drawing of the detector. The thin white veto region is seen at the edge of the sphere. The small circles represent PMTs (main) or PMT clusters (veto).

remove imperfections in the sphericity of the tank. Additionally, the strut-tank coupling allows for angular adjustment of the strut so that its end can be leveled at the desired height before the lat pipe is attached.

3.3 Oil

The detector is filled with ExxonMobil Marcol 7 mineral oil. The oil serves both as the target for neutrino interactions and as the light-producing medium for resulting charged particles. The properties of ten different oils from several vendors were measured and compared against pre-established criteria before Marcol 7 was chosen [64]. Table 3.1 lists some properties of MiniBooNE’s oil. Understanding the oil’s optical properties is intimately tied to the development of the detector simulation, and a detailed description is saved until Chapter 8.

density:	0.855 g/cm^3	($T=17.2 \text{ }^\circ\text{C}$)
thermal expansion:	$8.9 \times 10^{-4} \text{ K}^{-1}$	(volumetric, $T=14.2 \text{ }^\circ\text{C}$)
index of refraction:	1.484	($\lambda=400 \text{ nm}$, $T=17.5 \text{ }^\circ\text{C}$)
extinction length:	18 m	($\lambda=400 \text{ nm}$)
detected photons:	4.2 MeV^{-1}	(100 – 1000 MeV electrons)

Table 3.1: Properties of MiniBooNE’s Marcol 7 mineral oil.

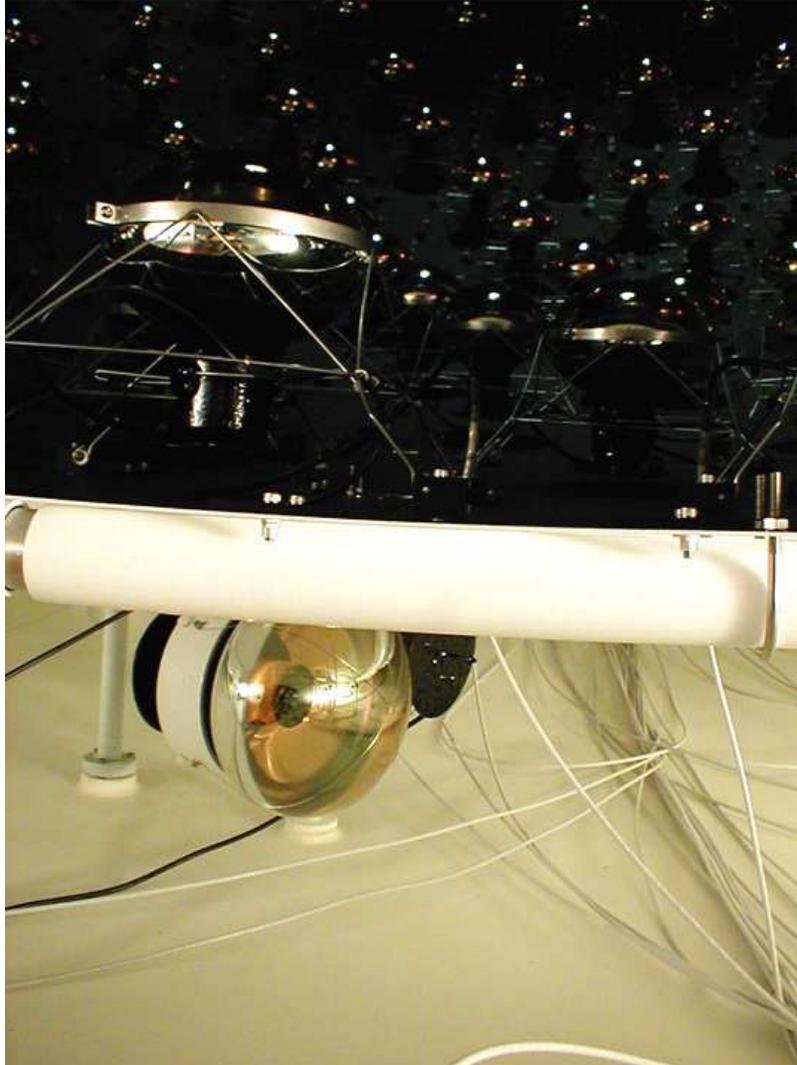


Figure 3.2: A view of the main and veto detector regions during construction.

The oil was delivered to FNAL via a series of rail cars over several months. A food-grade tanker truck shuttled the oil from the railhead to the detector. The oil is typically left static, although it can be circulated and chilled. (The chiller has never been used.) The system includes an overflow tank to handle thermal expansion of the oil. Also, nitrogen is bubbled into the tank at three locations (bottom, middle, top) to remove dissolved oxygen and to maintain a nitrogen atmosphere above the oil surface.

3.4 PMTs

Of the 1,520 PMTs in the detector, 1,198 are Hamamatsu R1408 PMTs and 322 are Hamamatsu R5912 PMTs. The R1408 PMTs were inherited from the LSND experiment while the R5912 PMTs were purchased new. Table 3.2 gives an overview of these tubes.

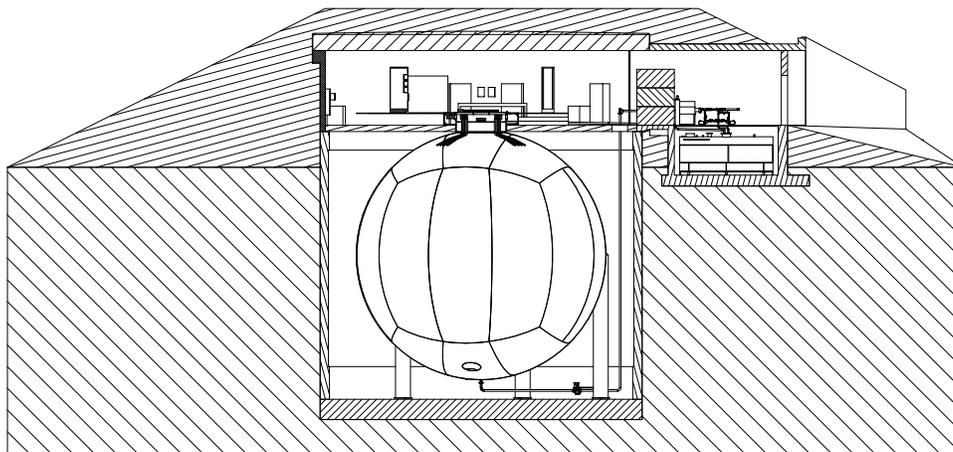


Figure 3.3: Detector hall. The personnel entrance (top right) leads to an anteroom containing oil equipment. The oil overflow tank sits beneath the floor of this room. Further in is the main electronics room. A floor hatch and curved stairway (not shown) provide access to the vault below. From Ref. [63].

	Hamamatsu R1408	Hamamatsu R5912
count:	1,198	322
operating voltage:	~ 2 kV	~ 2 kV
timing resolution:	1.7 ns	1.1 ns
1-PE charge resolution:	140%	50%
dark rate:	1.0 kHz	1.4 kHz

Table 3.2: The photomultiplier tubes in MiniBooNE.

Tube placement throughout the detector was chosen by tube type and measured dark rate [65]. The veto region was instrumented with the quietest of the R1408 PMTs. The older tubes have less chance of failure and the low dark rates keep the veto threshold down. The remaining tubes were distributed throughout the main region.

3.5 Electronics and DAQ

[*Note: This section, together with Chapter 4, may be skipped without loss of continuity.*]

PMT anode signals are carried to preamp boards in the electronics room via 100-ft coaxial cables which also supply each PMT with high voltage. The preamp (Figure 3.5) is based on an AD9617 op-amp and provides $\sim 20\times$ amplification. Each preamp board includes sixteen channels but is fed by a single high voltage supply. Therefore, the PMTs assigned to a given board have comparable operating voltages. Individual voltages are selected with resistive

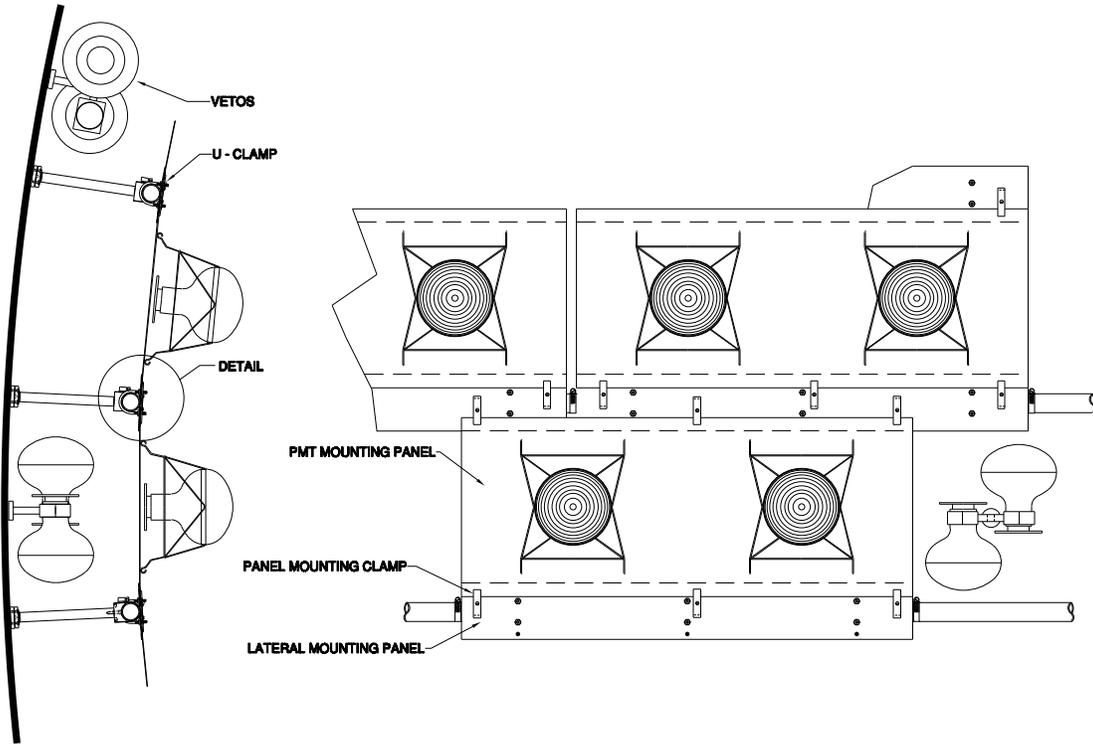


Figure 3.4: The PSS described in the text. Also shown are veto clusters attached to the outer tank. From Ref. [63].

voltage dividers on the boards.

The “QT boards” (Figure 3.6) convert the amplified anode pulses into more slowly varying signals that still contain the charge and time information desired. The anode pulse V_{pmt} feeds an integrating capacitor which bleeds off with a time constant $\tau \approx 1200$ ns. The voltage V_q across the capacitor is digitized every 100 ns, in step with a 10 MHz GPS clock. If an anode pulse is large enough to fire the on-board discriminator¹, a separate voltage V_t begins ramping linearly away from baseline. The ramp continues until two clockticks have passed, at which point it rapidly resets to baseline. V_t is also digitized every 100 ns. The data acquisition (DAQ) software looks at the stream of these V_q and V_t digitizations coming from each channel, and for each time ramp it detects (*i.e.*, for each firing of the discriminator), the DAQ records four V_q and V_t values – one before the discriminator fired and three after. This whole process is shown graphically in Figure 3.7.

The detector data stream, then, contains the following information from each hit:

- The PMT channel number.
- The clocktick, counted from the start of the event, that precedes the discriminator firing. For the hit in Figure 3.7, this would be i .

¹threshold: 0.1 – 0.2 photoelectrons

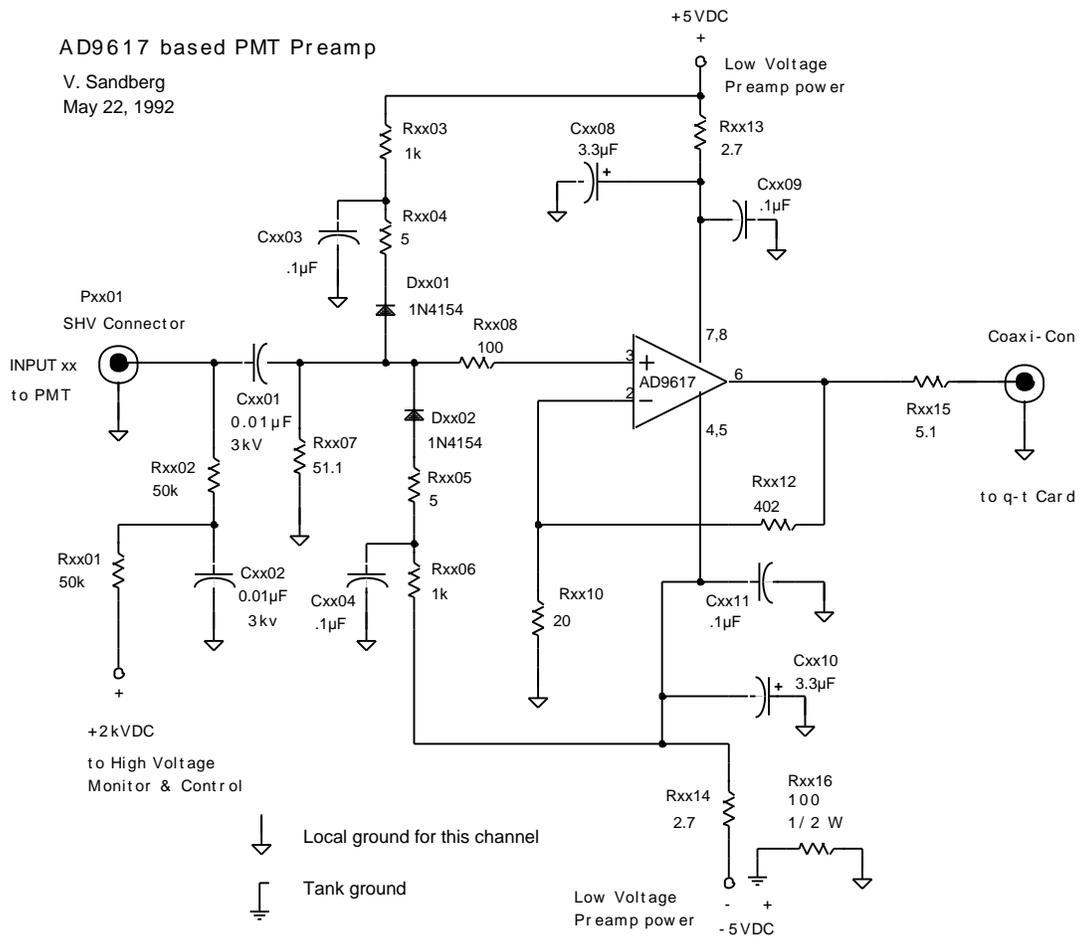


Figure 3.5: Preamp circuit for a single channel. Each preamp board has sixteen such channels. From Ref. [63].

- The four recorded V_q values in ADC counts (Q_ADC0 , Q_ADC1 , Q_ADC2 , Q_ADC3). Collectively, the “charge quad”.
- The four recorded V_t values in ADC counts (T_ADC0 , T_ADC1 , T_ADC2 , T_ADC3). Collectively, the “time quad”.

If a hit has more than ~ 20 photoelectrons, V_q becomes too large for the ADC. For these “saturated” hits, the DAQ writes out additional quads until V_q is back on-scale.² Chapter 4 describes how this data stream is used.

²More precisely: the DAQ keeps writing quads for a hit until it has written a completely on-scale quad.

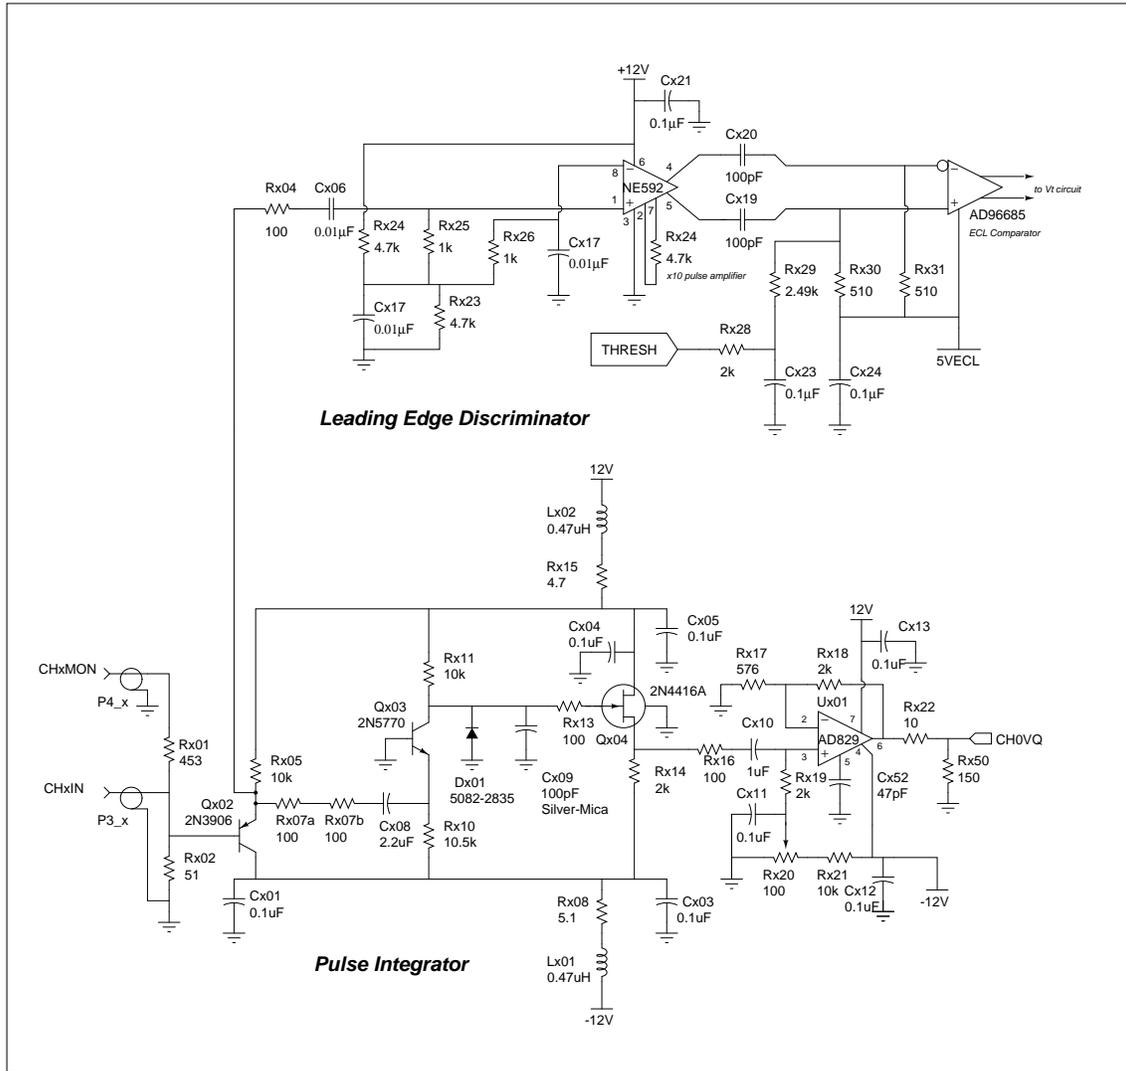


Figure 3.6: PMT front-end circuit. The top portion shows the pulse discriminator which feeds the V_t ramp circuit (not shown). The bottom portion shows the pulse integrator. From Ref. [63].

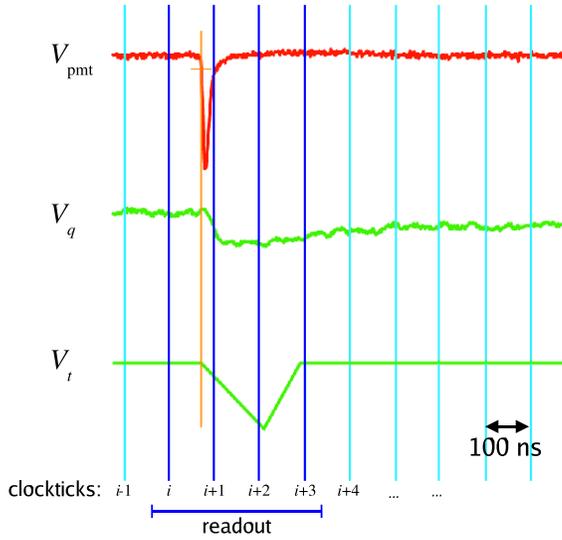


Figure 3.7: Pulses in the front-end electronics. V_{pmt} is the incoming anode signal. Its integral, convolved with an exponential decay, is V_q . The vertical orange line indicates when the discriminator fired. This firing induces the start of the time ramp V_t . V_t resets after two clockticks have passed. For this hit, the DAQ records the four V_q and V_t values digitized at clockticks i , $i + 1$, $i + 2$, and $i + 3$. (Adapted from Ref. [63].)

3.6 Calibration system

Laser flasks

Suspended within the oil volume are four 4-inch spherical light dispersion flasks filled with Ludox[®] colloidal silica [66] and fed by optical fibers from a diode laser.³ One flask is within a few centimeters of the tank center, and the others are distributed throughout the tank. The laser pulses sent to these flasks are used for detector calibration and monitoring. Figure 3.8 shows a flask prior to its installation.

A pulsed diode laser provides light bursts lasting ~ 100 ps each. The laser output follows an optical fiber into a mechanical switch box which selects the target flask. The switch box and laser driver are computer controlled. Under normal running conditions, only the detector-center flask is used, with pulses from a 397 nm laser head occurring at 3.33 Hz. (A 438 nm laser head is also available.)

Muon tracker and cubes

Immediately above the detector is a four-plane cosmic ray muon hodoscope, the “muon tracker”. The lower planes of the muon tracker rest directly on the tank’s top access hatch while the upper planes hang from the ceiling 1 m above the lower planes. Each upper plane

³Lasers and laser drivers are from PicoQuant GmbH.

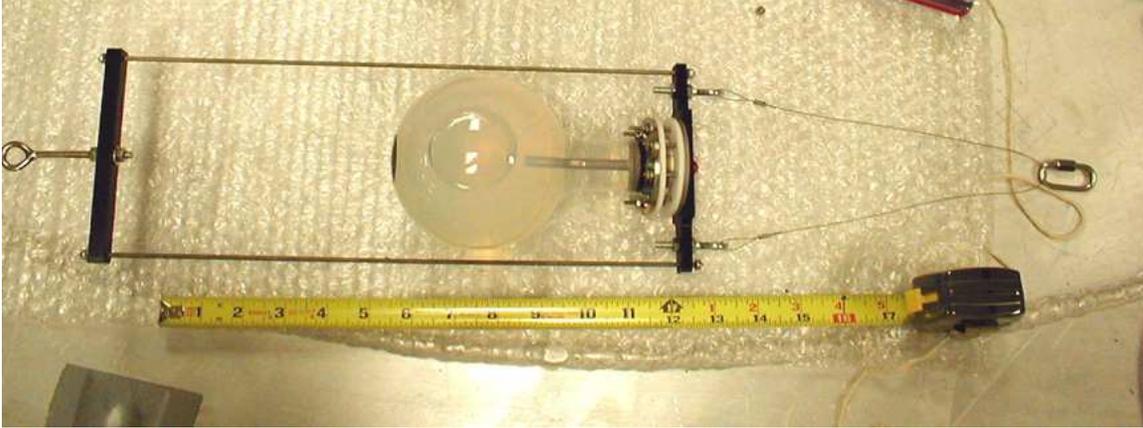


Figure 3.8: A laser flask on its side before installation. The thin metal tube running through the neck of the flask contains a light guide which extends just beyond the end of the tube. An optical fiber (not shown) is coupled to this light guide. The black spot on the bottom of the flask prevents undispersed laser light from hitting tubes directly below.

comprises 23 strips of Bicron BC-408 plastic scintillator. Each scintillator strip measures $10\text{ cm} \times 228\text{ cm} \times 1.59\text{ cm}$ and has an acrylic light guide and 2-in PMT attached at each end. The lower planes' strips measure $6\text{ cm} \times 170\text{ cm} \times 1.59\text{ cm}$ and are read out at one end only. The two upper planes are rotated 90° relative to one another to provide 2D positioning; likewise for the lower planes. Using the muon tracker, we obtain a sample of muons with known directions ($\sigma_\theta=2^\circ$) and positions in the main detector ($r_0=575\text{ cm}$; $\sigma_{x,y}\approx 10\text{ cm}$).

Suspended at depths ranging from 30 cm to 400 cm are seven small cubes of scintillator each housed in a black aluminum box and read out via an optical fiber coupled to a PMT in the electronics room. These cubes are 5 cm wide except for the deepest cube which is 7.6 cm wide. Of the two cubes at 30 cm, one sits beneath a 15 cm tall cylindrical “coffee can” light shield, resulting in 15 cm of visible muon pathlength. A few hundred hodoscope-tracked cosmic ray muons stop and decay in these cubes each month. These muons have known origins, directions, and (thanks to the cubes) pathlengths through the oil and are used to study energy reconstruction, scintillation light, and discriminator time slewing.

3.7 Trigger and readout

The DAQ records an event when one of several trigger conditions is met. Among these is the presence of a beam spill: any time a Booster batch is extracted and directed toward the MiniBooNE target (corresponding to the coincidence of FNAL ACNET signals 1D and 1F), a $19.2\text{ }\mu\text{s}$ window of detector activity is stored. Each event written to disk includes:

- global time information for synchronizing with external DAQ systems
- zero-suppressed PMT charge and time digitizations
- muon tracker and cube readouts

- information relating to the trigger itself (*e.g.*, the laser firing time for laser triggers)
- ACNET information (beam alignment and luminosity, horn and target conditions) when applicable.

A list of MiniBooNE triggers follows. Table 3.3 gives typical rates. Phrases in **bold** indicate the trigger names used in the table.

Beam

As mentioned, a proton batch destined for the **Booster neutrino beam** induces a trigger. The **NuMI neutrino beam** also triggers the DAQ. A beam trigger trumps all other triggers and disarms the laser driver.

Cosmic activity

Several triggers take advantage of PMT “sum cards” which provide a running tally of the number of PMTs that have fired in the last 200 ns. More specifically, the trigger receives a set of bits indicating whether certain hardware-set thresholds for main and veto PMT multiplicities have been reached. The thresholds are 10, 24, 60, 100, and 200 for main PMTs; 4 and 6 for veto PMTs.⁴ So, a simple **veto** activity trigger is constructed from the condition $N_{\text{veto}} \geq 6$, and a **main** activity trigger occurs whenever $N_{\text{main}} \geq 10$. A **Michel** trigger⁵ happens when all of the following are true: $N_{\text{veto}} < 6$, $N_{\text{main}} \geq 24$, and cosmic-muon-like activity ($N_{\text{veto}} \geq 6$ and $N_{\text{main}} \geq 100$) occurred sometime between 3 μs and 15 μs before now.

The rates are far too high for the DAQ to trigger on every instance of the above conditions, so only some of these events are actually taken. The scaling factors are listed in Table 3.3.

Calibration and monitoring

At a rate of 2.01 Hz, **random** detector activity is recorded. Other trigger conditions: the **laser** fires; the muon **tracker** has four-plane coincident activity; or, a **cube** shows activity while $N_{\text{main}} \geq 100$. An additional **beam-laser** trigger occurs when there is a Booster extraction for antiproton stacking (ACNET signals 14 and 1F). This provides laser events during Booster activity which can be compared to the usual laser events.

Other

The **big ν** trigger asks if activity is consistent with a high-energy neutrino event despite the lack of beam ($N_{\text{veto}} < 6$ and $N_{\text{main}} \geq 200$). The **follower** triggers write blocks of activity

⁴Typical multiplicities are discussed in Chapter 5. Also, N_{main} and N_{veto} refer here to trigger-level sums in this section only. We use N_{main} and N_{veto} elsewhere to represent “subevent” multiplicities, which we define in Chapter 5.

⁵Events from this trigger are not used in the analysis. Michel electrons (stopped μ -decay electrons) can be collected in a more controlled fashion from beam and random triggers.

after beam triggers when certain multiplicity conditions are met in both the beam window and in the period after. The **supernova** trigger looks for contained activity ($N_{\text{veto}} \geq 6$ and $N_{\text{main}} \geq 60$) when no cosmic activity ($N_{\text{veto}} \geq 6$ and $N_{\text{main}} \geq 100$) has occurred in the past $15 \mu\text{s}$. None of these miscellaneous triggers plays a role in this analysis.

trigger	typical rate (Hz)	scaling (if not 1/1)
Booster ν beam	0 to 5	
NuMI ν beam	0.4	
main	0.3	1/90000
veto	0.5	1/5001
Michel	1.2	1/600
random	2.0	
laser	3.3	
cube	1.2	
tracker	0.7	scaling set in hardware
beam-laser	0.8	
big ν	0.8	
follower	1.0	
supernova	10.2	
total	~ 25	

Table 3.3: MiniBooNE triggers. Typical rates (post-scaling) are shown in the middle column. The right column indicates the fraction of triggers actually recorded.

Chapter 4

PMT charge and time

[*Note: This technical chapter may be skipped without loss of continuity.*]

As described in §3.5, each recorded event includes charge and time “quads” containing PMT hit information. We need to extract two things from each hit’s quads:

1. the time of the hit (relative to the times of the other hits, at least); and,
2. the number of photoelectrons (PE) in the hit.

4.1 Determining t

We calculate a hit’s time t with

$$t = T_{\text{raw}} + (100 \text{ ns}) i_{\text{clock}} + \Delta t_{\text{offset}} + \Delta t_{\text{slew}}(Q_{\text{raw}}) . \quad (4.1)$$

T_{raw} ranges from 0 ns to 100 ns and measures the time of the hit relative to the 10 MHz clocktick that immediately preceded it. (See Figure 4.1.) It is calculated from the $\{\text{T_ADC}n\}$ (the time quad) as described below. The second term adds the time of the preceding clocktick itself as measured from the start of the event. This term is clearly a multiple of 100 ns, and it contributes the coarse timing that complements the fine timing provided by T_{raw} . The third term, Δt_{offset} , is a channel-dependent calibration constant that removes channel-to-channel timing differences arising from, among other things, differing cable lengths and dynode structures. The final term, $\Delta t_{\text{slew}}(Q_{\text{raw}})$, accounts for the charge-dependent time slewing caused by discriminating a pulse with finite rise-time.

If we know the slope of the time ramp, we can calculate T_{raw} with

$$T_{\text{raw}} = (200 \text{ ns}) - \frac{\text{T_ADC0} - \text{T_ADC2}}{m} , \quad (4.2)$$

where m is the magnitude of the ramp’s slope in (ADC counts)/ns. (A typical slope is 1.2 ADC counts/ns.) T_{raw} is actually over-specified by the $\{\text{T_ADC}n\}$, and we could have

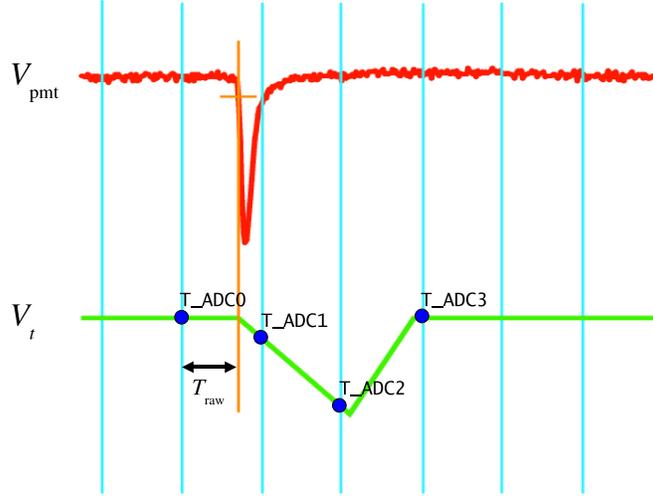


Figure 4.1: The time quad and T_{raw} . V_{pmt} is the hit's anode signal. The thin vertical line indicates when the discriminator fired. The $\{T_{\text{ADC}n}\}$ values recorded for this hit are identified by the solid dots. T_{raw} measures the time between the discriminator firing and the preceding clocktick. It is determined from the $\{T_{\text{ADC}n}\}$ and the (known) slope of the time ramp.

equally well used

$$T'_{\text{raw}} = (100 \text{ ns}) - \frac{T_{\text{ADC0}} - T_{\text{ADC1}}}{m} . \quad (4.3)$$

We might consider reporting T_{raw} as, say, the average of these two expressions. However, the start of the time ramp is delayed by 12.5 ns due to the finite response time of the QT board electronics, and T_{ADC1} consequently reports a baseline value for 12.5% of hits. For these hits, only Eq. (4.2) is meaningful. Thus, we use the average of Eq. (4.2) and Eq. (4.3) only when the following condition is met:

$$T_{\text{ADC0}} - T_{\text{ADC1}} > 20 \text{ ADC counts} . \quad (4.4)$$

Otherwise, we use Eq. (4.2) alone.

As we shall see, the charge calculation expects T_{raw} to fall in the range 0 ns to 100 ns (although some spillover is fine.) The time ramp delay causes the above calculation to produce T_{raw} 's in the range ~ 13 ns to ~ 113 ns, so our final expression includes a fixed shift Δt_{nudge} to remedy this:

$$T_{\text{raw}} = \begin{cases} \frac{1}{2} \left[300 \text{ ns} - \frac{1}{m} (2 T_{\text{ADC0}} - T_{\text{ADC1}} - T_{\text{ADC2}}) \right] + \Delta t_{\text{nudge}} & \text{if Eq. (4.4) true} \\ 200 \text{ ns} - \frac{1}{m} (T_{\text{ADC0}} - T_{\text{ADC2}}) + \Delta t_{\text{nudge}} & \text{if Eq. (4.4) false} \end{cases} \quad (4.5)$$

with $\Delta t_{\text{nudge}} = -12.5$ ns and with the channel-dependent slopes m determined beforehand (§4.3).

4.2 Determining q

We would like to know how many photoelectrons were in each hit. However, our tubes have a broad charge response (Figure 4.2), and we cannot determine that a given hit definitely had, say, 3 photoelectrons. We estimate the number of photoelectrons with

$$q \equiv \frac{Q_{\text{raw}}}{\langle Q_{\text{raw}}^{1\text{PE}} \rangle} \equiv \frac{Q_{\text{raw}}}{g}, \quad (4.6)$$

where Q_{raw} is any quantity proportional to anode charge and $g \equiv \langle Q_{\text{raw}}^{1\text{PE}} \rangle$ is the mean value of this quantity for single-PE hits. g is a channel-dependent calibration constant determined from low-intensity laser flash events (§4.3).

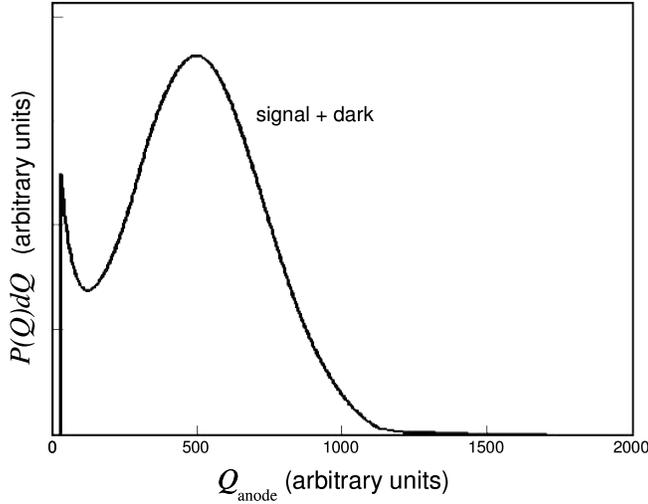


Figure 4.2: Single-photoelectron charge response of Hamamatsu R5912 PMTs. Adapted from Ref. [67].

4.2.1 The shape of V_q

Example anode pulses V_{pmt} are shown in Figure 4.3 for both tube types. The pulses exhibit overshoot/ringing features which impact the integrated pulse V_q on relevant time scales. (Recall the 100 ns digitization period.) Given V_{pmt} , one can compute the expected V_q with

$$V_q(t) \propto \int_{-\infty}^t V_{\text{pmt}}(t') e^{-\frac{(t-t')}{\tau}} dt'. \quad (4.7)$$

Figure 4.4 shows V_q so obtained for our two traces (taking $\tau = 1200$ ns, see §B.1). Figure 4.5 shows V_q as measured on the QT board. We use a reference shape V_q^{ref} like these to turn the $\{\text{Q_ADC}_n\}$ digitizations of V_q into a hit's Q_{raw} value.

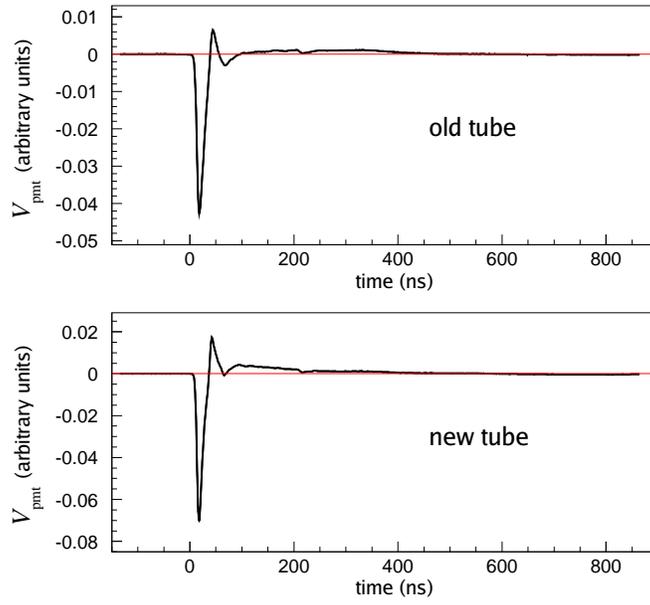


Figure 4.3: V_{pmt} for our tubes. The traces are average PMT signals from channels 1154 (top) and 1153 (bottom) as measured from the monitor port of their QT board with an oscilloscope. Channel 1154 (1153) is attached to an old (new) tube.

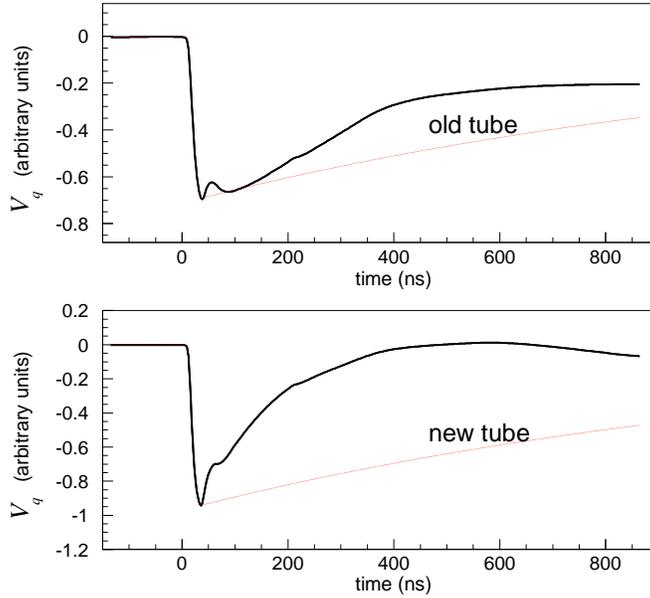


Figure 4.4: V_q as derived from V_{pmt} . The two thick traces are the results of Eq. (4.7) applied to the V_{pmt} signals shown in Figure 4.3. The top (bottom) plot is from channel 1154 (1153), an old (new) tube. The thin trace superimposed on each plot shows an idealized V_q created by artificially removing the ringing from the input signal V_{pmt} .

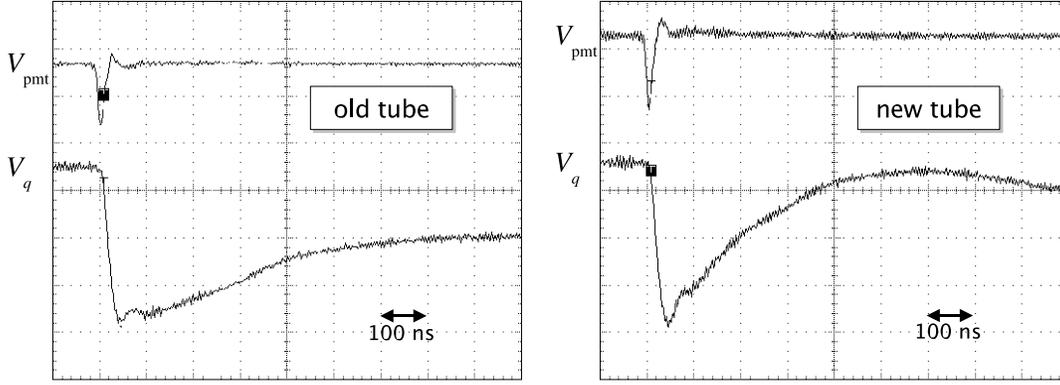


Figure 4.5: V_q measured directly. The signal V_{pmt} from the monitor port of a QT board was routed into a separate benchtop QT board. Each V_q trace shown is the average signal just upstream of the charge ADC on the benchtop board. The left (right) plot shows channel 1154 (1153), an old (new) tube. These V_q traces largely agree with the curves obtained using Eq. (4.7) and shown in Figure 4.4.

4.2.2 Calculating Q_{raw}

To good approximation, the shape of V_{pmt} (and, consequently, V_q) is independent of charge. That is, an N -PE signal looks just like a 1-PE signal except it is N times larger. Calculating Q_{raw} for each hit, then, amounts to determining the amplitude of the V_q signal that was digitized. We know which four points along V_q correspond to the four digitizations $\{\text{Q_ADC}n\}$ because we know T_{raw} . We fit the hit's three baseline-subtracted charge values

$$\text{Q_ADC}n' = \text{Q_ADC}n - \text{Q_ADC}0 \quad n = \{1, 2, 3\} \quad (4.8)$$

to αV_q^{ref} , where α is the amplitude we are after and V_q^{ref} is a fixed-amplitude reference curve. This is shown schematically in Figure 4.6.

The best-fit normalization is proportional to the anode charge, so it can serve as Q_{raw} :

$$Q_{\text{raw}} = \alpha_{\text{fit}} . \quad (4.9)$$

Giving the three $\{\text{Q_ADC}n'\}$ equal weight¹, we can write the result of the fit as²:

$$Q_{\text{raw}} = \frac{\sum_{n=1}^3 \text{Q_ADC}n' V_q^{\text{ref}}(n; T_{\text{raw}})}{\sum_{n=1}^3 [V_q^{\text{ref}}(n; T_{\text{raw}})]^2} , \quad (4.10)$$

where $V_q^{\text{ref}}(n; T_{\text{raw}})$ is the value of V_q^{ref} corresponding to $\text{Q_ADC}n'$ given the hit's T_{raw} . In

¹One could imagine doing otherwise.

²This expression is modified later.

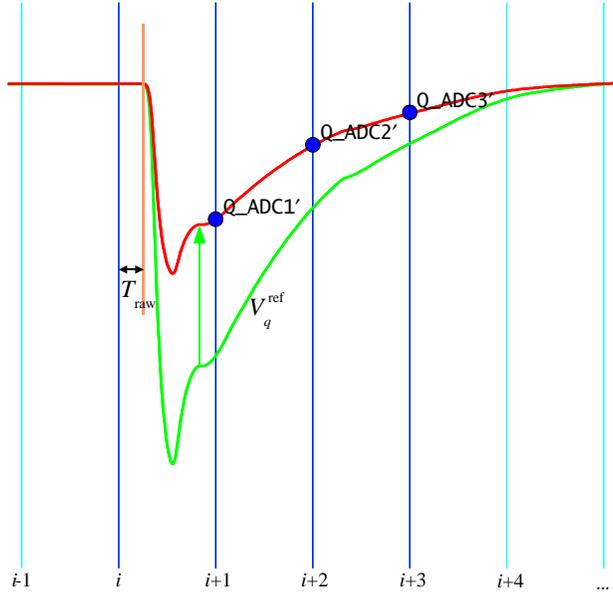


Figure 4.6: Fitting V_q^{ref} to the $\{\text{Q_ADC}n'\}$. Having calculated T_{raw} from the $\{\text{T_ADC}n\}$, we know how the reference curve V_q^{ref} should line up in time with the digitizations (solid dots). The curve running through the $\{\text{Q_ADC}n'\}$ represents αV_q^{ref} , with α determined from a fit to the points. The quantity α serves as Q_{raw} . For this hit, $\alpha \approx 0.5$.

practice, the reference curve V_q^{ref} is stored as a lookup table indexed by n and T_{raw} binned in 0.5 ns intervals.

4.2.3 Establishing V_q^{ref}

While we could establish V_q^{ref} using oscilloscope traces, as in §4.2.1, it is more convenient to extract the reference curve from the data stream itself, which we do as follows.

Low-intensity laser events provide low-charge hits at all possible values of T_{raw} . That is, over the course of a laser run, a channel will sample all parts of the V_q curve, reporting these samples as the $\{\text{Q_ADC}n'\}$. Though each hit has a different (unknown) charge Q , the *mean* charge $\langle Q \rangle$ is constant. Further, the mean charge is independent of T_{raw} – *i.e.*, independent of which part of V_q is being sampled. In each n and T_{raw} bin, we calculate the mean $\{\text{Q_ADC}n'\}$ seen over the run. This forms a full V_q curve whose normalization is representative of a hit with charge $\langle Q \rangle$. Declaring this (arbitrary) normalization fixed, we now possess a reference curve V_q^{ref} . We create two such curves, one for old tubes and one for new. Figure 4.7 shows the result. While the fine structure of the curves is washed out by the multi-PMT averaging, the important $\mathcal{O}(100 \text{ ns})$ structure agrees with what we see in the scope traces.

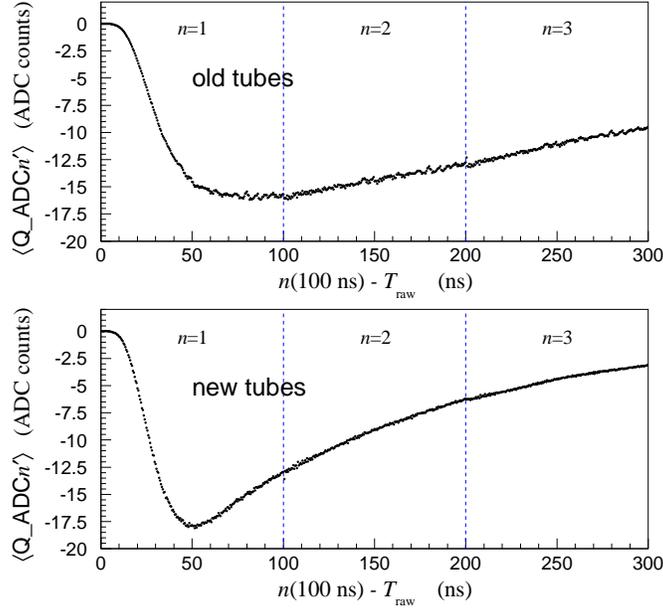


Figure 4.7: Forming V_q^{ref} from the data. For each of the 200 T_{raw} bins, we have plotted the mean baseline-subtracted quad value $\langle Q_ADCn' \rangle$ ($n = \{1, 2, 3\}$), where the mean is determined using all prompt hits in a low-intensity laser run. $\langle Q_ADCn' \rangle$ for a given T_{raw} bin is plotted with abscissa $n(100 \text{ ns}) - T_{\text{raw}}$, where T_{raw} is the low-edge of the bin, so that the continuous form of V_q is manifest. The vertical dashed lines separate the contributions from the three n values. All old (new) tubes contribute to the top (bottom) curve, excepting known bad tubes which are excluded.

4.2.4 Completing the picture

We have ignored “saturated” hits up to now. Recall that a saturated hit is one whose V_q is too large for the ADC. The DAQ recovers from this by writing additional quads for the hit until V_q is back on scale. If our reference curve extends out far enough, we can simply fit these “ $n > 3$ ” quad values.³ We obtain the extended reference curves using a special DAQ run (see §4.3.3), and we handle saturated hits by generalizing Eq. (4.10) to

$$Q_{\text{raw}} = \frac{\sum_{\{\text{all usable } n\}} Q_ADCn' V_q^{\text{ref}}(n; T_{\text{raw}})}{\sum_{\{\text{all usable } n\}} \left[V_q^{\text{ref}}(n; T_{\text{raw}}) \right]^2}, \quad (4.11)$$

where any usable Q_ADCn' is included in the fit. §B.2 outlines what makes a Q_ADCn' usable.

To summarize the entire charge and time calculation sequence for a hit:

1. Calculate T_{raw} from the $\{T_ADCn\}$ using Eq. (4.5).

³We introduce here the convention of using Q_ADC4 , Q_ADC5 , etc., to refer to a follower quad’s Q_ADC0 , Q_ADC1 , etc.

2. Given T_{raw} , calculate Q_{raw} using Eq. (4.11). That is, calculate Q_{raw} by fitting the reference curve to the $\{\text{Q_ADC}n'\}$.
3. Given Q_{raw} , report t using Eq. (4.1). (Recall that the slew correction depends on Q_{raw} .)
4. Given Q_{raw} , report q using Eq. (4.6).

The steps are slightly different for those channels recording diagnostic, non-PMT pulses (*e.g.*, trigger signals), as the above reference curves are not appropriate for these. The charge q reported for a non-PMT hit is just $\text{Q_ADC}2'$, and no slew correction is applied when calculating t .

4.2.5 Verification

We can examine whether the fit procedure actually works – that is, whether the $\{\text{Q_ADC}n\}$ behave as expected on a hit-by-hit and channel-by-channel basis. Figure 4.8 shows four typical Q_{raw} fits, and the $\{\text{Q_ADC}n\}$ indeed track the shape of the reference curve within a few ADC counts at all n . We can better quantify this agreement by looking at the fit residuals

$$\Delta_n \equiv \text{Q_ADC}n' - Q_{\text{raw}} V_q^{\text{ref}}(n; T_{\text{raw}}) , \quad (4.12)$$

where a residual Δ_n exists for each $\text{Q_ADC}n'$ used in the fit. For non-saturated hits, we have only $n = 2, 3$. Figure 4.9 shows the distributions of Δ_2 and Δ_3 for hits in a low-intensity laser run. The spiked structure is due to the intrinsic discreteness of the $\{\text{Q_ADC}n\}$. Of particular note are the widths of the distributions: $\text{RMS} \lesssim 0.5$ ADC counts. Given that a typical single-PE hit has $|\text{Q_ADC}n'| \sim 10$ ADC counts (for $n = 2, 3$), we can say that hit-to-hit and channel-to-channel variations in V_q are contributing $\lesssim 5\%$ smearing to our low charge calculations. The intrinsic PMT charge resolution is much larger, $\mathcal{O}(100\%)$.

4.3 Obtaining calibration constants

Five calibration tables are required:

- the time ramp slopes m ,
- the timing offsets Δt_{offset} ,
- the extended charge reference curves V_q^{ref} ,
- the charge scale factors g ,
- the slew corrections $\Delta t_{\text{slew}}(Q_{\text{raw}})$.

The charge and slew curves (one version for old PMTs, one for new) are extremely stable and, while checked periodically, are established only once for the MiniBooNE run. The

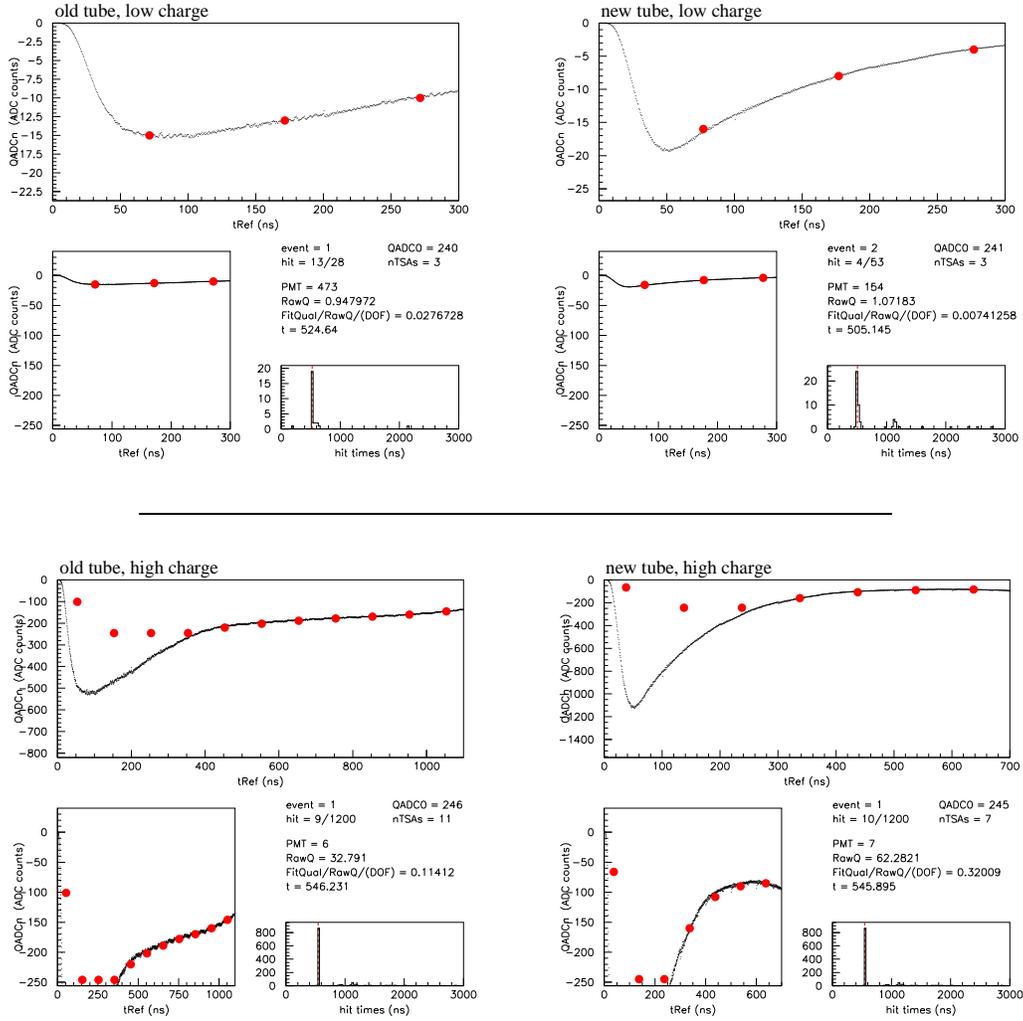


Figure 4.8: Four sample hits viewed with the “hit display”. The top panels show two low-charge hits, one from an old tube and one from a new tube. The bottom panels show two high-charge hits. In each panel, the large solid dots show the $Q_{\text{ADC}n}$ for the hit. The trace shows the result of the fit, *i.e.*, $Q_{\text{raw}} V_q^{\text{ref}}$. This information is shown twice in each panel, once with a variable vertical scale (top sub-panel) and once with a fixed vertical scale (lower left sub-panel). The fitted curve tracks the digitizations well. Note the ignored saturated digitizations in the bottom examples.

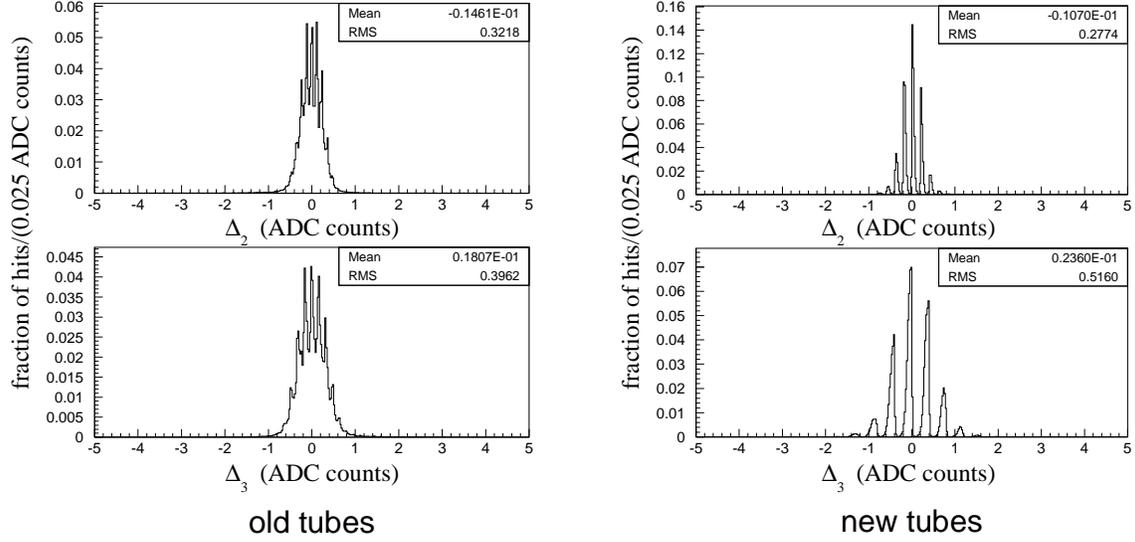


Figure 4.9: Q_{raw} fit residuals for low-intensity laser hits. The residuals for $Q_{\text{ADC2}'}$ and $Q_{\text{ADC3}'}$ are shown, with old (new) tubes on the left (right). The spiked structure is due to the intrinsic discreteness of the $\{Q_{\text{ADC}n}\}$.

others are updated every four days by an automated system which measures the calibration constants tube-by-tube while also performing many channel health checks.⁴

4.3.1 The slope m

The time ramp slope is measured by taking the average of $T_{\text{ADC2}} - T_{\text{ADC1}}$ over $\mathcal{O}(100\text{k})$ hits from laser triggers. Many of the hits are cosmic ray induced, but any hit works for determining m . A hit is only included in the average if $T_{\text{ADC1}} - T_{\text{ADC0}} > 20$ ADC counts. This is necessary due to the 12.5 ns time ramp delay mentioned previously.

4.3.2 The timing offset Δt_{offset}

Laser triggers are again used, but cuts are applied to remove events contaminated with cosmic activity.⁵ We begin by calculating each hit's time t using Eq. (4.1), except that (a) the Δt_{offset} term is left off, (b) the laser firing time (recorded separately by the DAQ) is subtracted, and (c) the photon time-of-flight from the laser flask to the PMT is subtracted (using $c/n = 19.48$ cm/ns; see §8.3.2. The flask used for calibration is 4 cm away from tank center.) An example t distribution is shown in Figure 4.10. The peak of the distribution is found via a fit to a δ function plus an exponential tail, all with Gaussian smearing.

⁴Channels which fail these checks are excluded from analysis in both data and simulation. Thirty-four main PMTs (2.7%) and zero veto PMTs are so excluded.

⁵ $N_{\text{veto}} \leq 11$ and $N_{\text{main}} \leq 80$. See Chapter 5 for definitions.

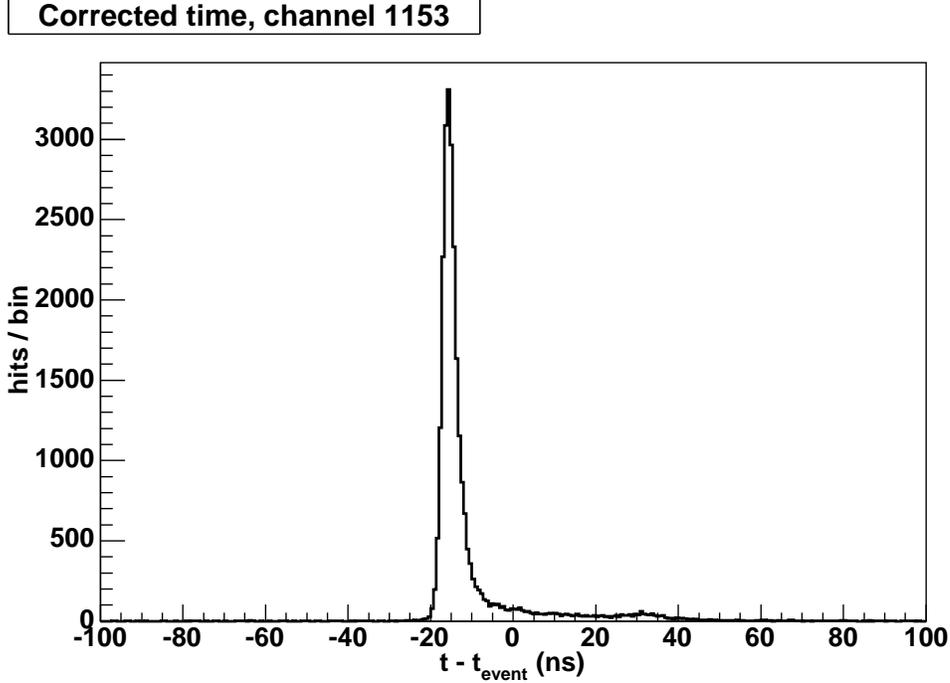


Figure 4.10: An example t distribution used to establish timing offset calibration constants.

The extracted peak times for all PMTs are shifted *en masse* so that the mean peak time is zero, thereby removing the arbitrary global offset. The negative of each PMT's shifted peak time is used as its Δt_{offset} :

$$\Delta t_{\text{offset}}^{\text{PMT } i} = - \left(t_{\text{peak}}^{\text{PMT } i} - \langle t_{\text{peak}} \rangle_{\text{all PMTs}} \right) \quad (4.13)$$

4.3.3 The extended V_q^{ref}

The procedure described in §4.2.3 yields the reference curves $V_q^{\text{ref}}(n; T_{\text{raw}})$ for $n \leq 3$. A special DAQ running mode exists that writes out *all* digitizations rather than only those associated with hits. Taking laser data in this configuration allows us to create a reference curve that goes beyond $n = 3$ using the methods of §4.2.3.

The shape of V_q is not completely independent of charge, though. Noting that the $n \leq 3$ part of $V_q^{\text{ref}}(n; T_{\text{raw}})$ is used only for lower charge (non-saturated) hits and that the $n > 3$ part is used only for higher charge (saturated) hits, we can account for some of the charge dependence by forming $V_q^{\text{ref}}(n; T_{\text{raw}})$ piecewise. We use single-PE hits for the $n \leq 3$ part and high-charge hits ($\langle q \rangle = 7$ PE) for the $n > 3$ part. We merge the two parts into a single curve by scaling one to the other such that they overlap in the transition region (250 ns to 300 ns). The result is shown in Figure 4.11.

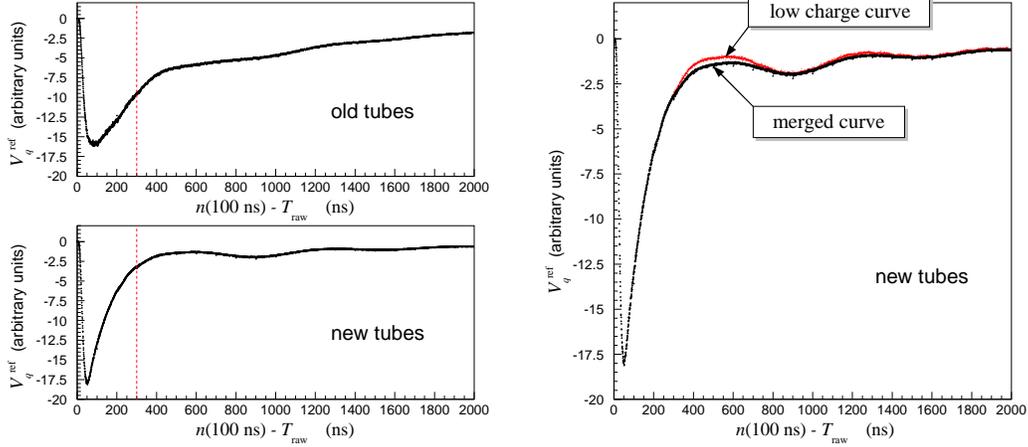


Figure 4.11: The full V_q^{ref} . (Left) The reference curves for old and new tubes obtained using the merging procedure. The regions left of the vertical lines were formed using low-charge data; right of the vertical lines, high-charge. (Right) In black is the new tubes’ merged curve. In red is the analogous curve obtained using only low-charge data. The shape change is apparent.

4.3.4 The charge scale factor g

Recall that $g \equiv \langle Q_{\text{raw}}^{1\text{PE}} \rangle$. We calculate these averages using hits from low-intensity laser events. There are two complications.

Dirty hits. Low-intensity laser events include all manner of hits (dark noise, pre-pulsing, electronics pickup). These represent backgrounds to the determination of g , as each source has, in general, a mean Q_{raw} that differs from the mean Q_{raw} of clean single-PE hits. Fortunately, these “dirty” hits occur infrequently enough that a tight timing cut eliminates their relevance. Figure 4.12 supports this claim, showing that hits near the prompt peak have a clean Q_{raw} distribution, devoid of low- and high-charge features seen away from the peak. We measure g using hits from the prompt window only. g is sensitive to the particular choice of window at the 2% level.

Multi-PE contamination. We can correct for multi-PE contamination using Poisson statistics by putting

$$g = \langle Q_{\text{raw}} \rangle \frac{-p}{\ln(1-p)}, \quad (4.14)$$

where p can be loosely interpreted as the probability per event of getting a hit on a particular tube. For our standard low-intensity laser runs, $p \approx 0.05$. The corresponding multi-PE correction is $\sim 3\%$.

4.3.5 The slew correction

The slew correction $\Delta t_{\text{slew}}(Q_{\text{raw}})$ in Eq (4.1) accounts for the fact that larger PMT signals fire the discriminator earlier than smaller ones. These slew corrections are stored as look-up

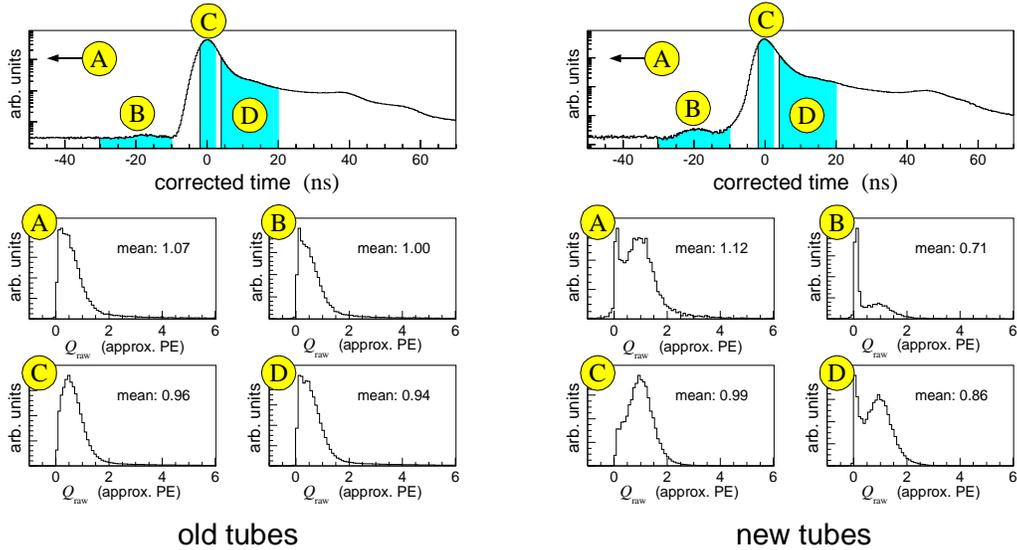


Figure 4.12: Q_{raw} distributions for hits in various time regions. The data are from a low-intensity laser run (Run 1659) with known bad tubes removed and with cosmics suppressed using the cuts $N_{\text{main}} < 60$ and $N_{\text{veto}} < 7$ (see §5.2). The left set of plots shows data for old tubes; the right set, new tubes. The corrected time distribution for each is shown at the top with a logarithmic vertical scale. Four time regions are identified:

- A. $-80 \text{ ns} \leq t < -60 \text{ ns}$ (off the left edge of the time plot)
- B. $-30 \text{ ns} \leq t < -10 \text{ ns}$
- C. $-1.75 \text{ ns} \leq t < 2.25 \text{ ns}$
- D. $4 \text{ ns} \leq t < 20 \text{ ns}$.

The Q_{raw} distributions for hits in these four time regions are shown in the bottom panels, with corresponding labels A through D. (These plots have linear vertical scales.) Note that the prompt region (C) exhibits the characteristic single-PE charge shape, while the off-peak regions (A, B, and D) have extraneous low- and high-charge hits. This is particularly noticeable for new tubes.

tables, with one table for old tubes and one for new. They are established as follows.

We set up N_{bins} bins of Q_{raw} . For hits in each bin, we form a histogram of hit times calculated using Eq (4.1) except without the slew term and (as in the offset determination above) with the laser firing time and photon time-of-flight subtracted. From the resulting distribution we extract the peak time, which varies from one Q_{raw} bin to the next. This Q_{raw} -dependent peak time function is the slew correction. An overall constant term is not relevant in the slew table, and we are free to shift the curve as we like. We choose a shift such that old tubes have $\Delta t_{\text{slew}}(Q_{\text{raw}}=1) \equiv 0$ ns.⁶

Single-PE hits only populate the slew tables out to a few Q_{raw} units, yet we must report times for higher-charge hits. We populate the high charge regions of the slew tables by mixing in higher intensity ($\langle q \rangle \approx 4$ PE) laser data. The final tables are shown in Figure 4.13.

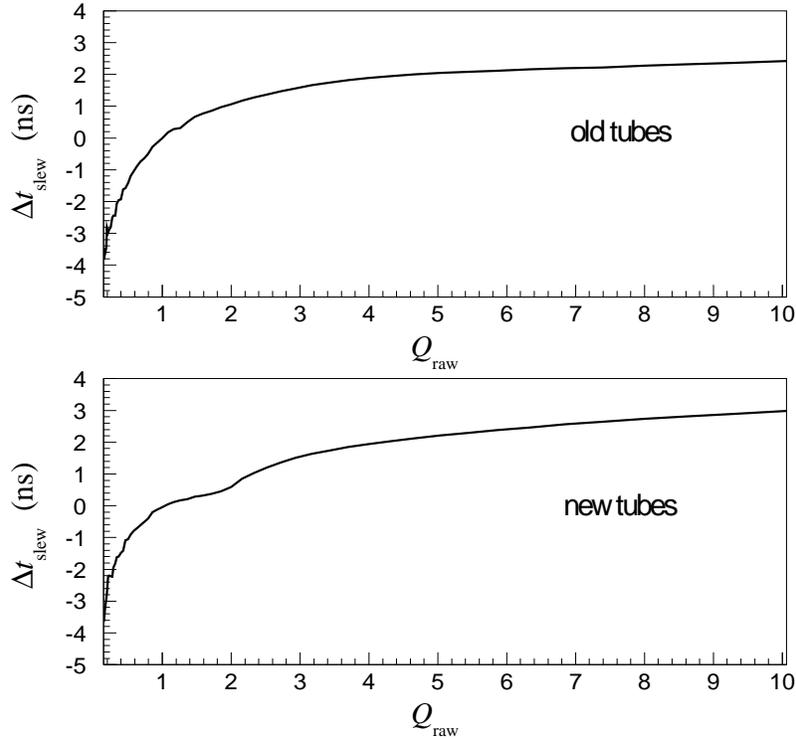


Figure 4.13: The slew tables. The correction $\Delta t_{\text{slew}}(Q_{\text{raw}})$ is plotted for old tubes (top) and new tubes (bottom). We show here only the $Q_{\text{raw}} < 10$ region; the tables continue out to $Q_{\text{raw}} = 50$. $Q_{\text{raw}}=1$ corresponds to about 1 PE.

⁶We give old and new tubes the same shift, letting the new tubes' $\Delta t_{\text{slew}}(1)$ fall where it may. If we wished, we could shift them independently, accounting for the difference with Δt_{offset} .

Variable-width charge bins are used for the slew tables. The bin number N is given by

$$N = \lfloor \log_b(Q_{\text{raw}}) \rfloor + N_1 . \quad (4.15)$$

This binning takes advantage of the higher statistics we have at lower charges while keeping the bin-to-bin change in Δt_{slew} small. We use $b = 1.08$, $N_1 = 50$, and total number of bins $N_{\text{bins}} = 100$.

Note that two effects actually contribute to time slewing:

1. *Discriminator slewing.* This is the familiar phenomenon whereby a large pulse crosses threshold sooner than a small pulse.
2. *“Pile-up” slewing.* When a hit has multiple photoelectrons, the threshold-crossing time depends on the creation times of the individual photoelectrons and (for nearly simultaneous photoelectrons) the jitter in transit times down the PMT dynode stack, which gives each additional photoelectron a shot at providing a (randomly) earlier anode pulse.

Pile-up slewing cannot be estimated without knowledge of the event topology. The reconstruction algorithm (Chapter 9) knows about event topology and accounts for this phenomenon.

Part II

Analysis

Chapter 5

Analysis overview

5.1 The task

The signature for LSND-like $\nu_\mu \rightarrow \nu_e$ oscillations in MiniBooNE is an excess of $\mathcal{O}(100)$ ν_e charged current events. This excess would sit above a comparable number of intrinsic ν_e events and a much larger number (1.7×10^6) of ν_μ events. The $\nu_\mu \rightarrow \nu_e$ appearance search, then, needs the following:

1. an event selection with efficient ν_μ rejection and ν_e acceptance
2. knowledge of the ν_μ and ν_e selection efficiencies
3. estimates of the ν_μ and intrinsic ν_e fluxes through the detector
4. estimates of the relevant cross sections.

The strong ν_μ rejection and high ν_e acceptance of item (1) must be reached without generating unacceptably large systematic errors in (2). Items (2)–(4) are often realized with the support of a near detector, with far-to-near ratios eliminating the largest systematic errors. MiniBooNE does not have a near detector, so we use a combination of available measurements, constraints obtainable from our own data, and Monte Carlo simulations.

5.2 Subevents

For each “spill” of neutrinos sent through the detector, the DAQ records all PMT activity in a $19.2 \mu\text{s}$ window starting $4.6 \mu\text{s}$ before beam arrival time. A given beam spill might have a distribution of PMT hit times like the example in Figure 5.1. Two clusters of activity are visible in the figure – one from a ν_μ charged current interaction and another from the delayed decay of the μ^- that was produced in the event and which came to rest in the detector.

These clusters of time-related PMT hits, or “subevents”, are extracted from each beam spill as follows: a subevent is any group of at least ten hits in which no two consecutive hits are separated by more than 10 ns, with the exception that up to two time gaps of 10 ns – 20 ns are allowed. Note:

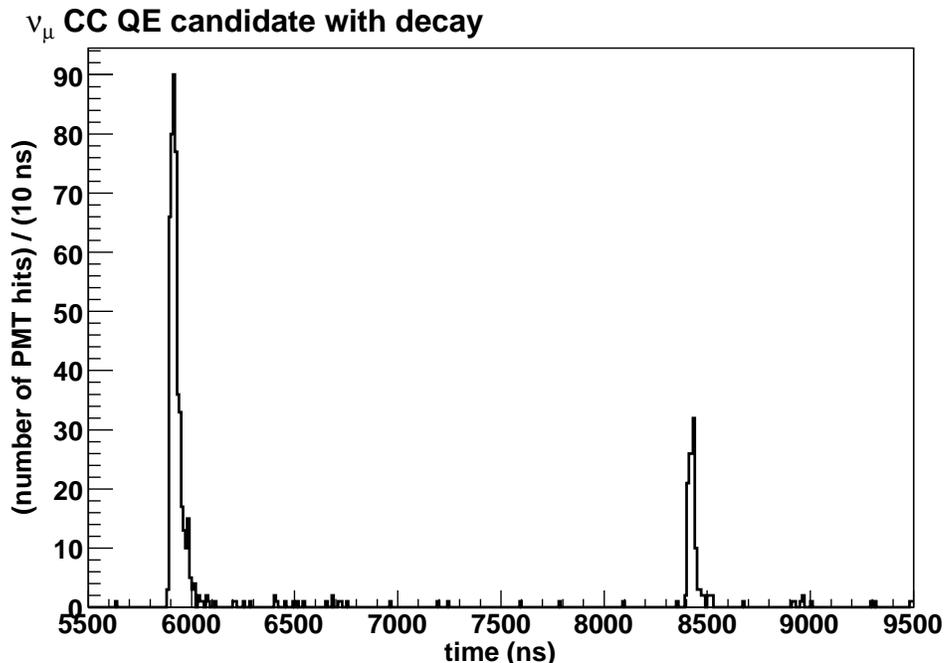


Figure 5.1: A portion of a beam spill containing a candidate ν_μ charged current event. The first cluster of hits is from the primary interaction, and the second is due to the delayed decay of the muon ($\tau_\mu \approx 2000$ ns).

- If a subevent includes six or more PMT hits in the veto region, the beam spill likely involves an exiting track or an incoming cosmic ray. (See Figure 5.2.) Spills with any such “uncontained” subevents are excluded from the oscillation analysis.
- If a beam spill has exactly two contained subevents, the second subevent is typically due to the electron from the decay of a stopped muon from the first subevent.

The top panel of Figure 5.3 demonstrates the latter point, showing that the reconstructed energies of second subevents follow the expected “Michel” electron spectrum [68]. The bottom panel shows the corresponding N_{main} distribution. Cosmogenic Michel subevents are readily eliminated by requiring that candidate neutrino subevents have more than 200 main PMT hits: $N_{\text{main}} > 200$.

Figure 5.4 provides insight into the effects of the $N_{\text{main}} > 200$ and $N_{\text{veto}} < 6$ cuts by showing subevent times¹ with these two cuts variously present and absent. With both cuts on, essentially no out-of-time activity remains.

¹subevent time = average time of the subevent’s PMT hits

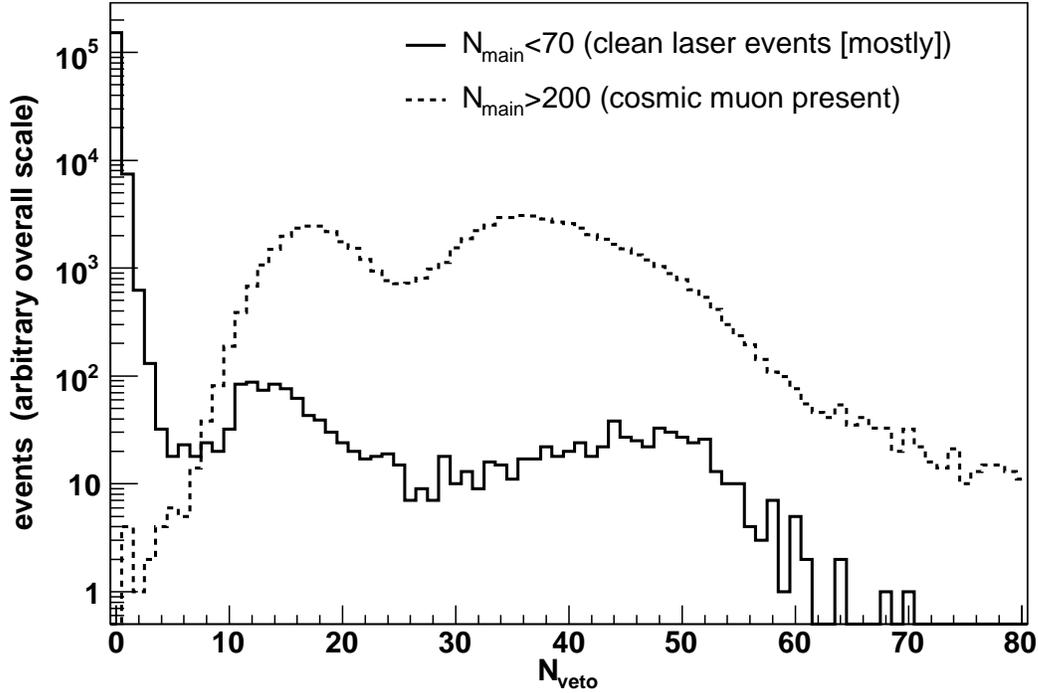


Figure 5.2: Distribution of veto multiplicity N_{veto} for two classes of low-intensity laser flask events. (Solid) Clean, cosmic-free laser events have fewer than 70 hits in the main region, and an $N_{\text{main}} < 70$ cut leaves behind mostly uncontaminated laser events. The solid histogram thus shows a large spike for $N_{\text{veto}} < 6$. (Note the logarithmic vertical scale.) (Dashed) Too much activity in the main region indicates that a cosmic muon is present, and the dashed histogram shows that the veto is typically active when an $N_{\text{main}} > 200$ cut is applied. The two peaks correspond to one veto-crossing or two (*i.e.*, whether the entering muon stops in the main region or continues out the other side). These peaks appear in the solid histogram as well since muons that stop soon after entering the tank (first peak) or that pass obliquely through the tank (second peak) can leave behind few enough hits to satisfy the nominally clean sample's $N_{\text{main}} < 70$ requirement. In the oscillation analysis, cosmic muon contamination is eliminated by requiring $N_{\text{veto}} < 6$.

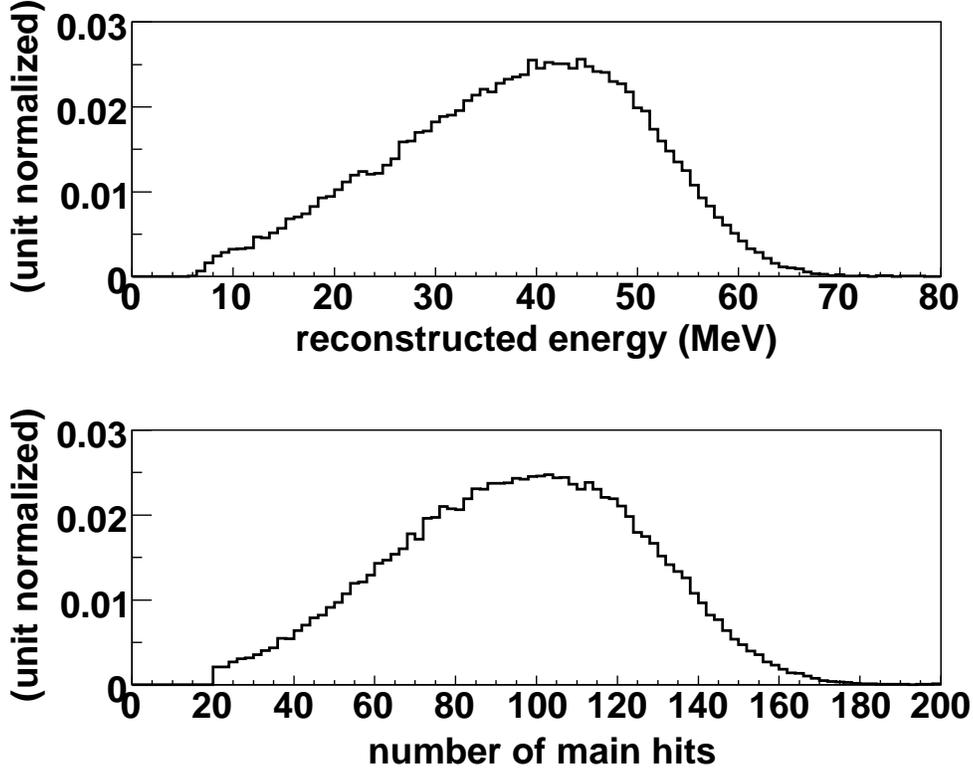


Figure 5.3: Michel electrons (stopped $\mu \rightarrow e\nu\nu$). (Top) Reconstructed electron energies for Michel electron candidates, each of which is the second subevent in a beam spill with exactly two contained subevents. The Michel endpoint of $E_0 = 53$ MeV can be seen (resolution @ E_0 : 12%). (Bottom) The number of main PMT hits present in these subevents. An analysis cut of $N_{\text{main}} > 200$ is applied in the $\nu_\mu \rightarrow \nu_e$ search to eliminate stray Michel electron subevents.

5.3 Neutrino events

If we are to learn the incoming neutrino’s flavor, we must select charged current interactions. If we are to observe any oscillation-induced spectral distortion, we must select events for which we can calculate the incident neutrinos’ energies. Charged current quasi-elastic (CC QE) scattering:

$$\nu_l + n \rightarrow l^- + p \quad (5.1)$$

satisfies both of these criteria and accounts for $\sim 40\%$ of the total ν_μ and ν_e cross sections at MiniBooNE energies.

The CC QE final state involves a charged lepton and a recoil proton. While both emit some scintillation light, the lepton’s Cherenkov emission dominates the event. (The proton is typically below its 870 MeV/ c Cherenkov threshold.) This Cherenkov light is the main source of event information: lepton direction, energy, flavor, and the interaction vertex.

Contained ν_μ CC events are usually flavor tagged by the observation of the μ decay.

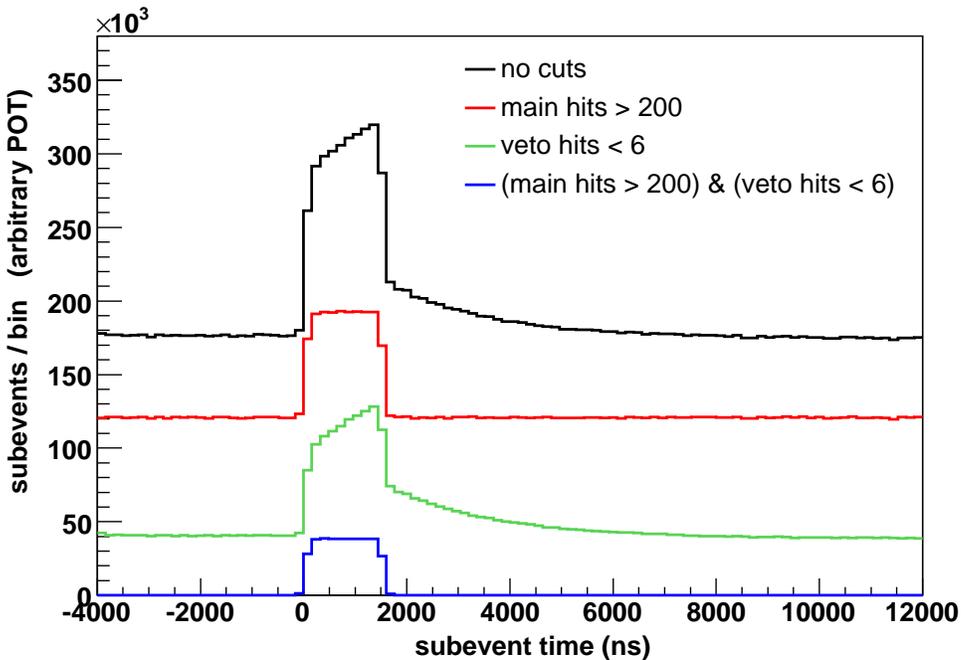


Figure 5.4: Subevent times. The topmost histogram (black) shows the times, relative to the start of the beam spill, of all subevents from a portion of the MiniBooNE run. With no cuts applied, the beam-on excess is apparent. The rising slope within the beam spill and the decaying feature immediately after are due to beam-induced Michel electrons. The second histogram (red) includes only those subevents that sit above the Michel electron endpoint – that is, those subevents satisfying $N_{\text{main}} > 200$. The sloped Michel features are now absent. The third histogram (green) shows only contained subevents, $N_{\text{veto}} < 6$. The bottom histogram (blue) includes both cuts and reveals a clean beam-only distribution with negligible out-of-time activity. A timing cut of $4000 \text{ ns} < (\text{subevent time}) < 7000 \text{ ns}$ is applied to all candidate neutrino interactions in the analysis, but it has no appreciable effect given the multiplicity cuts $N_{\text{main}} > 200$ and $N_{\text{veto}} < 6$ already used. We mention the timing cut nowhere else.

When this tag fails², the μ flavor must be learned from the pattern of Cherenkov light. Since the muon travels many times the PMT spacing before it builds up appreciable angular deflection from multiple Coulomb scattering, it creates a ring of Cherenkov light at the PMT surface with a sharp outer edge that fills in as the muon approaches the tank wall. In contrast, electrons from ν_e interactions lose energy mostly via bremsstrahlung (critical energy $E_c \approx 80 \text{ MeV}$ [9]). Thus, the electron signature is a diffuse ring due to the electromagnetic shower induced. Simulation is used to quantify the differences in the electron and muon light patterns. An example muon ring is shown in Figure 5.5. Data and simulation are compared in Figure 5.6.

²causes are: μ^- capture on carbon (8%, §8.2), μ/e subevent overlap (8%), and below-threshold decay electrons (2%)

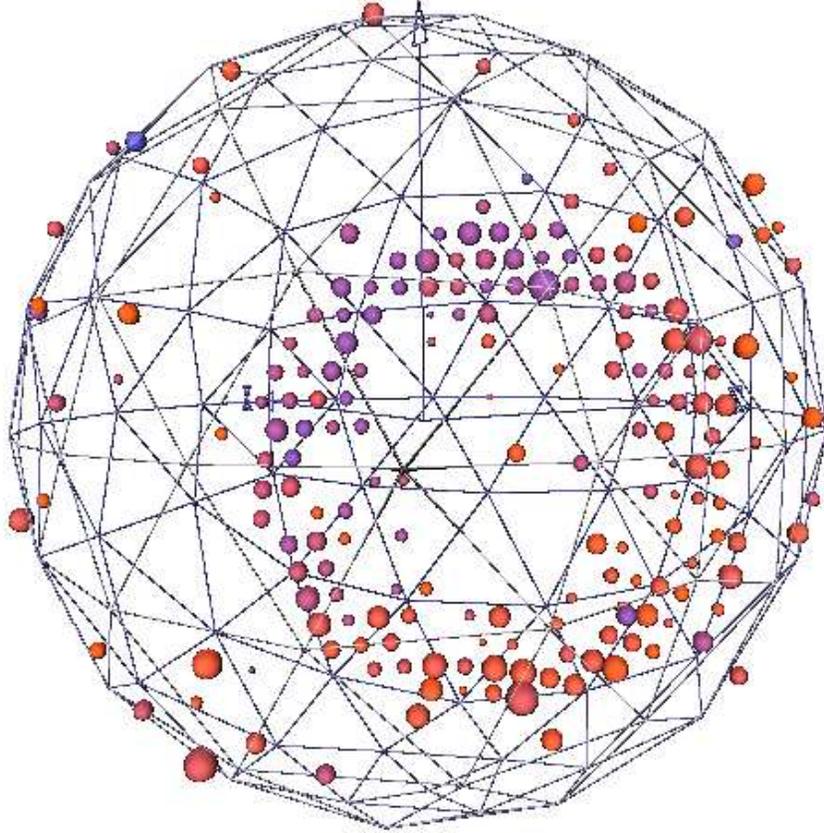


Figure 5.5: A candidate ν_μ CC QE event. A timing cut is applied to remove delayed fluorescence and scintillation light so that the prompt Cherenkov ring can be seen. Each ball represents a PMT hit: size \Leftrightarrow charge, color \Leftrightarrow time.

The electromagnetic shower induced by an electron is indistinguishable from one induced by the e^+e^- pair from γ conversion. Consequently, events with γ 's in the final state constitute backgrounds to the ν_e search. Indeed, neutral current π^0 production (with the subsequent decay $\pi^0 \rightarrow \gamma\gamma$) is the largest source of misidentified ν_μ events. The creation of Δ resonances sometimes leads to the radiative decay $\Delta \rightarrow N\gamma$, resulting in another misidentification background. Figure 5.7 reviews the signal and backgrounds discussed so far.

5.4 Blind analysis

Understanding the neutrino fluxes and cross sections, developing the particle identification methods, and predicting final event rates and spectra were all carried out without the possibility of knowing whether an LSND-like oscillation signal was present in the MiniBooNE data. This “blindness” was enforced by embargoing any event that satisfied conservative ν_e selection criteria. As needs and understanding progressed, so did the aggressiveness of the criteria. Aspects of this blindness appear in subsequent chapters.

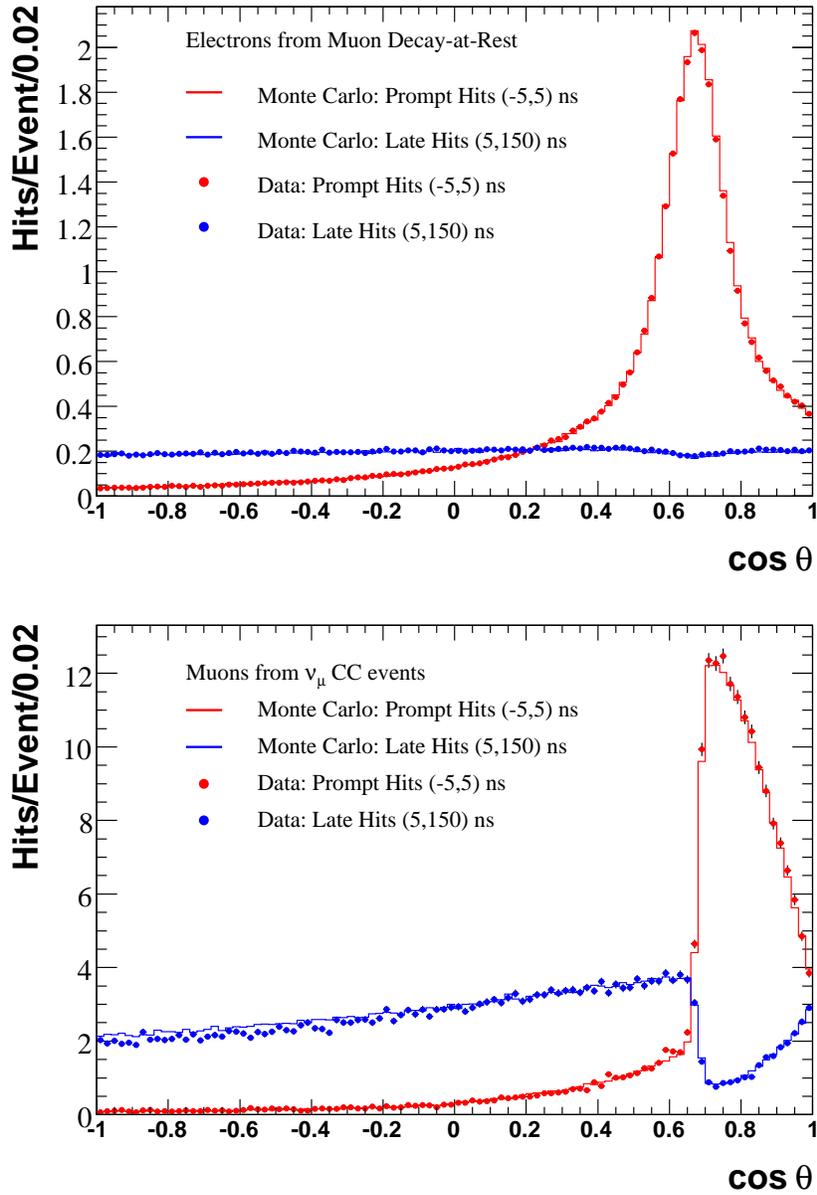


Figure 5.6: Ring profiles. (Top) Electrons from stopped- μ decay. The best-fit electron direction and starting vertex are found, and PMT hits are histogrammed according to their angular location about the reconstructed vertex relative to the reconstructed direction. Prompt (red) and delayed (blue) hits are shown separately. The Cherenkov ring can be seen as prompt light. (Bottom) ν_μ CC QE events. The cavity in the delayed light histogram is due to electronics lockout in channels that received prompt hits. It is more pronounced in the muon plot due to the higher PMT occupancy (evident in the vertical axes.) Also, these higher energy muons result in extended tracks which hide the isotropic nature of the delayed light (*cf.* the flat $\cos \theta$ distribution in the top panel). Data and simulation are both shown.

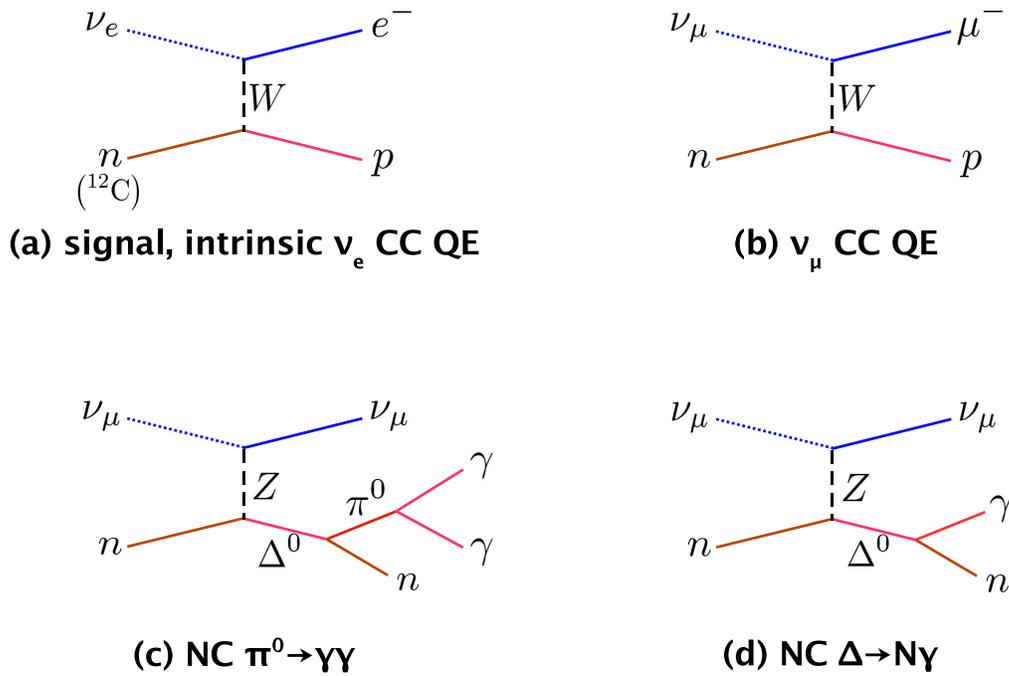


Figure 5.7: Feynman-esque diagrams of four important processes in the ν_e appearance search. (a) Signal and intrinsic ν_e events are indistinguishable. (b) ν_μ CC QE events, when not tagged by their decay electron, must be identified by the muon’s Cherenkov pattern. (c) π^0 ’s produced in neutral current interactions typically decay to two photons (shown). The two γ -induced EM showers are sometimes mistaken for a single EM shower, resulting in misidentification. (d) Neutral current production of a Δ resonance can result in de-excitation to $N\gamma$ (shown), leaving behind a single EM shower.

5.5 Simulation chain

Signal and background estimates are produced using an extensive Monte Carlo simulation chain. A variety of external and internal (MiniBooNE) data is used to tune the simulation, the principal components of which we list here:

1. GEANT4 [69] simulation of the target and decay regions, including secondary hadron production, propagation, and decay (Chapter 6);
2. Nuance v3 [70] neutrino cross section model (Chapter 7);
3. GEANT3 [71] detector simulation (Chapter 8);
4. simulation of the PMTs, electronics, and DAQ (Chapter 8).

5.6 Systematic error propagation

Most MiniBooNE systematic uncertainties can be cast as uncertainties in the above simulation components. These systematic uncertainties are propagated by varying simulation parameters according to their underlying probability density functions (p.d.f.'s) and observing how relevant downstream quantities (often histogram bin contents) change. Doing this many times allows one to build a joint p.d.f., usually encapsulated in a covariance matrix, for these downstream quantities. This Monte Carlo method for propagating uncertainties is discussed in context in Chapter 12.

5.7 Likelihood ratios and $M_{\gamma\gamma}$

The maximum likelihood event reconstruction described in Chapter 9 plays an important particle identification role. Best-fit track parameters and the corresponding maximum likelihoods are extracted under three hypotheses: electron, muon, and π^0 . Ratios of the maximum likelihoods provide differentiation among the three hypotheses. If we label the three likelihoods \mathcal{L}_e , \mathcal{L}_μ , and \mathcal{L}_{π^0} , we can write two particle identification discriminants used in the analysis:

$$\log(\mathcal{L}_e/\mathcal{L}_\mu) = F_\mu - F_e \quad (5.2)$$

$$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0}) = F_{\pi^0} - F_e , \quad (5.3)$$

where the F quantities are those introduced in Chapter 9: $F = -\log(\mathcal{L})$. The sign convention in Eqs. (5.2) and (5.3) is such that electron-like events have more positive values than μ - or π^0 -like events. The expectation is that signal and background roughly straddle zero ($\mathcal{L}_e=\mathcal{L}_\mu \Leftrightarrow \log(\mathcal{L}_e/\mathcal{L}_\mu)=0$). Events are also fit under a generic $\gamma\gamma$ hypothesis with no specified invariant mass. (The π^0 hypothesis enforces a mass of M_{π^0} .) The resulting fitted mass $M_{\gamma\gamma}$ is used to select/reject π^0 s. Cuts on these quantities are established and studied with simulation.

Chapter 6

Neutrino flux prediction

6.1 GEANT4 beamline description

MiniBooNE uses a GEANT4-based [69] Monte Carlo simulation of the target and decay regions. Simulated pions and kaons produced in p -Be interactions are propagated through the target pile (beryllium target slugs, aluminum horn, iron shielding, concrete support structures, iron collimator) and the horn’s magnetic field. Downstream, the particles encounter the decay pipe, the 25 m and 50 m absorber cavities, the 50 m absorber plates, and the overburden.

6.1.1 Proton beam

The process begins with 8 GeV kinetic energy protons 1 cm upstream of the target. The transverse positions (x, y) and directions (θ_x, θ_y) of the protons are chosen from correlated Gaussian distributions with the parameters listed in Table 6.1. The correlations are set such that the beam gets focused to a waist halfway down the target. Reasonable variations in the beam’s profile and convergence properties produce negligible change in the predicted flux (0.8%) [72].

Beam protons stepped through the target will undergo elastic, quasi-elastic, and inelastic interactions. The cross sections for these processes are discussed below. The total p -Be interaction cross section at 8 GeV is 286 ± 15 mb, implying that 17% of incident protons traverse the target without interacting.

quantity	Gaussian σ
x	0.151 cm
y	0.075 cm
θ_x	0.66 mrad
θ_y	0.40 mrad

Table 6.1: Beam profile parameters. Each quantity is drawn from a Gaussian of mean 0 and width σ as indicated. The x and θ_x distributions are anticorrelated ($\rho=-1$), as are the y and θ_y distributions. The values come from beam optics measurements.

6.1.2 Hadron production

Each inelastic p -Be interaction has the potential to produce secondary hadrons. For each type of hadron, the number of particles produced in a given inelastic interaction is chosen from a Poisson distribution with mean

$$N_{\pi^+} = \frac{\iint \frac{d^2\sigma}{dp d\Omega} dp d\Omega}{\sigma_{\text{INE}}}, \quad (6.1)$$

where $\frac{d^2\sigma}{dp d\Omega}$ is the doubly differential production cross section for the hadron and where σ_{INE} is the total p -Be inelastic cross section. This Poisson production is carried out separately for π^\pm , K^\pm , K^0 , n , and p .

Sanford-Wang (π^\pm , K^0)

The differential cross sections used for π^+ , π^- , and K^0 production are described by the Sanford-Wang (SW) [73] parametrization:

$$\frac{d^2\sigma}{dp d\Omega} = c_1 p^{c_2} \left(1 - \frac{p}{p_B - c_9}\right) \exp \left[-c_3 \frac{p^{c_4}}{p_B^{c_5}} - c_6 \theta (p - c_7 p_B \cos^{c_8} \theta) \right], \quad (6.2)$$

where p_B is the proton (“beam”) momentum, p is the outgoing meson momentum, and θ is the outgoing meson direction relative to the proton direction. For π^\pm production, we fix $c_9 \equiv 1$ GeV/ c [74].

Fits to available measurements of $\frac{d^2\sigma}{dp d\Omega}$ are used to determine the SW parameters for each meson. The following χ^2 expression is minimized in each fit:

$$\chi^2 = \sum_{i,j} (D_i - N_{k(i)} F_i) V_{ij}^{-1} (D_j - N_{k(j)} F_j) + \sum_k \frac{(N_k - 1)^2}{\sigma_k^2}, \quad (6.3)$$

where i (or j) labels the individual cross section data points D_i , k labels individual experiments (with $k(i)$ representing the experiment to which data point i belongs), F_i is the value of the SW function at the (p, θ) coordinate of data point i , \mathbf{V}^{-1} is the inverse of the covariance matrix for the data points, N_k is a fit parameter adjusting the normalization of experiment k , and σ_k is the normalization uncertainty quoted by experiment k .

Data from BNL E910 [75] and HARP (CERN) [76] are used for the π^+ cross section fit [74]. The E910 collaboration measured $\frac{d^2\sigma}{dp d\Omega}$ at beam momenta of 6.4, 12.3, and 17.6 GeV/ c and for outgoing π^+ kinematics of 0.4 GeV/ $c < p < 5.6$ GeV/ c and $0.018 < \theta < 0.4$. HARP provided data at $p_B = 8.9$ GeV/ c (MiniBooNE’s beam momentum) over the pion kinematic range of 0.75 GeV/ $c < p < 6.5$ GeV/ c and $0.030 < \theta < 0.21$. These experiments cover the ranges of pion production angles and momenta relevant for MiniBooNE, as shown in Figure 6.1.

The parameters c_3 and c_5 exhibit strong correlations in the π^+ fit [74]. Setting $c_3 \equiv 1$ has negligible impact on the ability of the SW function to represent the data, and this is done for the final π^+ fit. The minimum of χ^2 changes by only 0.02 from the free- c_3 fit.

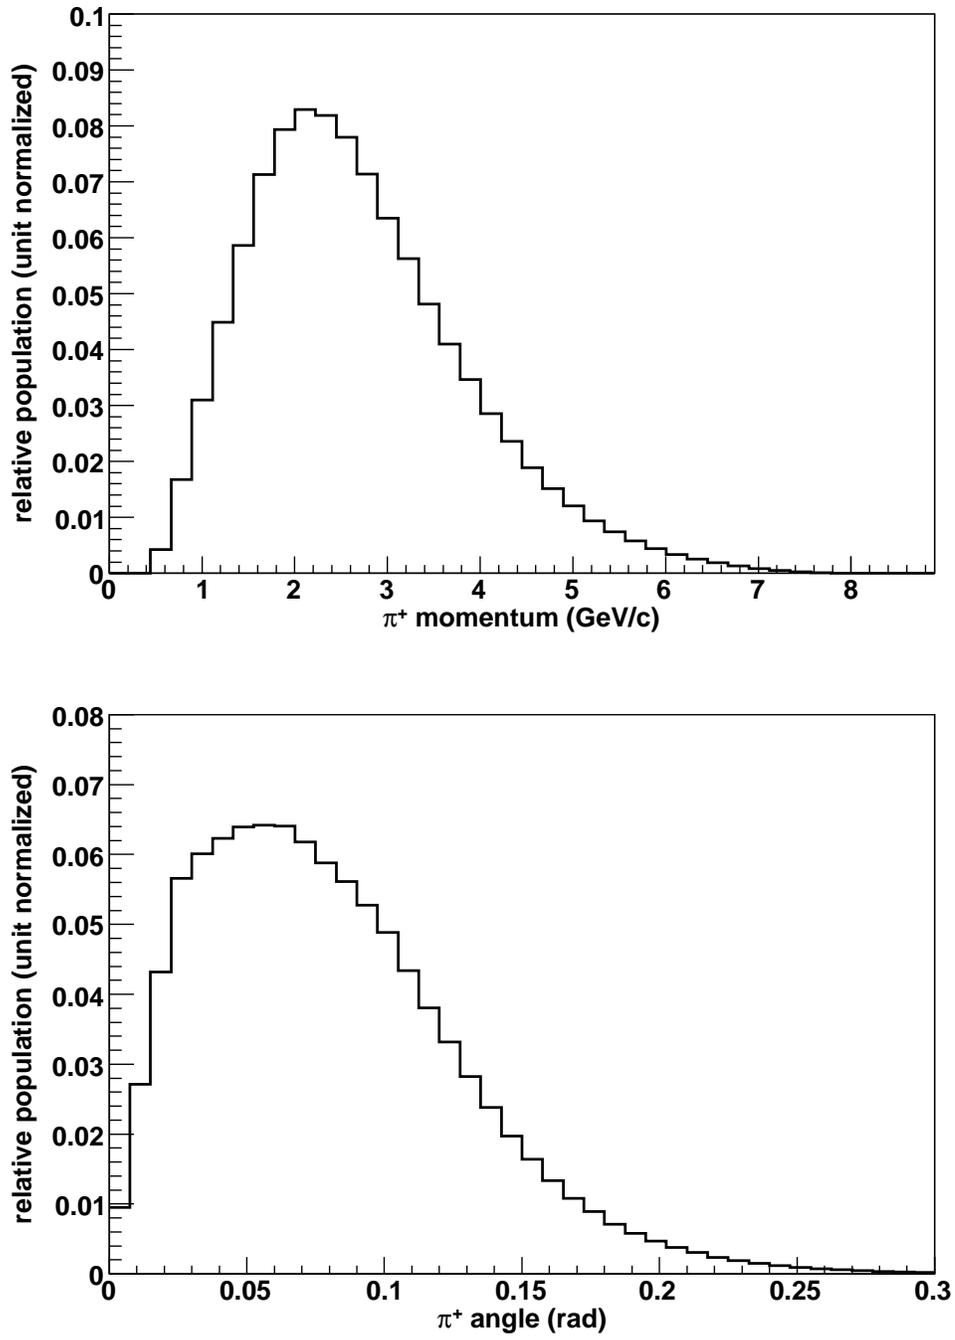


Figure 6.1: π^+ momentum and angle distributions for those pions which produce neutrinos in the MiniBooNE detector. The distributions for other mesons cover similar ranges.

The π^+ fit has $\chi^2/N_{\text{dof}} = 298/165 = 1.8$ and gives fitted experimental normalizations of $N_{\text{E910}} = 1.038$ and $N_{\text{HARP}} = 0.974$ (to be compared with the quoted normalization uncertainties of $\sigma_{\text{E910}} = 5\%$ and $\sigma_{\text{HARP}} = 4\%$). The high value of 1.8 for χ^2/N_{dof} implies that parameter errors, if derived from the $\Delta\chi^2$ surface, will likely be underestimated. Thus, the input covariance matrix \mathbf{V} is scaled by $\chi^2/N_{\text{dof}} = 1.8$ before the correlated errors for the $\{c_n\}$ are extracted.¹ Figure 6.2 shows the best-fit SW function along with HARP data.

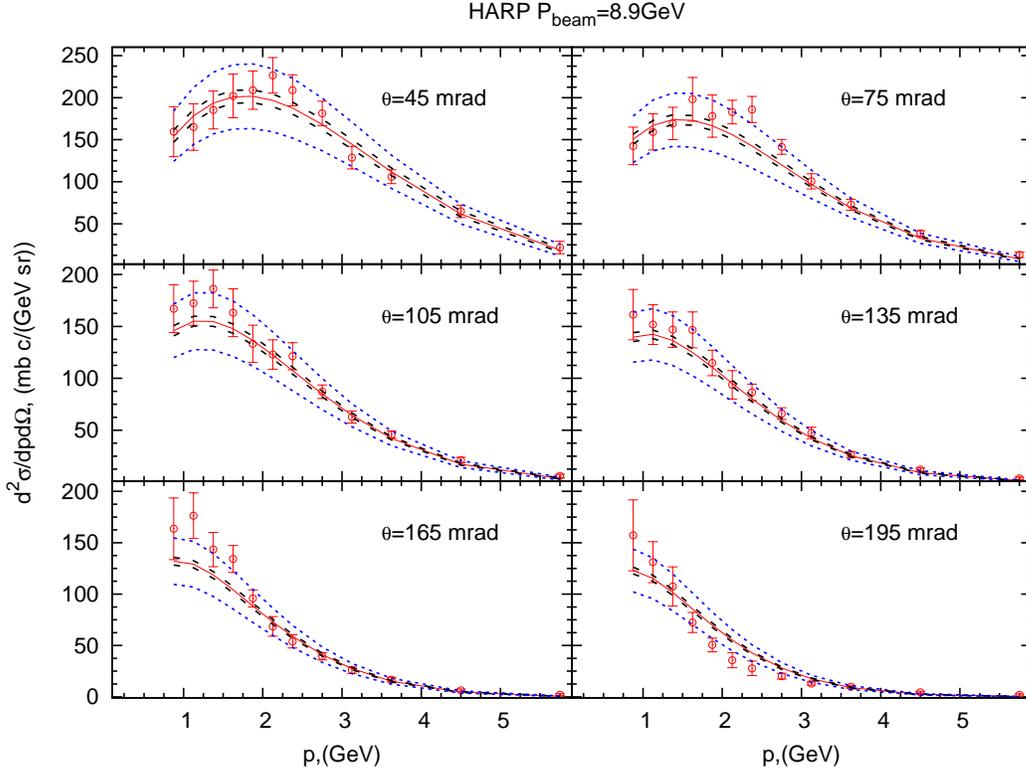


Figure 6.2: π^+ Sanford-Wang function (red curve) and HARP data (red points). The blue curves represent the uncertainties assigned. To form these uncertainty “bands”, 1000 random draws were chosen from the Gaussian p.d.f. implied by the 7×7 parameter covariance matrix. The band indicates the RMS cross section so obtained at each (p, θ) . From Ref. [74].

For π^- production, E910 ($p_B = 6.4, 12.3$ GeV/ c) and HARP ($p_B = 8.9$ GeV/ c) data are again used, but c_3 stays a free parameter in the fit. For K^0 production, data from E910 ($p_B = 12.3, 17.6$ GeV/ c) and from an experiment by Abe *et al.* at KEK ($p_B = 12.0$ GeV/ c) [78] are used. The Abe K^0 momentum range (4 GeV/ $c < p < 6$ GeV/ c) makes their $\theta = 0.061$ and $\theta = 0.087$ measurements most relevant, and higher angles are not included in the fit [77]. Table 6.2 summarizes the extracted SW parameters.

¹If this “solution” to the poor χ^2 is worrisome, note that the π^+ errors end up having little influence on the $\nu_\mu \rightarrow \nu_e$ result thanks to the constraints of Chapters 10 and 11.

	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9
π^+ :	220.7	1.080	1	1.978	1.32	5.572	0.08678	9.686	1
π^- :	237.2	0.8986	4.521	1.154	1.105	4.224	0.06613	9.961	1
K^0 :	15.13	1.975	4.084	0.9277	0.7306	4.362	0.04789	13.3	1.278

Table 6.2: Sanford-Wang parameters for π^+ , π^- , and K^0 production [74] [77]. The units are such that p and p_B expressed in GeV/c yields a differential cross section in (mb c GeV⁻¹ sr⁻¹).

Feynman scaling (K^+)

For the fit to K^+ production data, a parametrization was developed that obeys Feynman scaling (FS) [79], whereby the invariant cross section $E \frac{d^3\sigma}{dp^3}$ depends only on the transverse momentum p_t of the outgoing meson in the center-of-mass frame and on the Feynman scaling parameter

$$x_F = \frac{p_{\parallel}^{\text{CM}}}{p_{\parallel}^{\text{CM,max}}} , \quad (6.4)$$

where $p_{\parallel}^{\text{CM}}$ is the longitudinal momentum of the meson in the center-of-mass frame and $p_{\parallel}^{\text{CM,max}}$ is the maximum possible value of the same [80]. The FS parametrization, given by

$$\frac{d^2\sigma}{dp d\Omega} = \frac{p^2}{E} c_1 \exp(-c_3 |x_F|^{c_4} - c_7 |p_t x_f|^{c_6} - c_2 p_t - c_5 p_t^2) , \quad (6.5)$$

exhibits good agreement with K^+ data over a larger kinematic range than does the Sanford-Wang parametrization.

Table 6.3 lists the data used in the K^+ FS fit. Measurements outside of $1.2 \text{ GeV}/c < p < 5.5 \text{ GeV}/c$ were not included.² Table 6.4 gives the best-fit FS parameters, and Figure 6.3 shows the resulting function with data. As with the Sanford-Wang fits, the FS parameter errors were extracted from the $\Delta\chi^2$ surface only after the input covariance matrix was scaled by $\chi^2/N_{\text{dof}} = 2.28$. Additionally, the resulting errors were doubled as a precaution against normalization disagreements historically observed by neutrino experiments utilizing similar hadron production fits. This latter increase has a minor impact on oscillation sensitivity.

Miscellany

- The cross sections for K^- , n and p production are taken from the MARS15 hadron Monte Carlo package [89].
- All meson production cross sections are artificially increased at high p , with subsequent events weighted appropriately to nullify the enhancement, so that useful numbers of neutrinos at high energy can be collected with reasonable CPU time.

²This is the approximate range of relevance for MiniBooNE (*cf.* Figure 6.1 for π^+). Extending the fit down to $p = 0 \text{ GeV}/c$ and up to $p = 6 \text{ GeV}/c$ adds 32 data points to the existing 119, yet it increases the minimum χ^2 by 337 units, perhaps suggesting that the FS function would have limited utility as an extrapolation tool. Ref. [80] provides further detail.

K^+ Production Data and Fit (Scaled to $P_{\text{beam}} = 8.89 \text{ GeV}$)

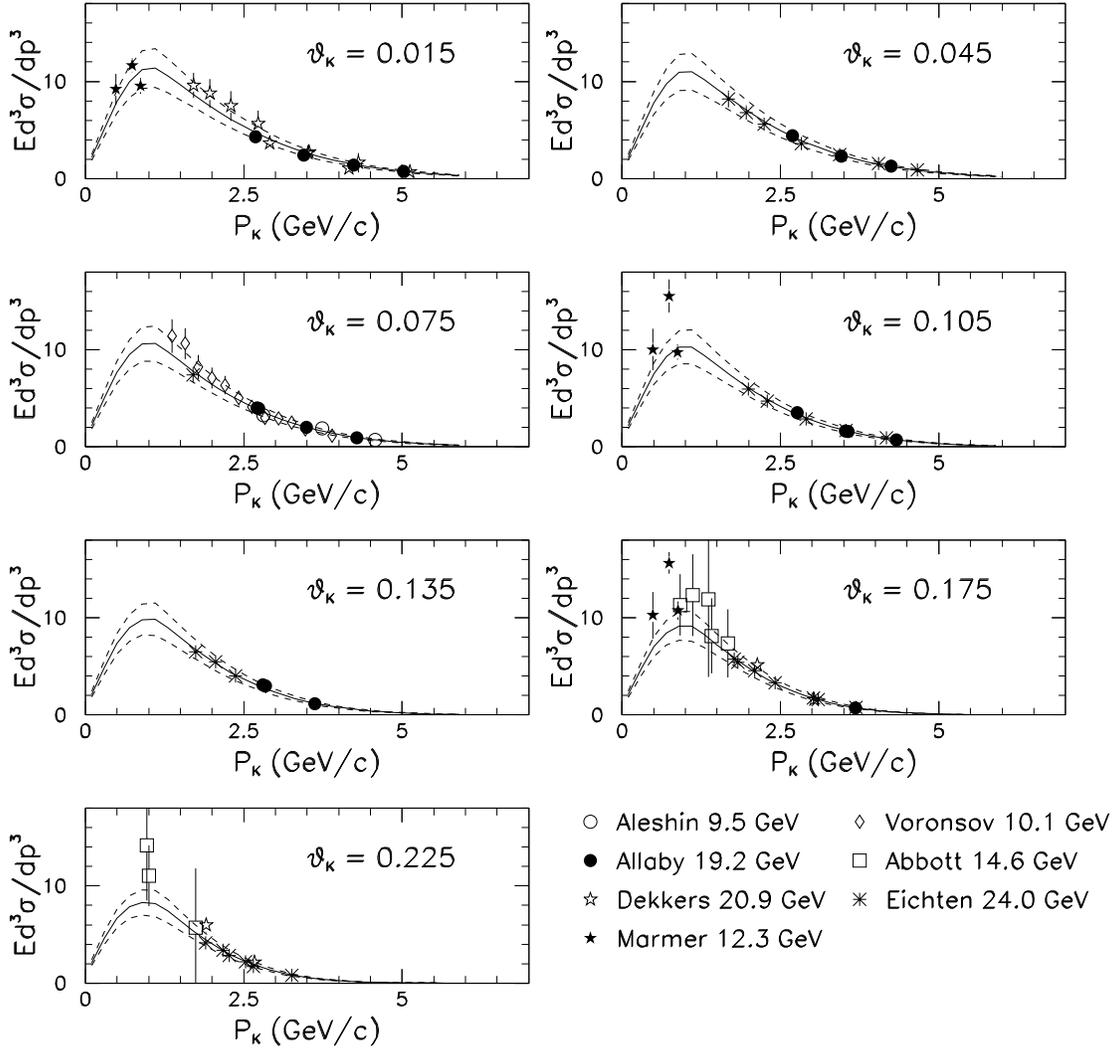


Figure 6.3: FS parametrization (solid curve), error bands (dashed curves), and data (points) for the K^+ production cross section. The data have been scaled to $p_B = 8.9 \text{ GeV}/c$ by preserving the FS quantities p_t and x_F .

data set	p_B (GeV/c)	p (GeV/c)	θ (degrees)	σ_k
Abbott <i>et al.</i> [81]	14.6	2 – 8	20° – 30°	15%
Aleshin <i>et al.</i> [82]	9.5	3 – 6.5	3.5°	10%
Allaby <i>et al.</i> [83]	19.2	3 – 16	0° – 7°	15%
Dekkers <i>et al.</i> [84]	18.8, 23.1	4 – 12	0°, 5°	20%
Eichten <i>et al.</i> [85]	24.0	4 – 18	0° – 6°	15%
Vorontsov <i>et al.</i> [86]	10.1	1 – 4.5	3.5°	$\sim\infty$

Table 6.3: K^+ production data for the FS fits [80]. Only those data satisfying $1.2 \text{ GeV}/c < p < 5.5 \text{ GeV}/c$ were used. While the Vorontsov *et al.* paper reports $\sigma_k=25\%$, their measurements show a large normalization offset relative to the others, so the Vorontsov normalization information is not used (achieved by making σ_k large). These data obey Feynman scaling fairly well (Figure 6.3). An additional data set, Lundy *et al.* [87] ($p_B = 13.4 \text{ GeV}/c$), was anomalous in this regard and was excluded. Another, Marmer *et al.* [88] ($p_B = 12.3 \text{ GeV}/c$), had no measurements satisfying $p > 1.2 \text{ GeV}/c$.

	c_1	c_2	c_3	c_4	c_5	c_6	c_7
K^+ :	11.70	0.88	4.77	1.51	2.21	2.17	1.51

Table 6.4: Feynman scaling parameters for K^+ production [80]. The units are such that E expressed in GeV and p and p_t expressed in GeV/c yield a differential cross section in $(\text{mb } c \text{ GeV}^{-1} \text{ sr}^{-1})$.

6.1.3 Particle propagation

Inelastic scattering

The p -Be and p -Al inelastic cross sections – beryllium for the target, aluminum for the horn – have been measured by Gachurin *et al.* [90] and Bobchenko *et al.* [91] for nucleon momenta of 1 – 9 GeV/c. These data are used in the simulation through the parametrization [92]

$$\sigma_{\text{INE}} = a_1 + a_2 p^{a_5} + a_3 (\log p)^2 + a_4 \log p . \quad (6.6)$$

Parameter values are given in Table 6.5, and the functions are plotted along with data in Figure 6.4. By isospin symmetry, the n and p cross sections are taken to be the same.

Inelastic cross sections for π^\pm -Be and π^\pm -Al are similarly parametrized using data from the above Gachurin and Bobchenko references as well as from Ashery *et al.* [93] and Al-

	a_1	a_2	a_3	a_4	a_5
p -Be:	-126.08	420.02	-6.802	115.47	-0.579
p -Al:	470.94	-0.259	48.86	-87.18	2.429

Table 6.5: Inelastic cross section parameters for Eq. (6.6). The units are such that p expressed in GeV yields a cross section in mb.

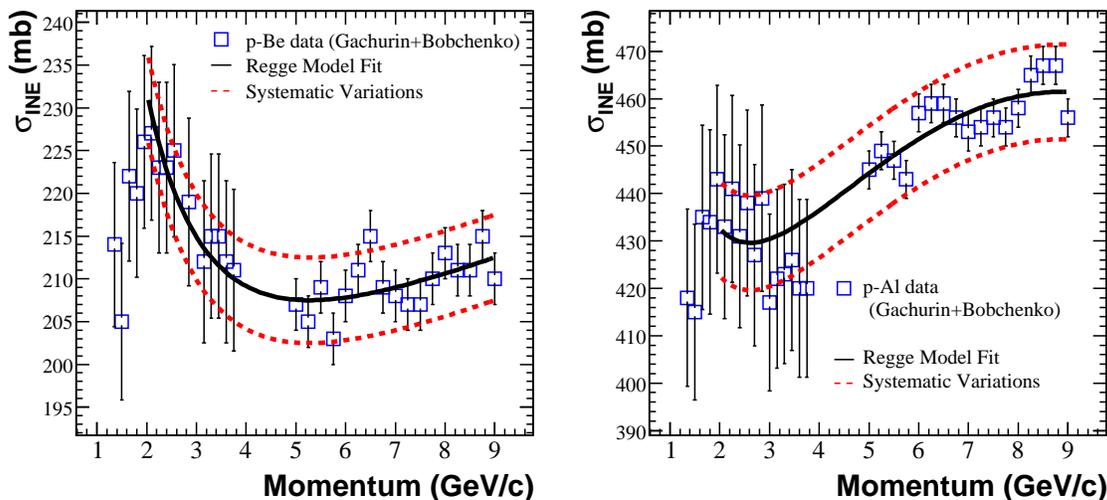


Figure 6.4: The inelastic cross section of Eq. (6.6) (solid line), its assigned uncertainty (dashed lines), and the fitted data for p -Be (left) and p -Al (right). The vertical scales are zero-suppressed. The disagreement (5 – 15%) below 2 GeV/ c has negligible impact on neutrino event rates. From Ref. [92].

lardyce *et al.* [94]. An empirical A scaling was applied to any data taken with non-Be/Al targets. The final states for π^\pm inelastic interactions are determined using default GEANT4 routines.

(Quasi)-elastic scattering

The Glauber model [95] is used to calculate hadron-nucleus elastic (coherent) scattering cross sections, following the example of Franco [96]. In the Glauber model, elastic scattering off a nucleus is treated by summing the scattering amplitudes from each constituent nucleon. The underlying hadron-nucleon cross sections come primarily from Particle Data Group compilations [9].

Hadrons may also interact quasi-elastically (incoherently), scattering off individual nucleons within the nuclear target. Compiled hadron-nucleon cross sections are again used, now within a nuclear shadowing model [95]. Shadowing lowers the quasi-elastic cross section from a simple $A \times \sigma_{\text{nucleon}}$ expectation. Ref. [92] has more.

6.1.4 Production weighting and decays

During propagation, and usually in the decay pipe, the produced mesons decay. Muons from $\pi^+ \rightarrow \mu^+ \nu_\mu$ lead to ν_e 's via $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$. To enhance the Monte Carlo statistics for this important flux component, each π^+ (or π^-) decay is repeated 20 times to produce additional muons; subsequent daughter neutrinos are weighted by $\frac{1}{20}$ to compensate.

The detector's solid angle is small when viewed from the decay region, so it is advanta-

geous to decay all neutrino parents multiple times. The GEANT4 stage of the simulation ends with a list of potential neutrino parents. This list is passed to the REDECAY utility which decays each particle 1000 times, reporting any neutrino with a trajectory that crosses the detector. REDECAY uses GEANT3 decay routines (*e.g.* GDECA2.F) modified to properly handle matrix elements and particle polarizations.

6.2 Neutrino history

The `nuance` neutrino event generator (next chapter) uses neutrino energy spectra – not individual neutrinos – from the beam simulation. Thus, neutrino parent information (*e.g.*, the energy of the π^+ that produced a given ν_μ) is lost at the `nuance` stage of the Monte Carlo chain, and this information must be reintroduced after neutrino interactions have been generated. A history is assigned to each interaction based on the neutrino energy and flavor. The histories come from a large database populated by the very same beam events that fill the neutrino energy histograms used by `nuance`. Since the `nuance`-chosen neutrino energy will not exactly match any of the entries in the database, a tolerance of 0.001 GeV (ν_μ) or 0.01 GeV (other) is allowed. The assigned history specifies the neutrino’s trajectory through the detector, and the detector simulation respects this information when choosing interaction vertices, thereby accounting for the small energy-angle correlations present in the flux.

6.3 Resulting flux

Figure 6.5 shows the predicted fluxes of ν_μ , $\bar{\nu}_\mu$, ν_e , and $\bar{\nu}_e$ at the detector. The vertical scale is the same in all four panels to facilitate comparisons. Table 6.6 gives a breakdown of the integrated flux. Some comments:

- The plots do not reflect the event rate as a function of E_ν since the neutrino cross section has not been applied.
- 0.6% of the flux is ν_e or $\bar{\nu}_e$. The dominant source is μ^+ (from π^+) decay. K^+ decay is a close second.
- 13% of the total ν_e and $\bar{\nu}_e$ flux comes from K_L^0 decay.
- As expected, the $\bar{\nu}_\mu$ and ν_e fluxes have matching contributions from $\pi^+ \rightarrow \mu^+ \rightarrow \nu_e$.

6.4 Systematic uncertainties

The following summarizes the beam-related systematic uncertainties assessed in the $\nu_\mu \rightarrow \nu_e$ analysis. Ref. [72] offers additional detail.

- Correlated errors on the SW parameters for π^+ , π^- , and K^0 production are obtained from the fits described above. Similarly, the FS fits provide K^+ production errors.

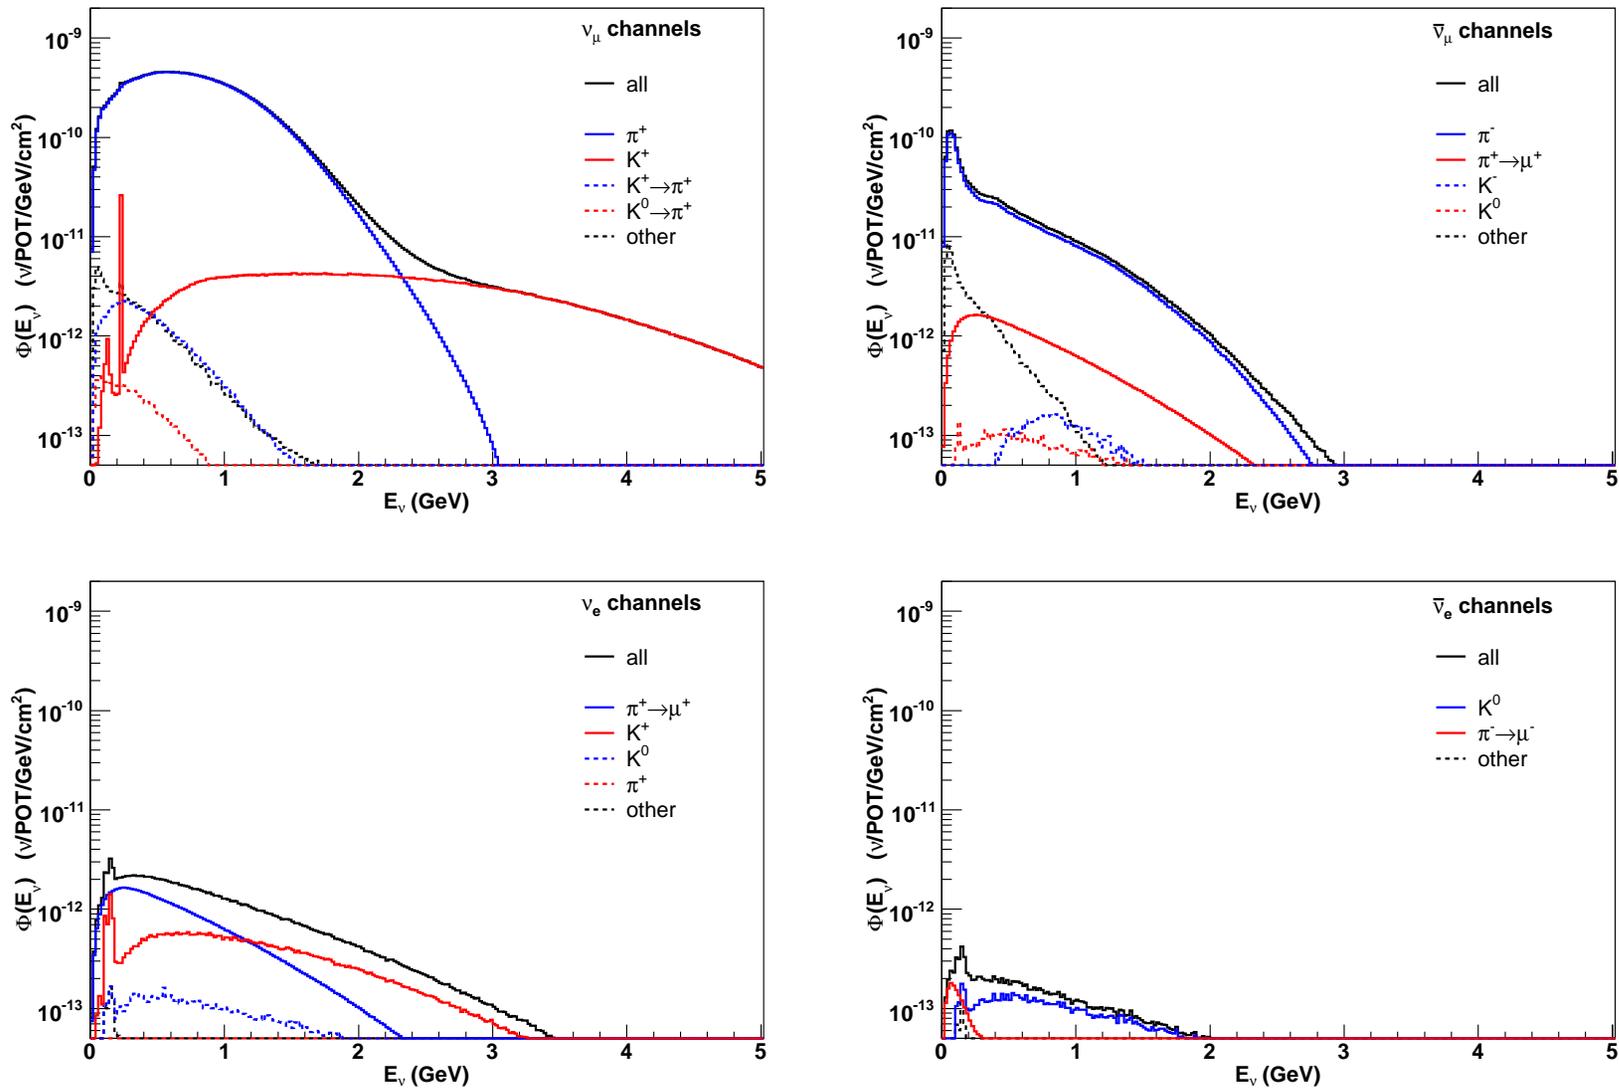


Figure 6.5: Flux of neutrinos at the detector, separated by flavor and production channel. The solid black curve in each panel shows the total flux for that flavor. The production channels listed are all nucleon induced. Arrows denote decay. Neutrinos from mesons that were produced by other mesons are counted in the black dashed spectra (“other”) along with negligible miscellaneous nucleon-induced channels.

total: $2.22 \times 10^{-7} \nu/\text{cm}^2/\text{POT}$			
ν_μ : 93.6%	$\bar{\nu}_\mu$: 5.86%	ν_e : 0.52%	$\bar{\nu}_e$: 0.05%
π^+ : 96.72%	π^- : 89.74%	$\pi^+ \rightarrow \mu^+$: 51.64%	K^0 : 70.65%
K^+ : 2.65%	$\pi^+ \rightarrow \mu^+$: 4.54%	K^+ : 37.28%	$\pi^- \rightarrow \mu^-$: 19.33%
$K^+ \rightarrow \pi^+$: 0.26%	K^- : 0.51%	K^0 : 7.39%	K^- : 4.07%
$K^0 \rightarrow \pi^+$: 0.04%	K^0 : 0.44%	π^+ : 2.16%	π^- : 1.26%
K^0 : 0.03%	$K^0 \rightarrow \pi^-$: 0.24%	$K^+ \rightarrow \mu^+$: 0.69%	$K^- \rightarrow \mu^-$: 0.07%
$\pi^- \rightarrow \mu^-$: 0.01%	$K^+ \rightarrow \mu^+$: 0.06%	other: 0.84%	other: 4.62%
other: 0.30%	$K^- \rightarrow \pi^-$: 0.03%		
	other: 4.43%		

Table 6.6: The fraction of the total flux due to each neutrino type, broken down by production channel. The nucleons that induce these channels are not indicated in the labeling. Arrows denote decay. The “other” category includes (and is dominated by) channels involving meson production by other mesons (rather than by nucleons).

- Errors on the inelastic, quasi-elastic, and total cross sections for nucleon-nucleus and pion-nucleus interactions come from the fits used to establish the parametrizations.
- The horn current is assigned a 1 kA uncertainty (the precision with which it is measured) [97].
- The skin depth of the horn at the frequency of the sinusoidal current pulse is $\delta = 1.4$ mm. Flux histograms are obtained for this skin depth as well as for $\delta = 0$ mm. The flux difference between the two cases (each flavor treated separately) is taken as a one standard deviation flux uncertainty due to skin depth.

Chapter 12 explores how these uncertainties impact the oscillation result.

Chapter 7

Neutrino cross section model

We review neutrino cross sections relevant in MiniBooNE and their implementations in version 3 of `nuance` (our Monte Carlo event generator). Ref. [70] and its citations offer additional detail on `nuance`. Ref. [98] presents some comparisons between available data and `nuance` predictions.

7.1 Charged current quasi-elastic scattering

As Figure 7.1 shows, charged current quasi-elastic (CC QE) scattering

$$\nu_l + n \rightarrow l^- + p \quad (7.1)$$

is the dominant charged current process for MiniBooNE ($E_\nu \sim 0.1 - 2$ GeV). Its cross section can be written¹ [100]

$$\frac{d\sigma}{dQ^2} = \frac{M^2 G_F^2 V_{ud}^2}{8\pi E_\nu^2} \left[A(Q^2) \pm B(Q^2) \frac{(s-u)}{M^2} + C(Q^2) \frac{(s-u)^2}{M^4} \right], \quad (7.2)$$

where M is the target nucleon mass (taking $M \approx M_n \approx M_p$), $-Q^2 = q^2$ is the square of the four-momentum transfer, G_F is Fermi's constant, V_{ud} is the relevant CKM matrix element, E_ν is the incident neutrino energy in the neutron rest frame, $(s-u) = 4ME_\nu - Q^2 - m^2$ involves the usual Mandelstam variables [102], m is the mass of the outgoing lepton l , the $(-)+$ sign corresponds to (anti)neutrino scattering, and the Q^2 -dependent coefficients A ,

¹averaging over spins and assuming a conserved vector current (CVC) and the absence of second-class currents [99]

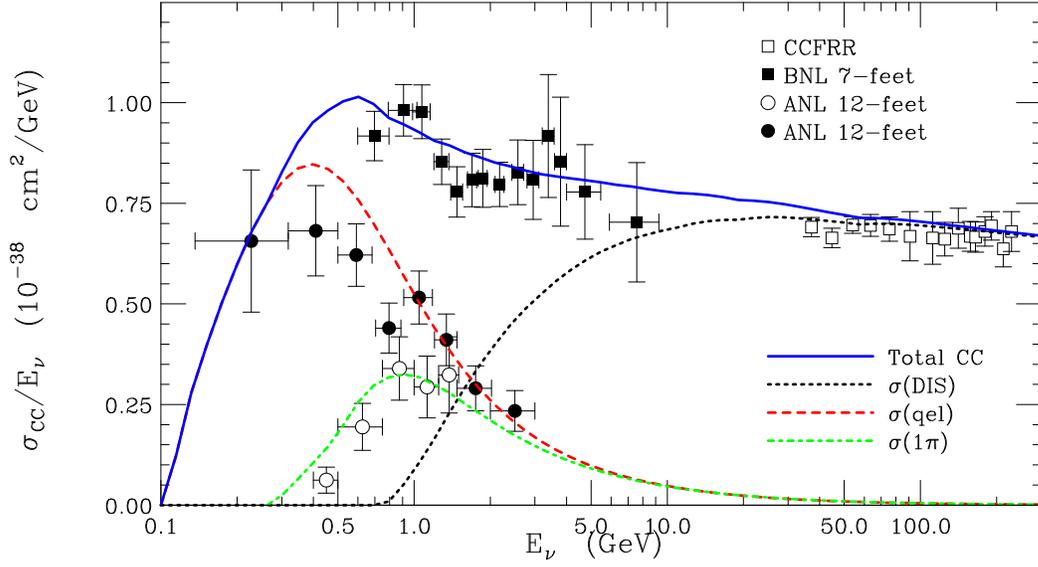


Figure 7.1: Charged current neutrino cross sections versus energy. Shown are total (solid), deep inelastic scattering (dotted), quasi-elastic (dashed), and single pion production (dot-dashed) cross sections. From Ref. [101].

B , and C are given by

$$A = \frac{m^2 + Q^2}{M^2} \left[(1 + \tau)F_A^2 - (1 - \tau)F_1^2 + \tau(1 + \tau)F_2^2 + 4\tau F_1 F_2 \right. \\ \left. - \frac{m^2}{4M^2} [F_1^2 + (F_A + 2F_P)^2 - 4(1 + \tau)F_P^2] \right]$$

$$B = \frac{Q^2}{M^2} F_A (F_1 + F_2)$$

$$C = \frac{1}{4} (F_A^2 + F_1^2 + \tau F_2^2)$$

with $\tau \equiv \frac{Q^2}{4M^2}$.

For a conserved vector current, the vector form factors F_1 and F_2 can be related to the electromagnetic form factors. These, in turn, have a dipole dependence on Q^2 if the radial distribution of charge in the nucleus is exponential, giving [100]:

$$F_1(Q^2) = \frac{1 + \tau(1 + \eta)}{(1 + \tau) \left(1 + \frac{Q^2}{M_V^2}\right)^2} \quad (7.3)$$

$$F_2(Q^2) = \frac{\eta}{(1 + \tau) \left(1 + \frac{Q^2}{M_V^2}\right)^2} . \quad (7.4)$$

Here, $\eta = \left(\frac{\mu_p}{\mu_N} - 1\right) - \left(\frac{\mu_n}{\mu_N} - 0\right) = 3.706$ is the difference between the proton and neu-

tron anomalous magnetic moments [9], and the parameter $M_V = 0.84$ GeV is determined from electron scattering data [103]. Although the dipole assumption is common, **nuance** was upgraded in MiniBooNE to use non-dipole vector form factors developed from recent fits [104].

F_P is related to F_A through [100]

$$F_P(Q^2) = \frac{2M^2}{m_\pi^2 + Q^2} F_A(Q^2) , \quad (7.5)$$

where m_π is the pion mass. **Nuance** uses a dipole form for F_A :

$$F_A(Q^2) = \frac{F_A(0)}{\left(1 + \frac{Q^2}{M_A^2}\right)^2} , \quad (7.6)$$

where $F_A(0) = -1.27$ is determined from neutron decay [9]. The single remaining free parameter M_A is set using our neutrino data (§10.2.1).

Nuclear effects

If the target nucleon is bound, three important conditions modify the above cross section:

1. The target nucleon is not at rest (Fermi motion).
2. The target nucleon has binding energy.
3. The outgoing nucleon must be in a state not already occupied by a spectator nucleon (Pauli blocking).

Item (1) causes an unavoidable smearing in the reconstruction of the neutrino energy. Items (2) and (3) suppress low Q^2 interactions.

Nuance uses the relativistic Fermi gas (RFG) model of Smith and Moniz [105] to incorporate these effects. In the RFG: target nucleons are assumed to have a uniform momentum density in the rest frame of the nucleus up to a cutoff Fermi momentum $k_F = 220$ MeV/ c (Item 1 above); if possible, the outgoing nucleon's energy is reduced by a binding energy $E_b = 34$ MeV, else the reaction is kinematically forbidden (Item 2); and, the momentum of the final state nucleon must be greater than the Fermi momentum (Item 3). These conditions are imposed by integrating, with respect to the initial nucleon momentum \mathbf{k} , the free nucleon cross section multiplied by a factor proportional to

$$\Theta(k_F - |\mathbf{k}|) \Theta(|\mathbf{k} + \mathbf{q}| - k_F) \delta(\epsilon(\mathbf{k}) - \epsilon(\mathbf{k} + \mathbf{q}) - E_b + \omega) , \quad (7.7)$$

where $\mathbf{q} = \mathbf{p}_\nu - \mathbf{p}_l$ is the 3-momentum transfer, $\omega = E_\nu - E_l$ is the energy transfer, and $\epsilon(\mathbf{k}) = (|\mathbf{k}|^2 + M^2)^{\frac{1}{2}}$. The first Θ function establishes the Fermi sea of target nucleons, the second enforces Pauli blocking, and the δ function enforces energy conservation.² In **nuance**, Pauli

²For MiniBooNE's $Z=N=6$ carbon target, a single k_F describes both the neutron and proton Fermi seas.

blocking is equivalently enforced via the lower bound of integration $E_{\text{low}} \Leftrightarrow \sqrt{k_{\text{low}}^2 + M^2}$, which we modify in MiniBooNE by a scale factor $\kappa = 1.024$,

$$E_{\text{low}} = \kappa \left(\sqrt{k_F^2 + M^2} - \omega + E_b \right), \quad (7.8)$$

the motivation and tuning of which is discussed in §10.2.1. In the Smith and Moniz RFG model, $\kappa \equiv 1$.

The need to constrain

The above cross section formalism is commonly used, but it is not particularly well tested in MiniBooNE's region of interest (Figure 7.2). Also, at $Q^2 < 0.2 \text{ GeV}^2$ where nuclear effects are important, the RFG model strains to match observed kinematic distributions [106]. Fortunately, MiniBooNE has a large sample of ν_μ CC QE events that helps constrain (Chapter 10) the ν_e CC QE cross section estimates for the $\nu_\mu \rightarrow \nu_e$ analysis.

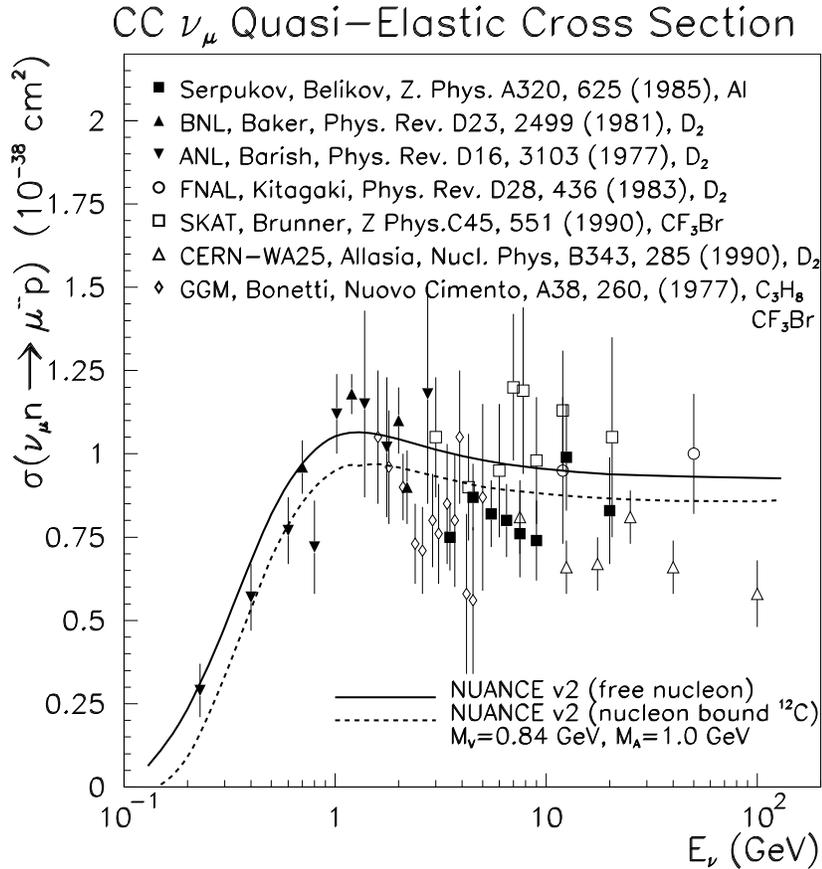


Figure 7.2: Nuance v2 prediction for the ν_μ CC QE cross section along with experimental data. From Ref. [107].

Reconstructing E_ν

The (approximately) two-body CC QE process allows one to reconstruct the incident neutrino energy from knowledge of the outgoing lepton kinematics and the incoming neutrino direction, assuming a stationary unbound target neutron:

$$E_\nu^{\text{rec}} = \frac{2M_n E_l + M_p^2 - M_n^2 - M_l^2}{2(M_n - E_l + \mathbf{p}_l \cdot \mathbf{u}_\nu)}, \quad (7.9)$$

where (E_l, \mathbf{p}_l) is the reconstructed lepton 4-momentum; $M_{\{n,p,l\}}$ are the masses of the neutron, proton, and lepton; and where \mathbf{u}_ν is the incident neutrino direction. E_ν^{rec} appears throughout later chapters.

7.2 Resonant processes

The excitation and subsequent decay of baryon resonances is the primary mechanism for π production, with the $\Delta(1232)$ resonance dominant and $N(1440)$ secondary. **Nuance** models all known non-strange resonances with invariant mass below 2 GeV using the harmonic oscillator quark wave function approach described by Rein and Sehgal [108], updated to reflect modern knowledge of the mass spectrum. The nucleon form factors for resonant scattering are taken to be identical to those for quasi-elastic scattering, though with different values for M_A , and nuclear effects are handled with an RFG prescription. For the latter, Pauli exclusion is enforced for the decay nucleon (*e.g.*, the neutron in $\Delta^+ \rightarrow n\pi^+$) rather than for the resonant baryon. Figure 7.3 shows the predicted cross section for charged current π^+ production along with available experimental data.

Radiative decay

The electromagnetic decay $\Delta \rightarrow N\gamma$ occurs with a branching ratio of 0.52%–0.60%, inferred from pion photoproduction on nuclei [9]. Though relatively rare, the $N\gamma$ final state is all but indistinguishable from the Ne final state sought in the oscillation search. Further, the branching ratio is comparable to the LSND oscillation probability of 0.26% [50]. Thus, radiative Δ decay is an important background in the analysis. We revisit it in Chapter 11 where we use measurements of π^0 production to adjust the **nuance** prediction for $\Delta \rightarrow N\gamma$.

7.3 Neutral current π^0 production

As Figure 7.4 indicates, cross section measurements for neutral current (NC) π^0 production at MiniBooNE energies are sparse. Chapter 11 describes how we isolate a high purity sample of these events to measure their production rate directly, thereby taking the pressure off the **nuance** model and available data.

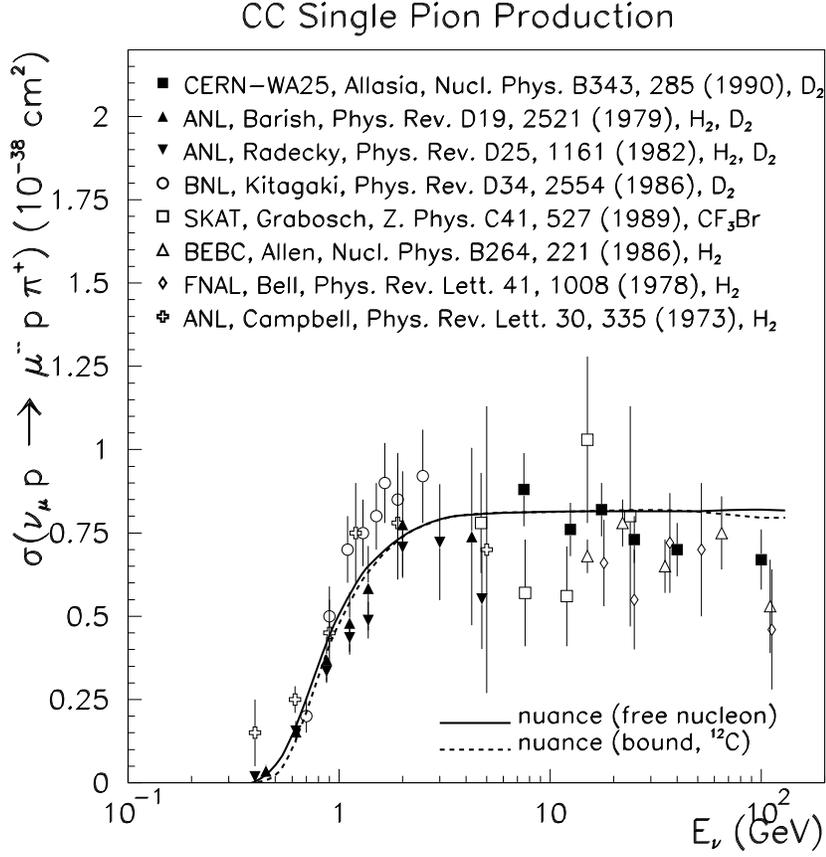


Figure 7.3: CC π^+ production, data and simulation. From Ref. [107].

Coherent π^0 production

The coherent reaction

$$\nu + C \rightarrow \nu + C + \pi^0 \quad (7.10)$$

is modeled in **nuance** following Ref. [111]. However, there is little experimental data for coherent π^0 production below 2 GeV, and calculations suggest that nuclear effects can significantly reduce the cross section [112] [113]. The forward-peaked angular distribution of coherently produced π^0 's allows us to use our π^0 event sample to tune the rate of this reaction in **nuance** (Chapter 11).

7.4 Final state interactions

A particle produced within a nucleus in a neutrino interaction may interact as it exits the nucleus, possibly modifying the final state. For example, a π^+ may undergo charge exchange with a neutron as it exits, appearing in the detector as a π^0 . A proton may induce the production of a pion in a collision with another nucleon. **Nuance** predicts rates

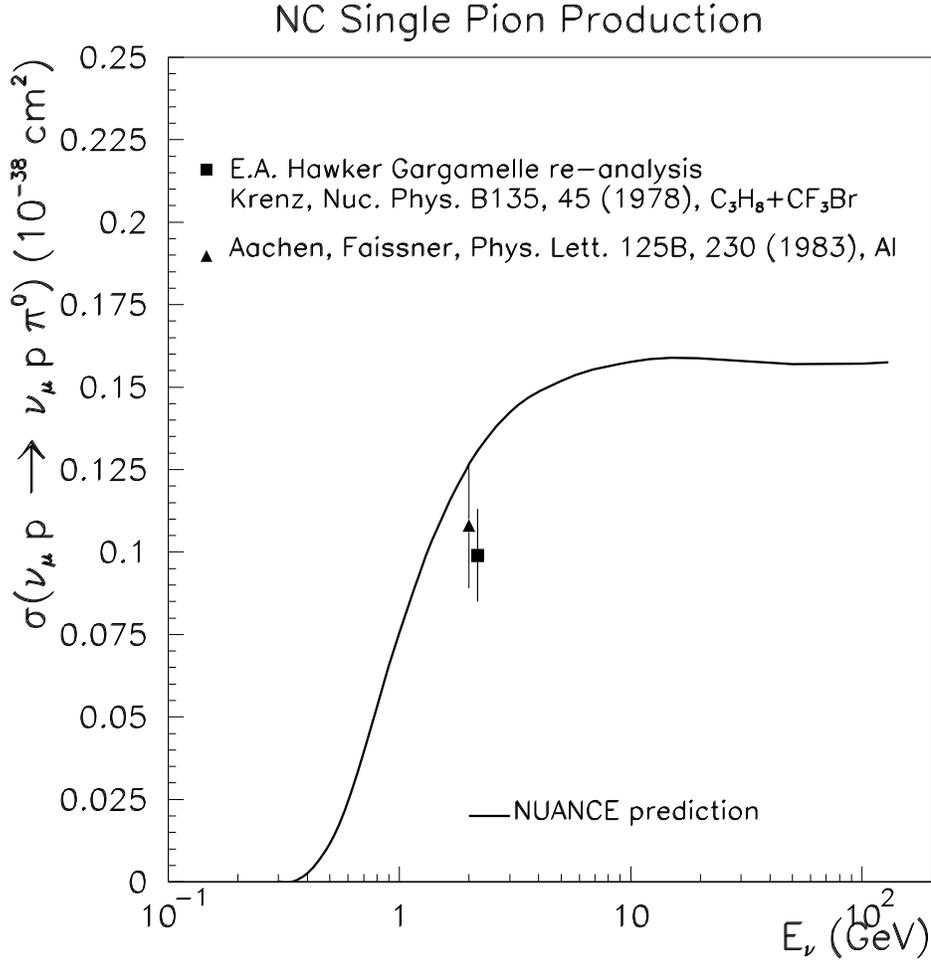


Figure 7.4: Measurements of the resonant π^0 production cross section and the free nucleon nuance prediction [109] [110]. No measurements exist below $E_\nu = 2$ GeV.

of $\sim 20\%$ for pion absorption and $\sim 10\%$ for charge exchange. Proton and neutron final state interactions lead to pions in 1 – 2% of originally pion-free events.

To simulate these effects, **nuance** tracks final state hadrons through the nuclear medium in 0.2 fm steps. Compiled π - N and N - N cross sections provide the interaction probabilities. The nucleon spatial distribution in **nuance** was updated for MiniBooNE using Ref. [114]. Additionally, roughly half of events include few-MeV de-excitation photons which are insignificant perturbations.

7.5 Miscellany

The remaining processes implemented in **nuance** play minor roles in the $\nu_\mu \rightarrow \nu_e$ analysis.

- NC elastic scattering ($\nu N \rightarrow \nu N$) is treated in analogy with CC QE scattering (bare

cross section embedded in an RFG model).

- Neutrino-electron scattering ($\nu e^- \rightarrow \nu e^-$, $\bar{\nu}_e e^- \rightarrow \bar{\nu}_e e^-$, etc.) is included at tree level.
- Deep inelastic scattering (DIS) and various multi- π channels are modeled.

All **nuance** calculations are performed under the assumption of a CH₂ target. Reaction rates are scaled with the density of the target material (*e.g.*, steel). This simplification has negligible effect on the analysis.

7.6 Systematic uncertainties

The systematic uncertainties assessed in the cross section simulation are as follows.

- Uncertainties in M_A and κ come from the fits described in Chapter 10. These parameters affect the normalization and Q^2 dependence of the quasi-elastic cross section. An additional 10% normalization uncertainty is also applied [115].
- Uncertainties in single- and multi-pion production come from available data at Mini-BooNE energies [116]. These uncertainties are assessed using the axial vector masses in **nuance** (M_A^π and $M_A^{\pi\pi}$). See Table 7.1 for values.
- Uncertainties in final state interactions enter through the absorption and charge exchange cross sections used in **nuance**. These uncertainties have the same effect on the analysis – namely, making the level of π^+ background in the ν_μ CC QE sample uncertain – as do uncertainties in pion propagation through the detector (§8.2). For simplicity, these are combined here, giving a total absorption (charge exchange) cross section uncertainty of 50% (35%).
- The parameter Δs , which governs the strange quark contribution to neutral current scattering, is set with external data: $\Delta s = 0.0 \pm 0.1$ [107].
- Deep inelastic scattering is given a 25% uncertainty [116]. Neither this nor the previous item plays an important role in the analysis.

M_A^{QE}	(1.15 ± 0.07) GeV
M_A^π	(1.1 ± 0.275) GeV
$M_A^{\pi\pi}$	(1.3 ± 0.5) GeV
M_A^{coh}	(1.03 ± 0.275) GeV

Table 7.1: Axial vector masses in **nuance**. $M_A^{\text{QE}} \equiv M_A$ and M_A^{coh} are established in Chapters 10 and 11; the others come from available pion production data [116]. The M_A^{coh} and M_A^π uncertainties are fully correlated. (See text.)

- Radiative Δ decay has a 9% uncertainty (§11.2.3) applied on top of the already assessed Δ production uncertainty.

Since we will use our ν_μ sample to constrain the ν_e quasi-elastic cross section (Chapter 10), uncertainties in the ν_e/ν_μ cross section ratio, which are driven by nuclear effects, are also relevant. This ratio was examined under nuclear model variations, including superscaling and spectral function calculations, to establish an energy-dependent ν_e/ν_μ uncertainty that is assessed on the ν_e quasi-elastic cross section. This contributes negligible error above $E_\nu \sim 400$ MeV and reaches 10% at $E_\nu \sim 200$ MeV [115]. Finally, nucleon form factor variations induce small changes in the energy dependence of the quasi-elastic cross section, motivating an additional energy-dependent uncertainty assessed on quasi-elastic events. The resulting uncertainty in the ratio $\sigma_{\text{QE}}(E_\nu=2.5 \text{ GeV})/\sigma_{\text{QE}}(E_\nu=0.3 \text{ GeV})$ is 13%. See Ref. [115] for more.

Neutral current π^0 's are discussed in Chapter 11 where we measure their momentum spectrum. The effects of the above uncertainties on π^0 momenta are removed during error assessment while maintaining effects on other π^0 distributions (in particular, the angular distribution). Additionally, the coherent π^0 production rate, controlled with M_A^{coh} , is taken to be fully correlated with M_A^π so that the resonant/coherent π^0 ratio stays fixed at the observed value (§11.2.2).

Chapter 8

Detector simulation

We use a GEANT3-based [71] Monte Carlo simulation of the detector up through photo-electron production. The PMT response, electronics, and DAQ are modeled in separate code. These simulation components are described below.

8.1 Geometry definition

The MiniBooNE tank is modeled as a spherical shell of iron capped with the “tophat” cylindrical hatch. This iron shell sits in an air-filled cylindrical vault which sits in a larger cylinder of dirt. Atop this larger cylinder is a conical frustum of dirt representing the overburden. The cylindrical electronics room sits inside the overburden and contains the muon tracker. Figure 8.1 shows these geometry components.

The tank volume is filled with oil which in turn contains the following.¹

- aluminum optical barrier
- four spherical laser flasks
- seven scintillator cubes and the single “coffee can” light shield
- lat pipes
- main PMTs
- veto PMTs

The PMT globes are created by joining two partial spheres symmetrically at the globe’s equatorial plane such that the final globe diameter and height are correct. Below each globe is a cylinder and thin disc representing the PMT neck and the base electronics.

Default GEANT3 material definitions are used when available (*e.g.*, air). Dirt is modeled as a silicon-oxygen-aluminum mixture with relative atomic contributions 20-65-15 and with density $\rho_{\text{dirt}} = 2.15 \text{ g/cm}^3$. The oil is modeled as CH_2 with density $\rho_{\text{oil}} = 0.855 \text{ g/cm}^3$ and with customized optical properties.

¹See Chapter 3 for definitions.

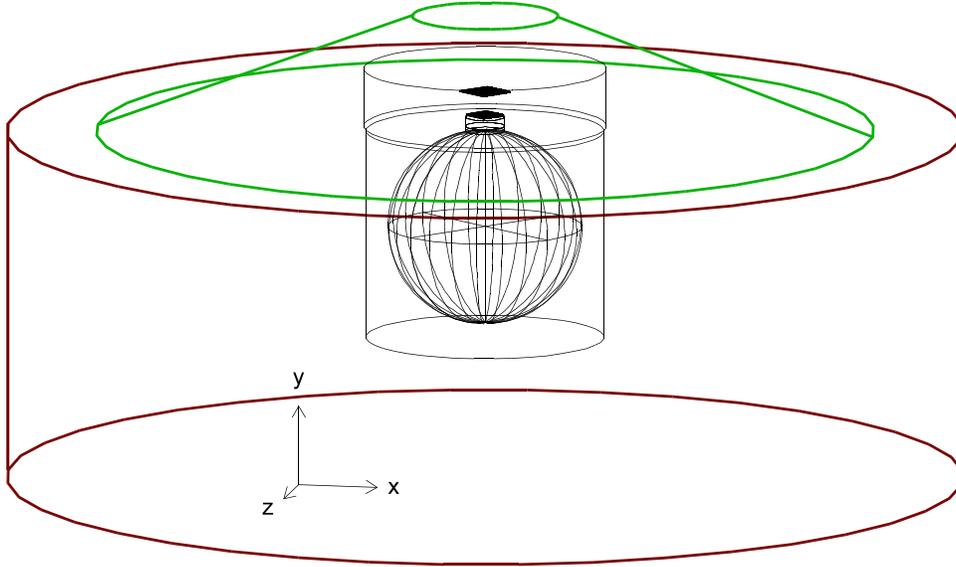


Figure 8.1: The GEANT3 geometry for the dirt and overburden (thick-lined cylinder and conical frustum), the electronics room and vault (thin-lined cylinders), the detector tank and tophat (sphere and small cylinder), and the muon tracker strips (small planes above tophat).

8.2 Particle propagation

Standard GEANT3 particle propagation and decay routines are used except as follows.

Decay matrix elements

The GEANT3 decay routines do not include matrix elements. Custom routines are used to correct this for two relevant decays: (1) $\pi^0 \rightarrow \gamma e^+ e^-$ (B.R.=1.2%), which is relevant as a ν_e background, and (2) $\mu \rightarrow e \nu \nu$, which is relevant as a muon tag and an energy calibration source.

μ^- capture

A stopped μ^- in mineral oil may be captured by carbon² via



The μ^- lifetime in carbon, $\tau_{\text{C}} = 2026.3 \pm 1.5$ ns [118], is thus smaller than the vacuum lifetime of $\tau_{\text{vac.}} = 2197.03 \pm 0.04$ ns [9]. These lifetimes imply a branching ratio for capture

$$\text{B.R.}(\text{capture}) = 1 - \frac{\tau_{\text{C}}}{\tau_{\text{vac.}}} = 0.0777 \pm 0.0007 . \quad (8.2)$$

²A μ^- that enters a hydrogen orbital upon stopping is quickly transferred to a carbon orbital, and thus μ^- capture on hydrogen is negligible [117].

The Monte Carlo simulation includes μ^- capture and the associated reduction in the μ^- lifetime. Additionally, the de-excitation of the induced B^* state is modeled, although the photons and neutrons emitted are rarely above the detection threshold, making μ^- capture essentially invisible.

Hadronic cross sections

The GCALOR [119] hadronic interaction package is used in lieu of the default GFLUKA package. The relevant differences are that

- π^+ charge exchange (CX) occurs too often in GFLUKA;
- π^+ absorption (ABS) is nearly absent in GFLUKA.

These statements are quantified in Table 8.1. Uncertainties in these cross sections, which affect π^+ backgrounds in the ν_μ CC QE sample (Chapter 10) but do little else, are handled through the nuance FSI model (§7.6).

8.3 Optical photons

8.3.1 Production

Charged particles propagating through the oil produce optical photons which are individually tracked in the Monte Carlo. The light is produced via three mechanisms.

1. **Cherenkov radiation.** This familiar phenomenon is implemented in GEANT3 as described in Ref. [120]. The requisite index of refraction is discussed below. A scale factor $f_{\text{Ch}}=1.106$ multiplies the Cherenkov production, although this should be thought of as an adjustment to propagation and detection efficiencies rather than as a scaling of the Cherenkov production itself.
2. **Scintillation.** Energy deposited in mineral oil by charged particles leads to scintillation light. Scintillation photons are created in the simulation according to

$$\frac{dN_{\text{sci}}}{dE} = \frac{31.64 \text{ MeV}^{-1}}{1 + B_1 \left(\frac{1}{\rho_{\text{oil}}} \frac{dE}{dx} \right) + B_2 \left(\frac{1}{\rho_{\text{oil}}} \frac{dE}{dx} \right)^2} \quad (8.3)$$

where the $\frac{dE}{dx}$ dependence handles saturation of the scintillator (Birks' law [121]) and where the coefficients are given by

$$B_1 = 0.014 \text{ g MeV}^{-1} \text{ cm}^{-2} \quad (8.4)$$

$$B_2 = 0.0 \text{ g}^2 \text{ MeV}^{-2} \text{ cm}^{-4} . \quad (8.5)$$

B_2 , though zero, is included so that it can be varied as part of the systematic error assessment. Scintillation emission is delayed exponentially with a time constant $\tau=34$ ns. Wavelengths are discussed in §8.3.2.

	GFLUKA	GCALOR	Ashery <i>et al.</i>
σ_{CX} :	118 ± 6 mb	59 ± 4 mb	45 ± 23 mb
σ_{ABS} :	< 3 mb	132 ± 7 mb	157 ± 37 mb

Table 8.1: Comparisons between GFLUKA, GCALOR, and measured π^+ hadronic cross sections at $KE_\pi = 205$ MeV. The experimental data come from Ref. [122].

- UV fluorescence.** Photons with wavelengths below $\lambda=250$ nm are not propagated in the simulation. (The PMTs are insensitive below $\lambda \approx 280$ nm.) However, a UV photon can induce fluorescence (see below), possibly resulting in a photon with a detectable wavelength. This possibility is accounted for by the creation of “UV fluorescence” photons, identical in spectrum and time constant to scintillation light, to represent the fluorescence daughters of the otherwise ignored UV light. Since Cherenkov radiation is the only source of $\lambda < 250$ nm light, these photons are generated in each Monte Carlo step in proportion to Cherenkov production:

$$dN_{\text{uvf}} = f_{\text{uvf}} dN_{\text{Ch}} \quad (8.6)$$

with $f_{\text{uvf}}=0.074$. Since attenuation lengths are < 1 cm at UV wavelengths, and because 1 cm is much smaller than the spatial scales relevant in the detector, UV fluorescence photons are produced along the track (rather than, more realistically, a short distance away).

8.3.2 Propagation

Each photon is stepped through the detector until it is absorbed. The absorption may occur within the oil, on various surfaces, or at a PMT photocathode, possibly creating a photoelectron. Figure 8.2 shows the rates of some of the relevant transport processes as a function of wavelength.

Scattering

Local variation in the index of refraction of a fluid due to thermal fluctuations leads to photon scattering. A complete description of the phenomenon includes the familiar Rayleigh scattering contribution (which assumes certain symmetries about the density fluctuations) and additional “depolarization” scattering components which stem from anisotropic fluctuations [123] [124]. We measured the rates and angular distributions of these scattering components at $\lambda=442$ nm and $\lambda=532$ nm using a Brookhaven BIS-200 goniometer³ under four permutations of incoming and outgoing polarization [125]. To obtain absolute scattering rates, the intensity of scattered light was calibrated using suspensions of NIST-certified

³We thank Wesley Kopacka for arranging use of the apparatus.

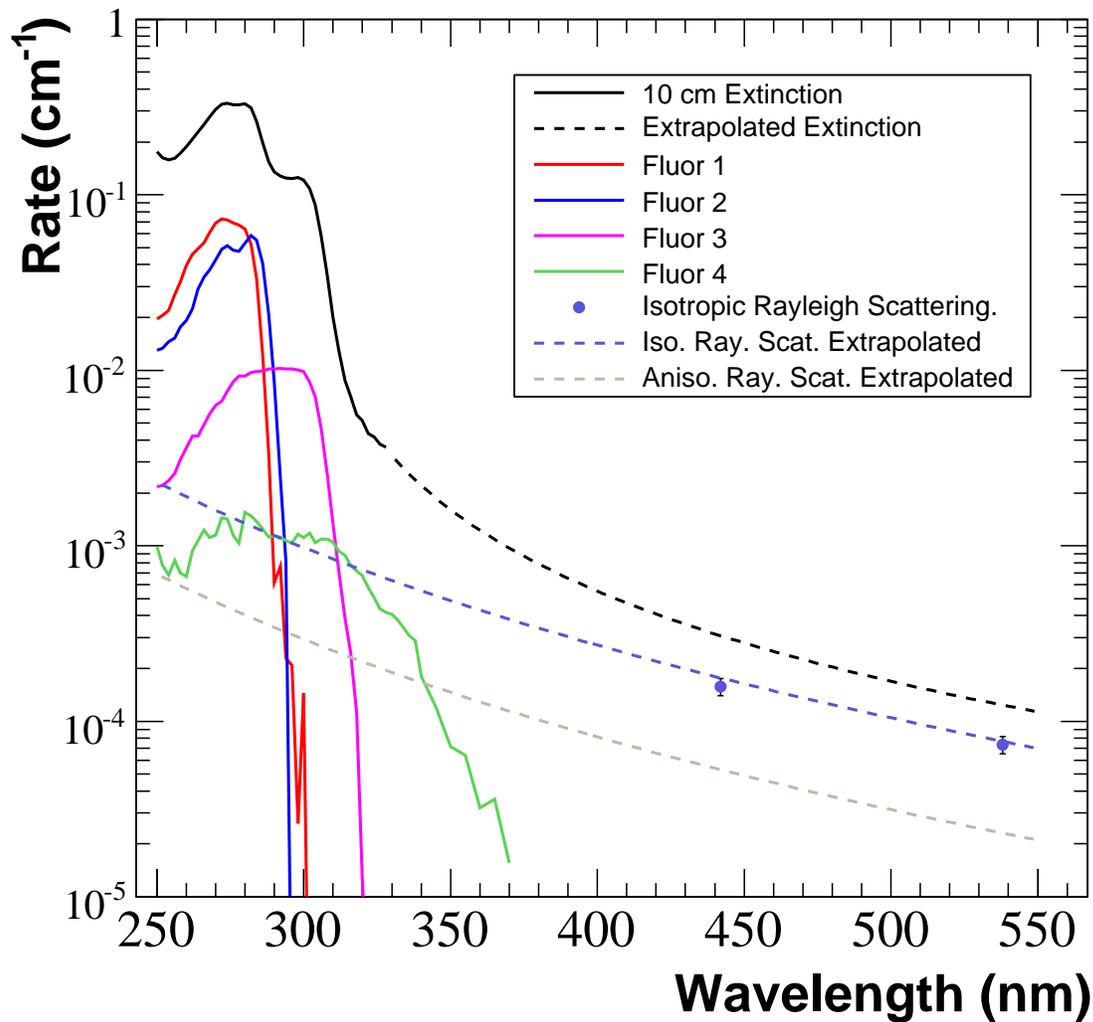


Figure 8.2: Rates of optical photon processes. The “10 cm extinction” curve shows the extinction rate measured in a 10 cm cell at FNAL. Consistent measurements were obtained from 1, 2, and 5 cm cells at FNAL and 1 cm cells at JHU. Extinction includes any process that interrupts a photon’s forward propagation (namely: absorption, scattering, and fluorescence.) The extrapolated extinction curve is derived using laser flask and Michel electron events. Items 3-6 in the legend show fluorescence rates. Finally, two Rayleigh scattering measurements are shown along with extrapolations of the isotropic and anisotropic components (§8.3.2).

50 nm polystyrene spheres⁴. Suspensions were prepared at various dilutions, and the Mie scattering rates were calculated to provide reference values. Wavelength extrapolation is carried out with a theoretically motivated λ^4 rate dependence. Figure 8.3 shows example angle and rate measurements.

Raman scattering was separately measured [126]. It contributes 5% to the total scattering rate and induces a small wavenumber shift of $\Delta\nu=2890 \text{ cm}^{-1}$. All of these scattering mechanisms are included in the simulation.

Fluorescence

Short wavelength photons can excite molecular states that decay to produce longer wavelength photons. A fluorescence analysis of Marcol 7 oil performed by Dmitri Toptygin (Johns Hopkins University) identified four distinct fluors, each with its own emission spectrum, excitation spectrum, and emission time constant, as shown in Figure 8.4. These fluors are simulated [127].

Returning to scintillation: we assume that scintillation light is produced by these same fluors. The relative contribution from each fluor is set by comparing data and Monte Carlo predictions for Michel electron and neutral current elastic scattering (proton) events. The data prefer a model in which scintillation comes from fluor 4 only (as labeled in Figure 8.4).

Absorption in oil

The scattering and fluorescence rates are subtracted from the total extinction rate [126] to obtain a residual that is attributed to absorption.

Reflections

Within the main detector region, most surfaces are modeled as perfectly absorbing. The two exceptions are (1) the lower halves of the PMT globes, which are silvered in reality and which produce perfect specular⁵ reflections in the simulation, and (2) the inner surface of the optical barrier, the reflection probability of which was measured (in air) to be $\sim 2.5\%$. We set it to 5% in the simulation to account for an additional $\sim 2.5\%$ reflection probability expected from the front faces of the PMTs, which cover 10% of the optical barrier’s surface. (The PMT faces themselves are modeled as perfectly absorbing.)

In the veto region, the white steel tank and white optical barrier surface had their wavelength-dependent albedos measured in air. These albedos are used in the simulation, though scaled by a factor $f_{\text{albedo}} = 0.905$ tuned by comparing predicted and observed veto hit multiplicities for through-going muons. Diffuse reflections (polish=0) are assumed for these surfaces. Also in the veto, the lower PMT globes are perfect specular reflectors and the PMT necks are perfect diffuse reflectors (to represent the white mountings that enclose them).

⁴Duke Scientific

⁵GEANT “polish” parameter set to unity

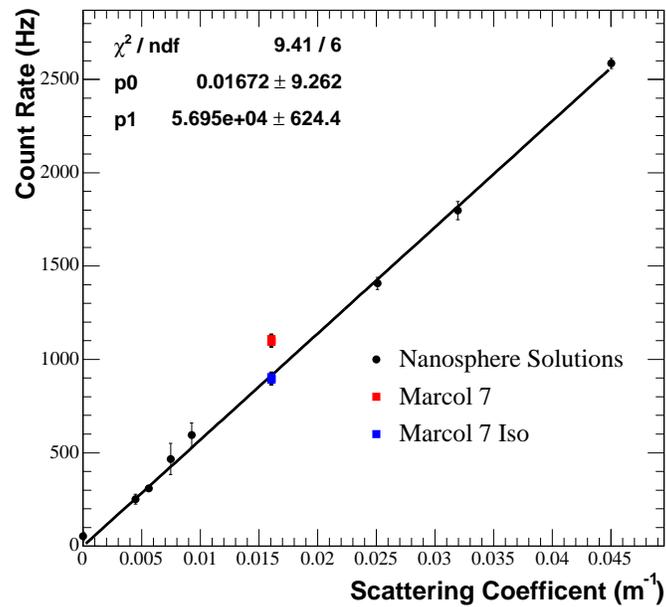
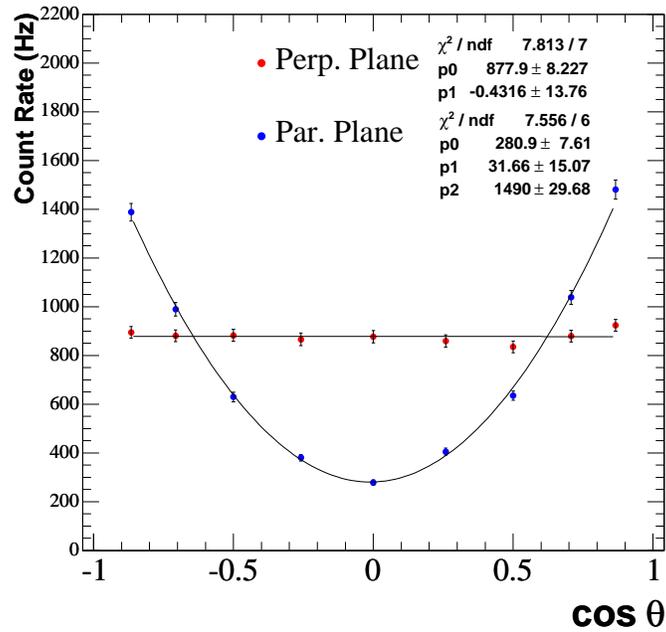


Figure 8.3: (Top) Angular dependence of the scattering rate of 442 nm light in planes perpendicular (red, flat) and parallel (blue, parabolic) to the incident polarization. (Bottom) Observed count rates for the “nanosphere” calibration suspensions (black), Marcol 7 oil (red), and Marcol 7 isotropic-only Rayleigh component (blue) plotted against the expected Mie scattering rate (nanospheres) or the expected isotropic-only Rayleigh rate (Marcol 7). The Marcol 7 “Iso” component is extracted from the total using the angular distributions above. The “Iso” point should (and does) lie on the nanosphere calibration curve. From Ref. [125].

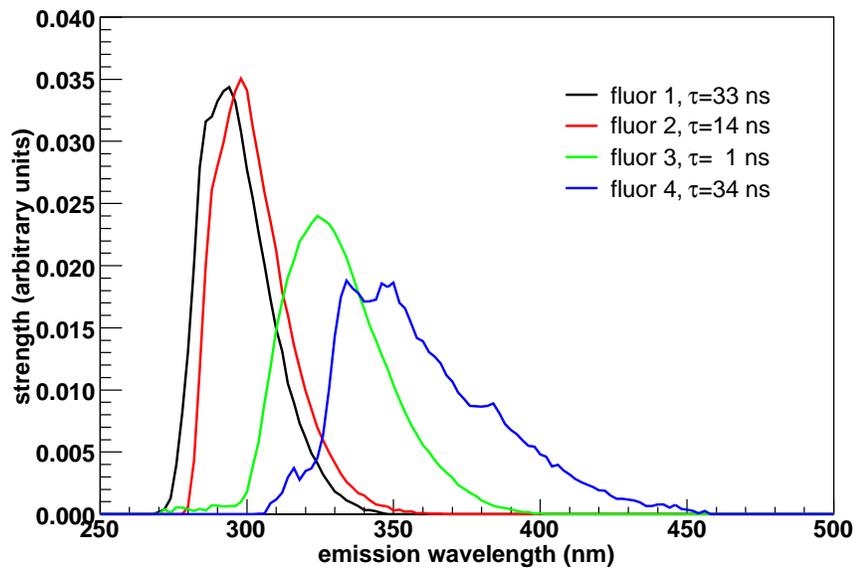
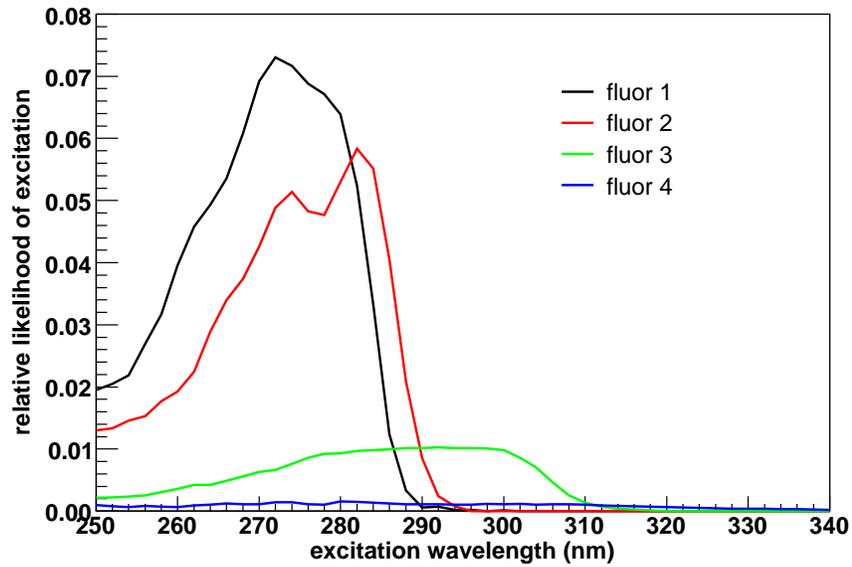


Figure 8.4: Fluorescence in Marcol 7 mineral oil. (Top) Relative excitation probabilities for the four principal fluors. (Bottom) Corresponding emission curves (arbitrarily normalized). Emission time constants are listed in the legend.

Index of refraction

The index of refraction of the oil was measured using an Abbé refractometer⁶ [128]. The data are well described by

$$n(\lambda, T) = \left[n_D + B \left(\frac{1}{\lambda^2} - \frac{1}{\lambda_D^2} \right) \right] [1 - \beta (T - T_0)] , \quad (8.7)$$

where λ is the wavelength and T is the oil temperature. Parameter values are given in Table 8.2. The consequent phase velocity c/n and dispersion relation $n(\omega)$, with $\omega \equiv \frac{2\pi c}{\lambda}$, are used to calculate group velocities relevant for photon propagation:

$$v_g = \frac{v_p}{1 + \frac{\omega}{n} \frac{dn}{d\omega}} . \quad (8.8)$$

At $\lambda=400$ nm and $T = 17.5$ °C, $v_g=(19.50 \pm 0.02)$ cm/ns.

8.3.3 Detection

The efficiency for photoelectron creation in the PMTs is shown in Figure 8.5. The curve comes from Hamamatsu measurements adjusted to account for reflections at the air/glass interface (which becomes an oil/glass interface in MiniBooNE). The R1408 efficiency is 0.83 times that of the R5912 efficiency, a ratio measured in the detector with Michel electron events. To minimize CPU time spent tracking photons that are destined to be undetected, the quantum efficiency is scaled in the Monte Carlo so that the peak efficiency is unity; the original peak efficiency of 22% is assessed at photon production.

The dependence of the PMT efficiency on the photon incident angle was measured using a sub-nanosecond pulsed LED light source⁷ directed at PMTs mounted on a rotatable support structure submerged in an oil bath [129]. The purely geometric contribution to the angular efficiency is removed and the residual gets applied to each incident photon in the simulation.⁸

n_D	1.4684 ± 0.0002
B	$(4240 \pm 157) \text{ nm}^2$
β	$(3.66 \pm 0.04) \times 10^{-4} \text{ (}^\circ\text{C)}^{-1}$
λ_D	589.3 nm
T_0	20.0°C

Table 8.2: Parameters describing the index of refraction of MiniBooNE’s mineral oil. The first three are the actual measured values. λ_D and T_0 are reference quantities [128].

⁶Xintian Fine Optical Instrument Corp., Model WY1A

⁷PicoQuant PLS 450.

⁸We use a simplified PMT geometry in the code, so the geometric contribution is not entirely removed here. Rather, the angular dependence is set such that it conspires with the actual simulated PMT shape to yield the correct efficiencies.

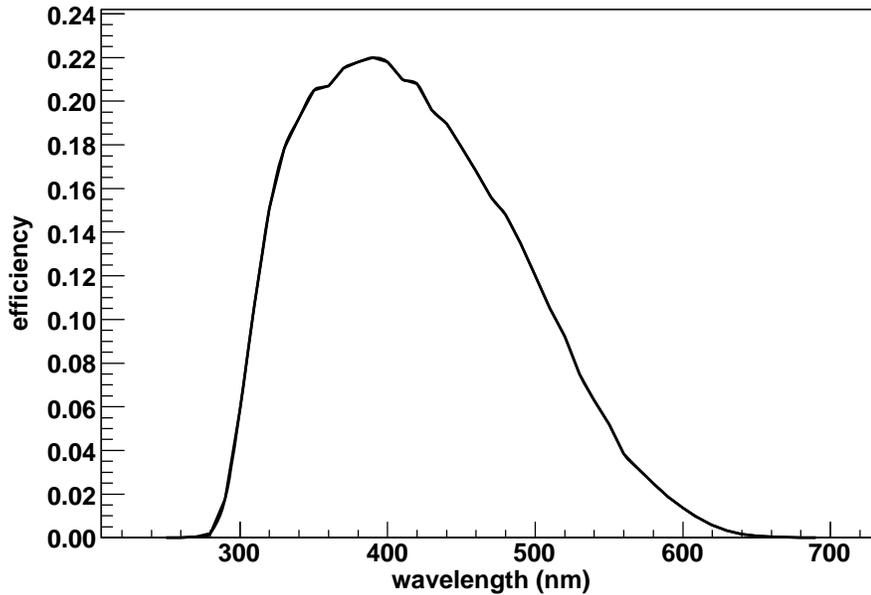


Figure 8.5: R5912 PMT quantum efficiency in oil.

8.4 PMT response

Time

An electron cascade induced in a PMT by a photoelectron (PE) travels down the dynode stack in a time interval that varies from one PE to the next. This transit time spread is the leading contribution to the PMT timing resolutions. Photoelectrons can also induce late-pulses ($\sim 5\%$) and pre-pulses (see below), in which the anode signal appears later or earlier than expected. Figure 8.6 shows these and other features.

The timing distributions for several PMTs were measured with the LED setup of §8.3.3. Figure 8.7 shows observed distributions overlaid with parametrizations that, with the following modifications, are used in the simulation to smear the times of PMT hits.

- The width of the prompt peak is tuned using Michel electron events.
- Due to differences in applied high voltage, the time gap between the prompt peak and the late-pulse “peak” varies from channel to channel. This time gap is measured for each main PMT using laser flask events, and the time response is adjusted in the simulation to reproduce these channel-to-channel variations.
- Pre-pulse timing features are included. Region B of Figure 4.12 reveals the existence of an early timing peak near $t = -18$ ns that contains hits with lower charges than those in the prompt peak (Region C). Not indicated in the figure is another low-charge region around $t = -2$ ns. Pre-pulsing is simulated by shifting the time and reducing

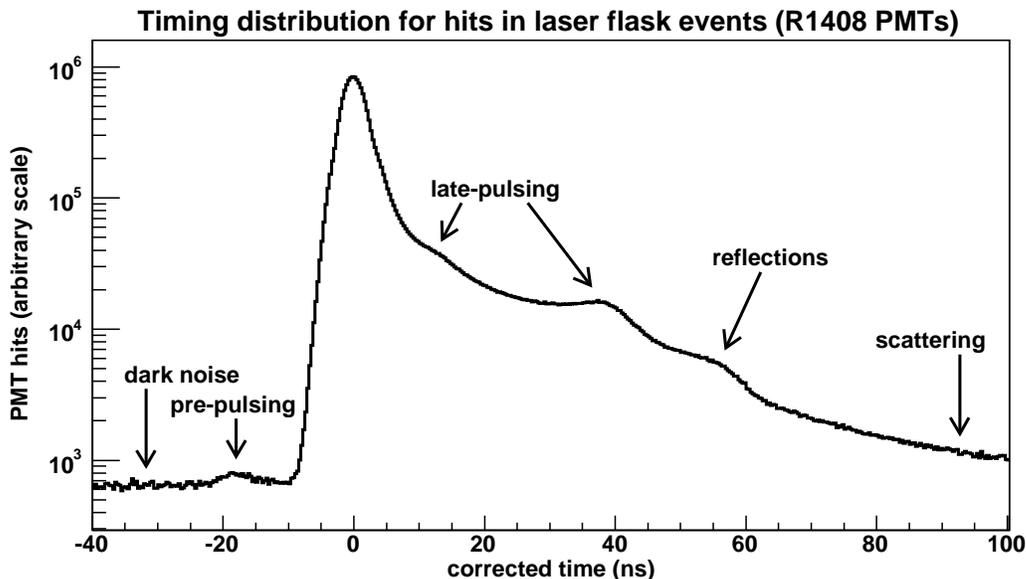


Figure 8.6: The time profile of R1408 PMT hits in low intensity laser flask events. Note the logarithmic vertical scale (*cf.* Figure 4.10). The prompt peak has been placed at $t = 0$ ns.

the charge of some fraction of photoelectrons. The -18 ns and -2 ns features are created using time shifts of those sizes, charge scalings of $\sim \frac{1}{5}$, and per-PE rates of 1.3×10^{-4} (former) and 0.08 (latter).

In addition to these smearing effects, each photoelectron's time is shifted by that channel's Δt_{offset} , as defined in §4.1.

Charge

The anode charge for each PE is drawn from a p.d.f. obtained from prompt laser flask PMT hits. R1408 and R5912 PMTs have different charge responses as Figure 8.8 indicates. The low-charge region of each distribution is artificially reduced to remove pre-pulse contamination that gets reintroduced by the Monte Carlo code.

8.5 Digitization

(Familiarity with §3.5 is assumed here.) Once each PE has acquired its smeared time and charge, a triangular anode pulse is created. The pulse has a width motivated by oscilloscope traces and an area driven by the charge drawn for the PE. The pulse shape also controls the phenomenon of time slewing. To achieve agreement between predicted and observed slew rates in cube-stopping cosmic muons, the pulse shape model was extended as follows. (R5912 values are in parentheses.) For PEs with drawn charge above $Q_{\text{co}}=0.87$ (1.2), the triangle's width at baseline is chosen from Gaussian(39.0 ns, 4.2 ns) (Gaussian(33.0 ns,

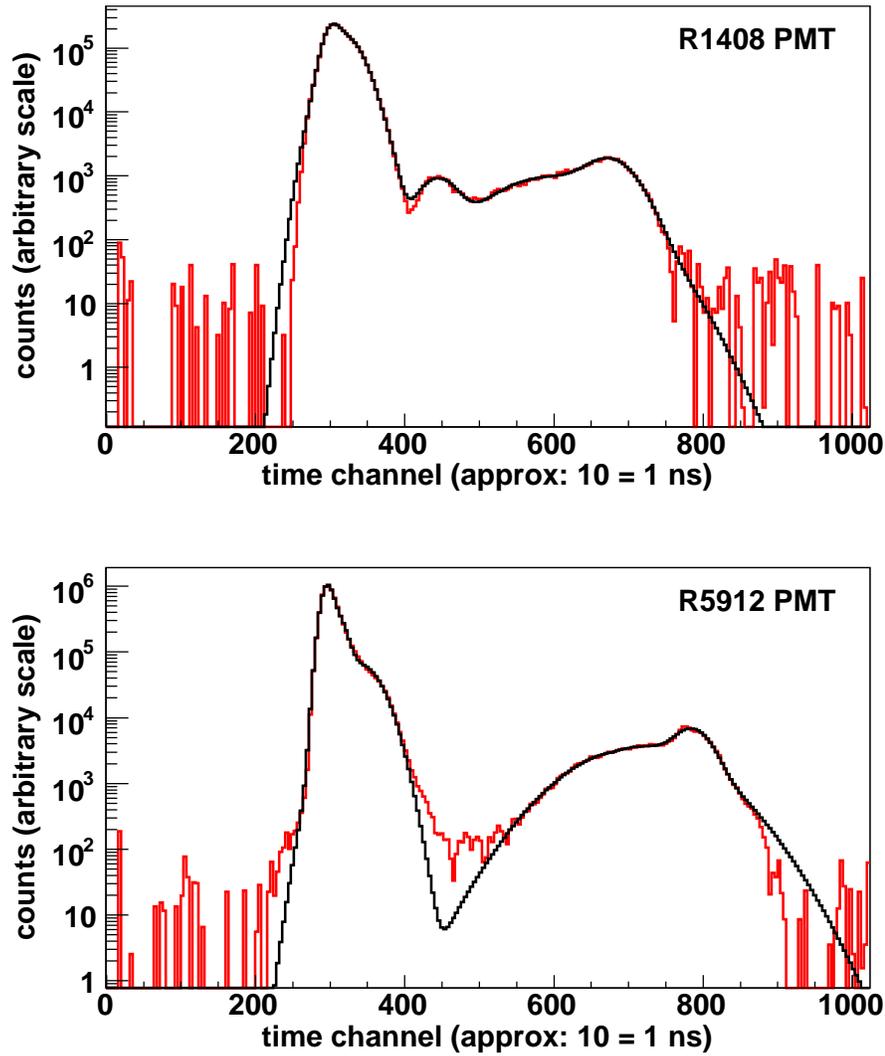


Figure 8.7: Measured PMT timing profiles. The parametrizations shown (black curves) are modified as described in the text to provide time smearing in the simulation. These measurements could not detect pre-pulsing, which gives charges typically below the apparatus's threshold.

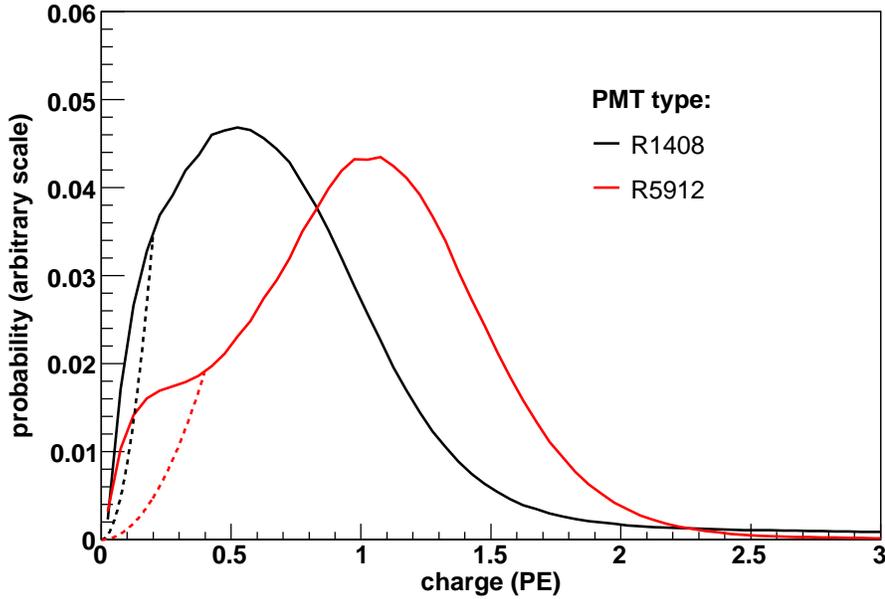


Figure 8.8: PMT charge response. The distributions indicated with the solid lines have means of 1 PE. For the older generation R1408 PMTs, a long tail in the charge distribution (reaching beyond 10 PE) balances the fact that the peak sits below 1 PE. The dotted lines represent the low-charge reduction mentioned in the text. Pre-pulsing restores this carved out portion.

3.6 ns)). The pulse height is set in proportion to the PE's charge and in inverse proportion to the chosen width. Below Q_{co} , the Gaussian mean varies linearly toward a $Q=0$ value of 14.0 ns (14.0 ns). The Gaussian RMS scales along with the mean, bottoming out at 1.51 ns (1.53 ns).

Despite this complexity, one is left with the simple situation of having a collection of triangular anode pulses of various sizes, one for each photoelectron. These pulses are used for two purposes:

1. A train of Q_ADC digitizations is assembled by scaling the charge reference curve V_q^{ref} (§4.2.3) by the chosen charge for each PE and by subsequently summing the results from all PEs on a channel.
2. Separately, all anode pulses on a channel are summed. The resulting complete anode signal is discriminated at the 0.1 PE level to initiate time ramps from which a train of T_ADC digitizations is formed.

This description requires two final modifications:

- Residual data/simulation disagreement in time slewing motivates an additional empirical slewing adjustment. The time of each PE is shifted by a small amount that depends on the total charge in the vicinity (“vicinity”= ± 60 ns).

- High charge PMT hits exhibit nonlinear response. Figure 8.9 shows the nonlinearity mapping obtained from laser flask data and applied in the simulation.

Various non-PMT channels (beam trigger signal, muon tracker strips, etc.) are also given appropriate activity. The entire event is then shifted in time to mimic the $1.6 \mu\text{s}$ neutrino beam spill. Finally, the trains of $\{Q_ADCn\}$ and $\{T_ADCn\}$ digitizations are sent through a zero-suppression, quad-defining algorithm functionally identical to that in the actual DAQ software.

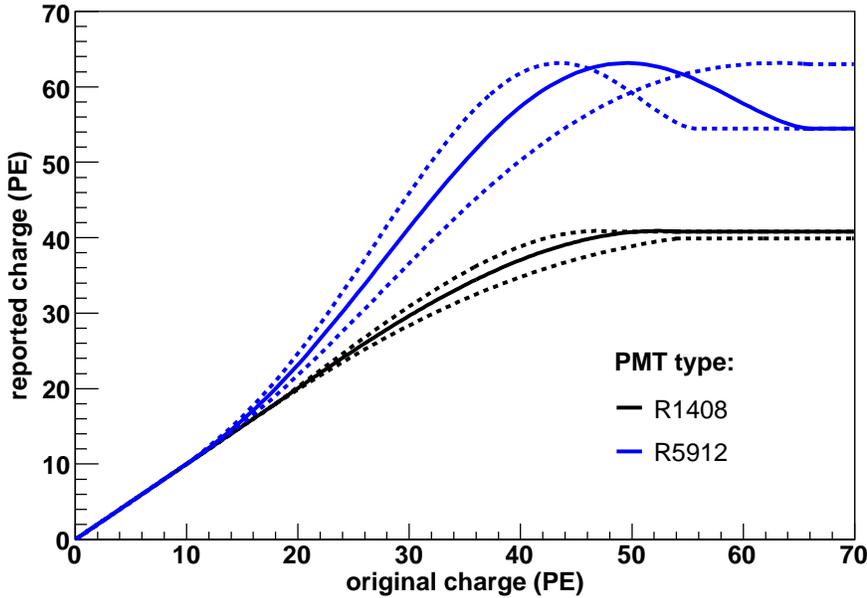


Figure 8.9: Nonlinear charge behavior. The solid curves show the mappings used to adjust hit charges in the simulation. These mappings were created using low- and high-intensity laser flask runs in which the laser’s intensity was throttled both via the injection current and through the use of neutral density filters. The extraction of these curves is sensitive to the assumed discriminator threshold values; the consequent uncertainty is indicated by the dashed curves ($\pm 1\sigma$).

8.6 Beam-off activity

As described in §3.7, “random” triggers are initiated at 2.01 Hz to record beam-uncorrelated activity in the detector, coming primarily from cosmic ray muons, their decay electrons, and PMT dark noise. To account for this activity in the simulation, PMT hits from the random triggers are merged with hits from the Monte Carlo events. The merging routine handles (rare) hit collisions via a quad-creating procedure akin to that used in the Monte Carlo itself for combining nearby PEs on a single channel.

8.7 Tuning and systematic uncertainties

The development of the detector simulation, and in particular its optical photon model, was guided by the measurements mentioned above and by comparisons of Monte Carlo predictions with detector data. For a handful of parameters, the laboratory measurements provide our best information. For most, the MiniBooNE detector provides the stronger (and sometimes only) parameter constraints.

A parameter that is more or less unrelated to any other can be adjusted by hand without difficulty. For example, the relative efficiencies of old and new PMTs can be set by counting hits in Michel electron events. Some observables, however, are affected by many parameters simultaneously. Reconstructed Michel electron energies, for instance, are governed by the Cherenkov yield, the extinction length, scattering rates, fluorescence rates, scintillation production, and more. One cannot arbitrarily adjust these to achieve Michel energy agreement since other observables – say, timing distributions – depend on them as well.

A careful by-hand approach did, over time, lead to reasonable default parameters. However, assigning correlated uncertainties to these is a more complex task for which the following procedure was developed.

Reducing parameter space

We begin by assigning initial parameter errors using the laboratory measurements alone – no detector data. Parameters for which no measurements are available are given enormous errors. For concreteness, the thirty-five parameters in play are:

- extinction length (5)
- Rayleigh/Raman scattering (3)
- refractive index (3)
- PMT angular efficiency (2)
- scintillation yield from each fluor (4)
- fluorescence yield from each fluor (4)
- UV fluorescence yield from each fluor (4)
- time constant of each fluor (4)
- Cherenkov scale factor (1)
- reflections (2)
- Birks' law coefficients (2)
- old/new PMT relative efficiency (1)

Other parameters exist in the code but do not affect the distributions below.

Our initial knowledge of these parameters can be cast as a multivariate Gaussian p.d.f. with mean $\hat{\mathbf{s}}$ and covariance matrix \mathbf{S} . Using this p.d.f., we draw $M = \mathcal{O}(3000)$ random Monte Carlo parameter sets $\{\mathbf{s}^{(1)}, \mathbf{s}^{(2)}, \dots, \mathbf{s}^{(M)}\}$, and we generate a sample of Michel electrons for each parameter set. Reconstructed energy distributions from four of these M samples are shown in the left column of Figure 8.10. The data/simulation agreement in these distribution varies from good to awful, indicating that our data can indeed constrain the model further.

For a perfect simulation, the χ^2 expectation from these comparisons is $N_{\text{dof}} = (\text{number of bins compared})$. An imperfect parameter set $\mathbf{s}^{(m)}$ will deviate from this, and $\Delta\chi^2 = \chi^2 - N_{\text{dof}}$ tells us how likely this parameter set is, relative to others, according to the data:

$$p = \exp\left(-\frac{1}{2}\Delta\chi^2\right). \quad (8.9)$$

The relative likelihood of obtaining parameter set $\mathbf{s}^{(m)}$ in the first place is

$$w = \exp\left(-\frac{1}{2}(\mathbf{s}^{(m)} - \hat{\mathbf{s}})^\top \mathbf{S}^{-1}(\mathbf{s}^{(m)} - \hat{\mathbf{s}})\right). \quad (8.10)$$

If our new information is more constraining than our old ($p < w$), we would like to adjust the relative likelihood of parameter set $\mathbf{s}^{(m)}$ from w to p . That is, we wish to take $w \mapsto \min(w, p)$, which we do by assigning an additional weight of

$$\eta = \min\left(1, \frac{p}{w}\right). \quad (8.11)$$

We construct from the M parameter sets a new p.d.f., with mean

$$\hat{\mathbf{s}}' = \frac{\sum_m \eta_m \mathbf{s}^{(m)}}{\sum_m \eta_m} \quad (8.12)$$

and covariance matrix

$$S'_{ij} = \frac{\sum_m \eta_m (s_i^{(m)} - \hat{s}'_i)(s_j^{(m)} - \hat{s}'_j)}{\frac{M-1}{M} \sum_m \eta_m}. \quad (8.13)$$

Using this new p.d.f., we repeat the above procedure with increased sample statistics and/or new test distributions (energy, timing, etc.) in the χ^2 calculation.⁹

About twenty iterations are carried out. We require $\sum_m \eta_m > 300$ after each iteration so that the new p.d.f. can be determined adequately. This sum can be increased by generating more random parameter sets or by comparing fewer distributions.

⁹Parameter bounds complicate this picture considerably, as the calculated means get biased away from the boundaries. This bias is compounded in each iteration. We recover the intended behavior by using a wholly different method for choosing the random parameter sets. Descriptions and toy Monte Carlo studies of all of this can be found in Ref. [130].

The right column of Figure 8.10 shows four reconstructed energy distributions drawn from a well-iterated parameter p.d.f. The discrepancies are noticeably reduced. Figure 8.11 shows a similar before/after comparison for a portion of the old PMT time distribution. Figure 8.12 shows several distributions at once, using a band to represent the spread of the M individual Monte Carlo histograms. The reduction in uncertainty is striking.

This procedure was performed using the following Michel electron observables:

- reconstructed energy spectrum
- late time histograms (50 ns to 120 ns after prompt) for old and new PMTs
- prompt and late angular profiles ($\cos \theta$)
- average reconstructed energy for inward-pointing, outward-pointing, and sideways-pointing electrons as a function of reconstructed radius.

A final timing distribution was built from neutral current elastic scattering events in order to disentangle scintillation and UV fluorescence light, which in our model cannot be distinguished with the highly relativistic Michel electrons.

Other parameters

- **Prompt timing.** The PMT prompt time response was not included above. Rather, the prompt timing parametrization established with laboratory measurements and adjusted with Michel electrons (§8.4) is given uncertainties based on residual sub-nanosecond discrepancies in the Michel electron and cube-stopping muon time distributions.
- **Particle propagation.** The cross section for pair production by high energy photons is given a 3% uncertainty based on limitations in the GEANT3 parametrization and its underlying data. Similar uncertainties are assigned to the bremsstrahlung (5%) and Compton scattering (5%) cross sections and to the average muon multiple scattering angle (4%).
- **GCALOR carbon excitation energy.** The energy given to low-energy gammas produced by the GCALOR nuclear excitation routines is given an uncertainty of 60% motivated by the observed ratio of prompt-to-late energy seen in neutral current elastic scattering events.

In the end, detector model errors have no significant effect on the oscillation sensitivity of this analysis.

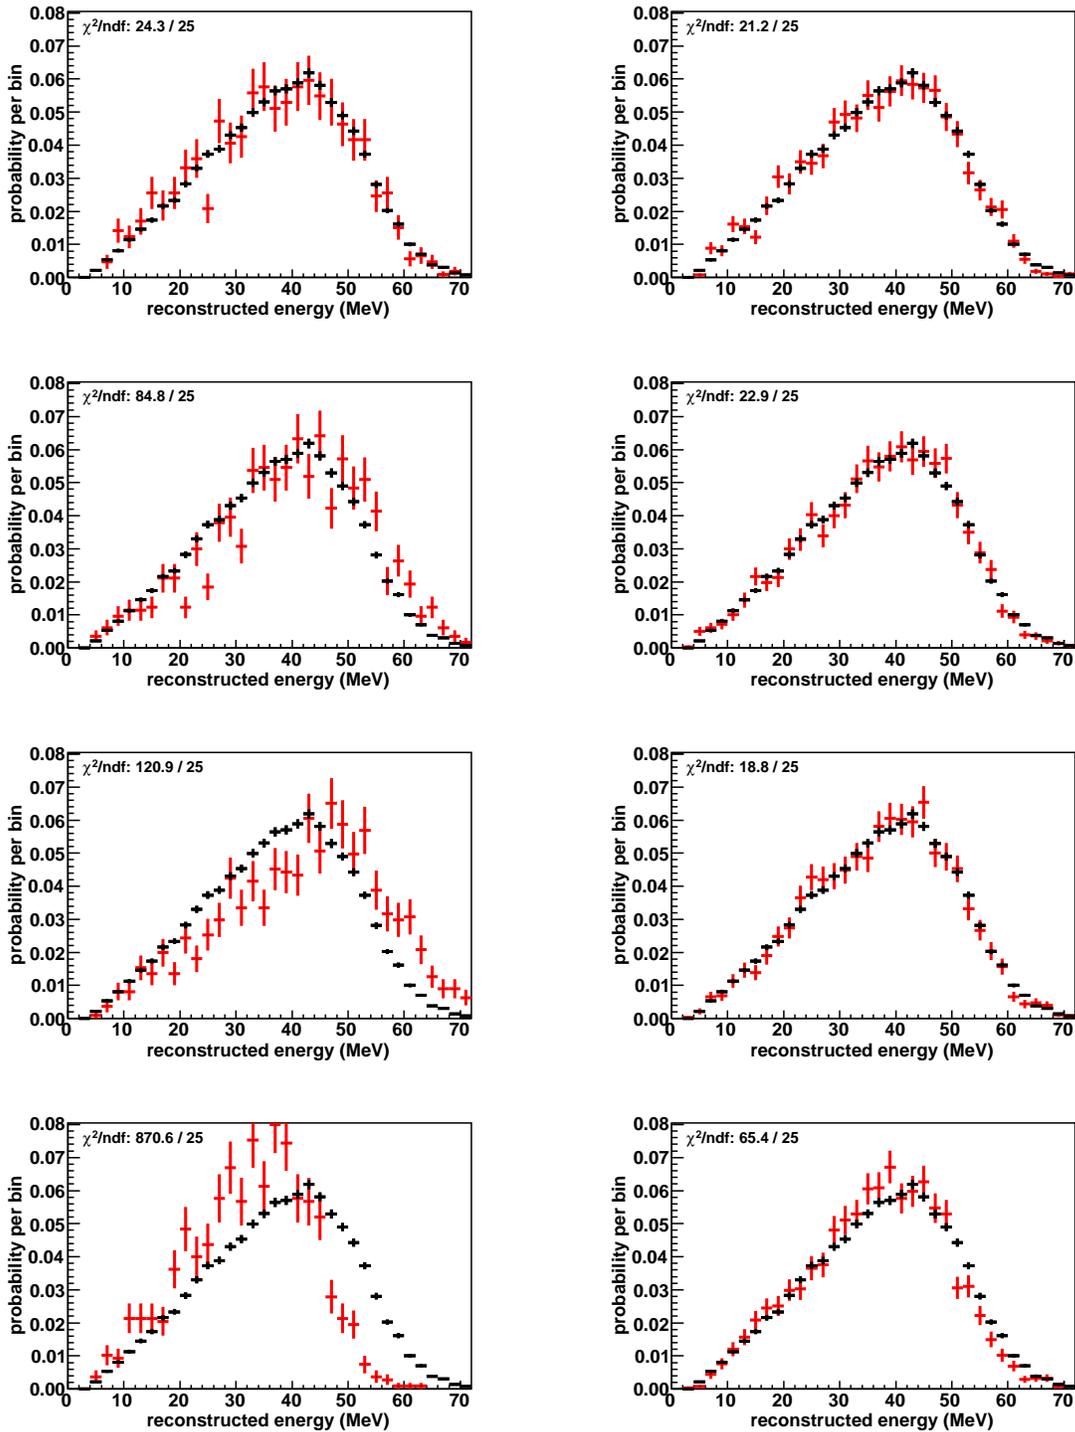


Figure 8.10: (Left column) Data/simulation comparisons (black/red) of the reconstructed Michel energy distribution for four Monte Carlo variants drawn from an early (not quite initial) parameter p.d.f. (Right column) The same after all iterations.

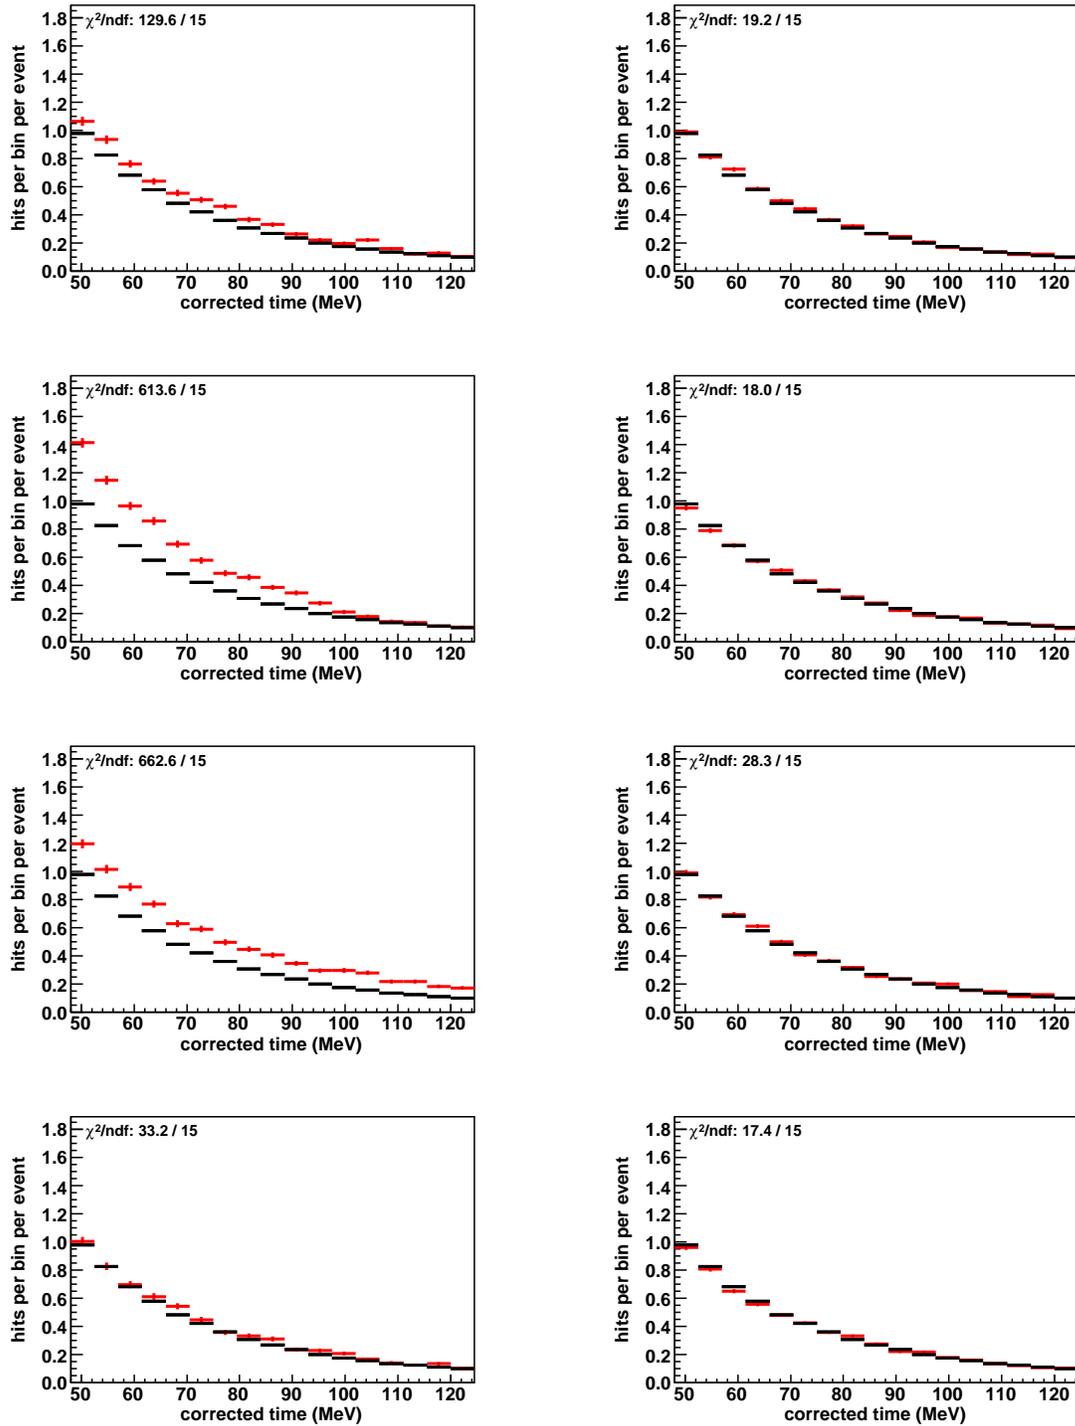


Figure 8.11: The timing distribution, from 50 ns to 120 ns, for old PMTs in Michel electron events. The plot conventions are as in Figure 8.10.

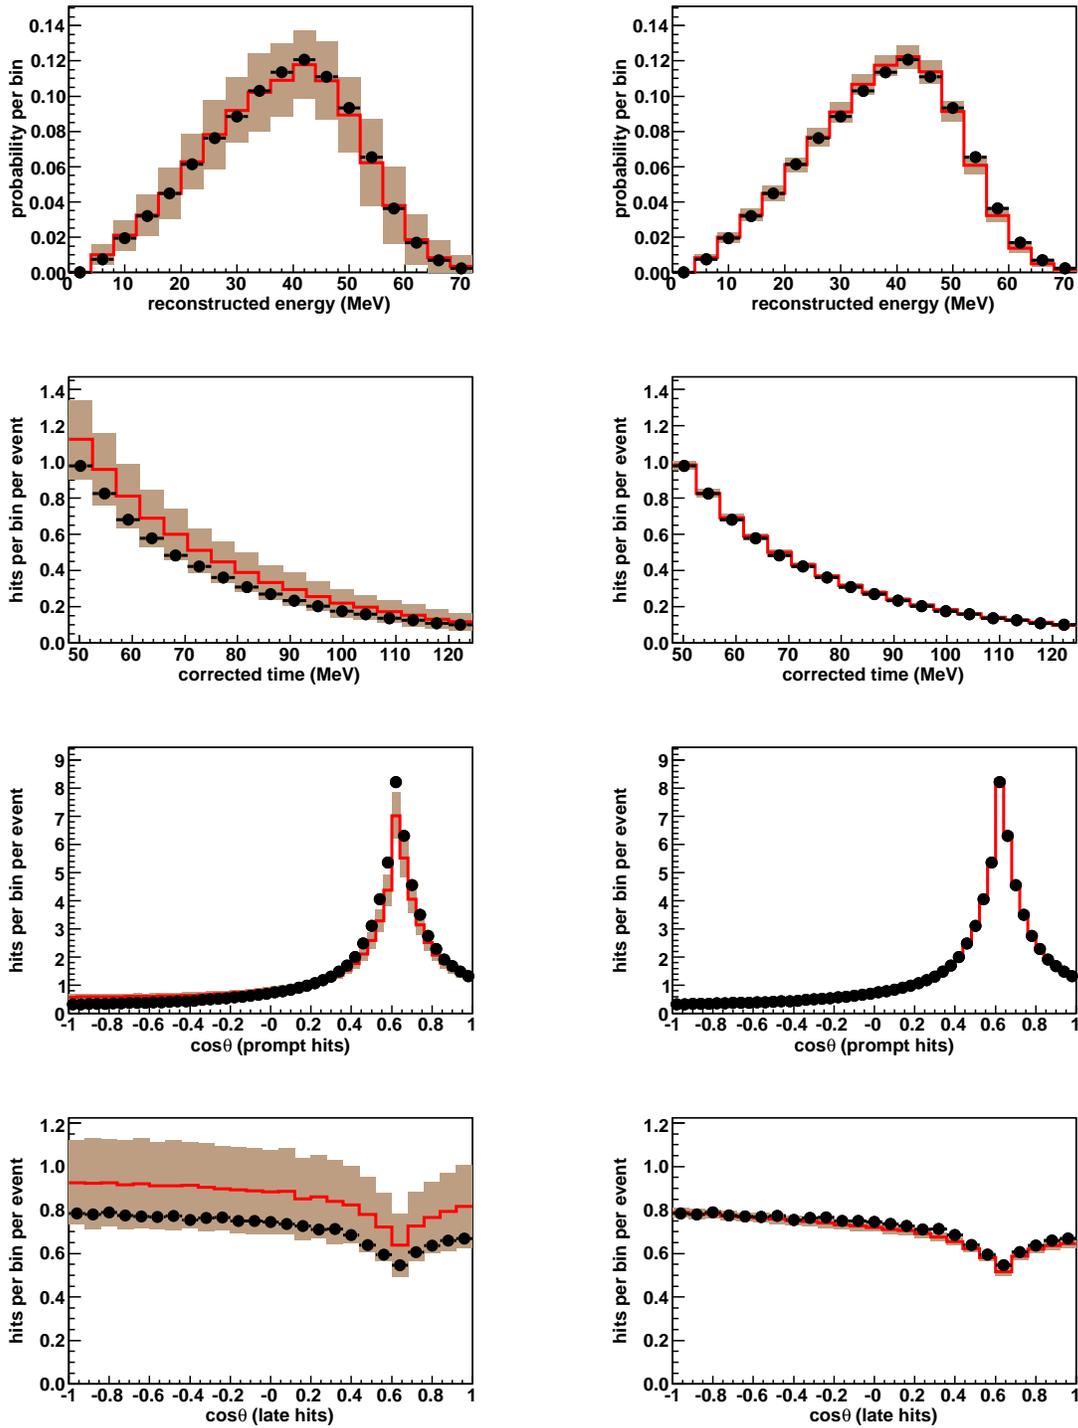


Figure 8.12: Data/simulation comparisons of four quantities using an early parameter p.d.f. (left column) and the final p.d.f. (right column). The error band indicates the spread of the Monte Carlo distributions given the parameter uncertainties.

Chapter 9

Event reconstruction

9.1 Introduction

The event reconstruction uses the method of maximum likelihood [9] [131] to estimate track parameters. We need seven parameters to describe a single particle track:

- starting point (x_0, y_0, z_0)
- starting time (T_0)
- direction (θ_0, ϕ_0)
- energy (E_0) .¹

We refer to this parameter set as \mathbf{x} . The observables in the likelihood are the PMT measurements. For each tube we have

- a bit indicating whether the tube was hit or unhit
- if the tube was hit, the charge of that hit
- if the tube was hit, the time of that hit.

Treating the PMTs as independent and treating each PMT's charge and time measurements as independent, we can immediately write down the likelihood that is to be maximized:

$$\mathcal{L}(\mathbf{x}) = \prod_{\text{unhit}} P(i \text{ unhit}; \mathbf{x}) \prod_{\text{hit}} P(i \text{ hit}; \mathbf{x}) f_q(q_i; \mathbf{x}) f_t(t_i; \mathbf{x}) , \quad (9.1)$$

where the products are taken over (un)hit PMTs and where

- $P(i \text{ (un)hit}; \mathbf{x})$ is the probability that, given the parameters \mathbf{x} , the i -th tube is (un)hit,
- q_i is the observed charge of the i -th tube's hit,
- t_i is the observed time of the i -th tube's hit,

¹In this chapter, "energy" is kinetic energy.

- $f_q(q_i; \mathbf{x})$ is the p.d.f. for the i -th tube's charge given \mathbf{x} evaluated at the observed value q_i , and
- $f_t(t_i; \mathbf{x})$ is the p.d.f. for the i -th tube's time given \mathbf{x} evaluated at the observed value t_i .

It is convenient to work with the negative logarithm of \mathcal{L} , and since the charge-related and time-related portions are somewhat distinct, we define F , F_q , and F_t :

$$F(\mathbf{x}) \equiv -\log(\mathcal{L}(\mathbf{x})) \equiv F_q(\mathbf{x}) + F_t(\mathbf{x}) , \quad (9.2)$$

with

$$F_q(\mathbf{x}) \equiv -\sum_{\text{unhit}} \log(1 - P(i \text{ hit}; \mathbf{x})) - \sum_{\text{hit}} \log(P(i \text{ hit}; \mathbf{x}) f_q(q_i; \mathbf{x})) \quad (9.3)$$

$$F_t(\mathbf{x}) \equiv -\sum_{\text{hit}} \log(f_t(t_i; \mathbf{x})) , \quad (9.4)$$

where we've replaced $P(i \text{ unhit}; \mathbf{x})$ with $1 - P(i \text{ hit}; \mathbf{x})$. For brevity, we usually call F_q and F_t the charge and time likelihoods, although they are actually negative logarithms of likelihoods.

9.2 The charge likelihood

If the number of photoelectrons n produced in a PMT is known, then the p.d.f. $f_q(q)$ for the observed charge q is fully specified, regardless of \mathbf{x} . Further, n is a Poisson variable whose mean μ_n is a function of \mathbf{x} . Finally, the probability of a hit actually occurring depends only on this Poisson mean:² $P(i \text{ hit}) = 1 - e^{-\mu_n}$. These facts motivate a two-step approach to calculating the charge likelihood.

Step 1

Using the known optical properties of the detector and the known propagation properties of particles, one can determine, for a given set of track parameters \mathbf{x} , the mean number of photoelectrons that a particular PMT should see. We call this quantity the PMT's "predicted charge" μ . When calculating μ , one must consider

- the quantity of light produced by the track,
- the angular profile of the produced light,
- the absorption, scattering, and fluorescence of the light in oil,
- the acceptance of the PMT,

²This ignores discriminator threshold effects which are small.

and anything else that influences the mapping $(x_0, y_0, z_0, T_0, \theta_0, \phi_0, E_0) \mapsto \mu$.

Step 2

The expression for F_q given in Eq. (9.3) can be rewritten in terms of these predicted charges (or Poisson means) $\{\mu_i\}$:

$$F_q(\mathbf{x}) \equiv - \sum_{\text{unhit}} \log(1 - P(i \text{ hit}; \mu_i)) - \sum_{\text{hit}} \log(P(i \text{ hit}; \mu_i) f_q(q_i; \mu_i)) , \quad (9.5)$$

where we've recast the symbols P and f as functions dependent on μ_i , making the mapping of Step 1 implicit. The two interesting quantities on the right hand side of Eq. (9.5), namely $P(\text{hit}; \mu)$ and $f_q(q; \mu)$, depend only on the properties of the PMTs, their downstream electronics, and the Poisson distribution. They do not depend at all on the happenings inside the detector volume. This stepwise approach decouples the fitter's track and optical photon models from its knowledge of PMT and electronics behavior. The latter is determined using laser flask data (§9.2.3). The former is described next.

9.2.1 Calculating predicted charge μ

We begin with a simplified scenario, bringing in complexities as we go. Throughout, we deal with a fixed PMT and fixed track parameters, allowing us to drop from our notation references to these quantities. We define many functions in this section without immediately discussing how they are established in practice. These discussions are saved for Appendix C.³

A simple case

Assume that we have a point-like event that emits scintillation light only. Also assume that all optical photon extinction in the oil is due to absorption (*i.e.*, there is no scattering or fluorescence.) Also assume that all materials are non-reflective.

We take a PMT situated as depicted in Figure 9.1. The distance between the event and the PMT is labeled r , and the angle of incidence of the light is labeled η , with $\eta=0$ corresponding to normal incidence. We write the predicted charge for this PMT as

$$\mu = \Phi \Omega(r) T(r) \epsilon(\eta) , \quad (9.6)$$

where:

- Φ is an event-energy-dependent light yield.⁴
- $\Omega(r)$ is a distance-dependent solid angle factor (§C.1).

³In general, the functions are built from the Monte Carlo, which includes all of the physical phenomena we are considering in the reconstruction algorithm.

⁴Since they show up in so many places, we suppress explicit E_0 dependences throughout this document. For example, we write here Φ instead of $\Phi(E_0)$. See §C.5 for a list of functions with full arguments.

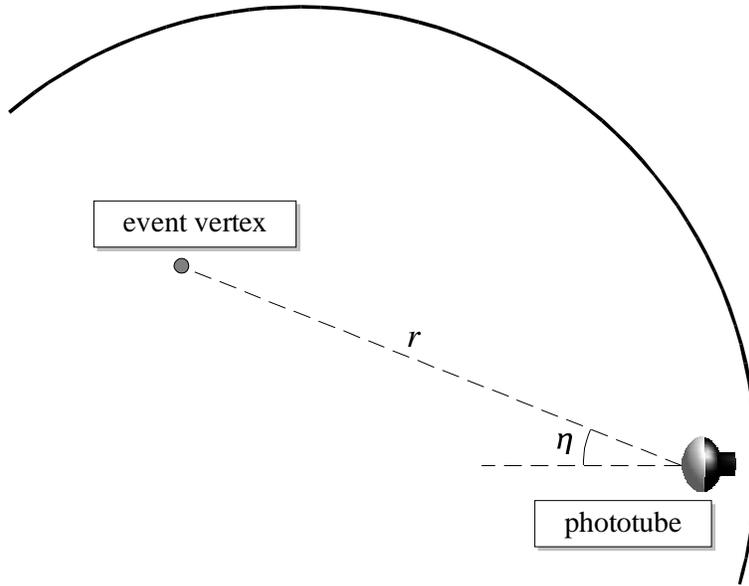


Figure 9.1: A simple geometry. A point source is located a distance r from a PMT, and the light impinges on the PMT at an angle η from the normal.

- $T(r)$ is the transmission of the oil and PMT glass as a function of the light propagation distance. This function includes the wavelength dependence of photon production, propagation, and detection (§C.2).
- $\epsilon(\eta)$ is the η -dependent acceptance of the PMT (§C.1).

Any overall constant factors can be taken as part of Φ .

An extended track

Now take the event to be an extended line source. The light production need not be uniform along the track, so we introduce $\rho(s)$, the probability density for light emission as a function of the distance s along the track.⁵ The track origin corresponds to $s=0$. (See Figure 9.2.) Examples of $\rho(s)$ are plotted in Figure 9.3. As $\rho(s)$ is a probability density, it satisfies $\int_{-\infty}^{\infty} \rho(s) ds = 1$.

The predicted charge for our PMT becomes an integral along the track,

$$\mu = \Phi \int_{-\infty}^{\infty} ds \rho(s) \Omega(r(s)) T(r(s)) \epsilon(\eta(s)) , \quad (9.7)$$

where the s -dependences of r and η have been made explicit. For ease of notation, we

⁵This is another instance of a suppressed E_0 dependence: $\rho(s; E_0)$. See §C.5 as needed.

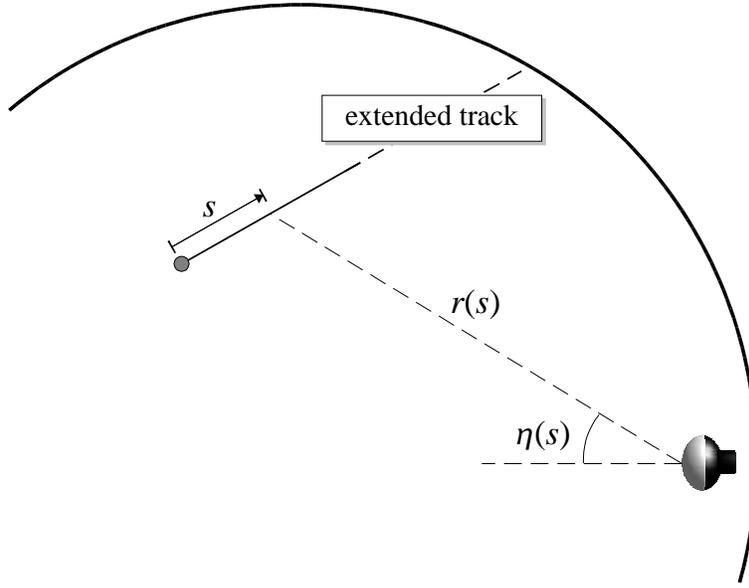


Figure 9.2: An extended source of light. The distance r and angle η are now functions of s , the distance along the track. Implicitly, the s dependences of r and η require knowledge of the geometric track parameters. $s=0$ occurs at the track's starting point.

redefine Ω , T , and ϵ to be functions of s directly rather than functions of r and η ,

$$\mu = \Phi \int_{-\infty}^{\infty} ds \rho(s) \Omega(s) T(s) \epsilon(s) , \quad (9.8)$$

although this requires a fixed \mathbf{x} . Since we are about to turn on Cherenkov light, we add “sci” subscripts to those expressions that are light-type specific:

$$\mu_{\text{sci}} = \Phi_{\text{sci}} \int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s) \Omega(s) T_{\text{sci}}(s) \epsilon(s) . \quad (9.9)$$

Cherenkov light

Figure 9.4 introduces θ , the angle to the PMT from the track. With this angle, we can write the Cherenkov predicted charge μ_{Ch} as

$$\mu_{\text{Ch}} = \Phi_{\text{Ch}} \int_{-\infty}^{\infty} ds \rho_{\text{Ch}}(s) \Omega(s) T_{\text{Ch}}(s) \epsilon(s) g(\cos \theta(s); s) . \quad (9.10)$$

This expression differs from its scintillation counterpart only by the presence of an angular emission profile $g(\cos \theta(s); s)$ (which, in the scintillation case, amounts to a constant factor which is absorbed into the definition of Φ_{sci} .) The “; s ” part of $g(\cos \theta(s); s)$ indicates that the emission profile $g(\cos \theta)$ changes as the track propagates and loses energy, while the “ $\theta(s)$ ” part reminds us that the angle θ depends geometrically on which part of the track

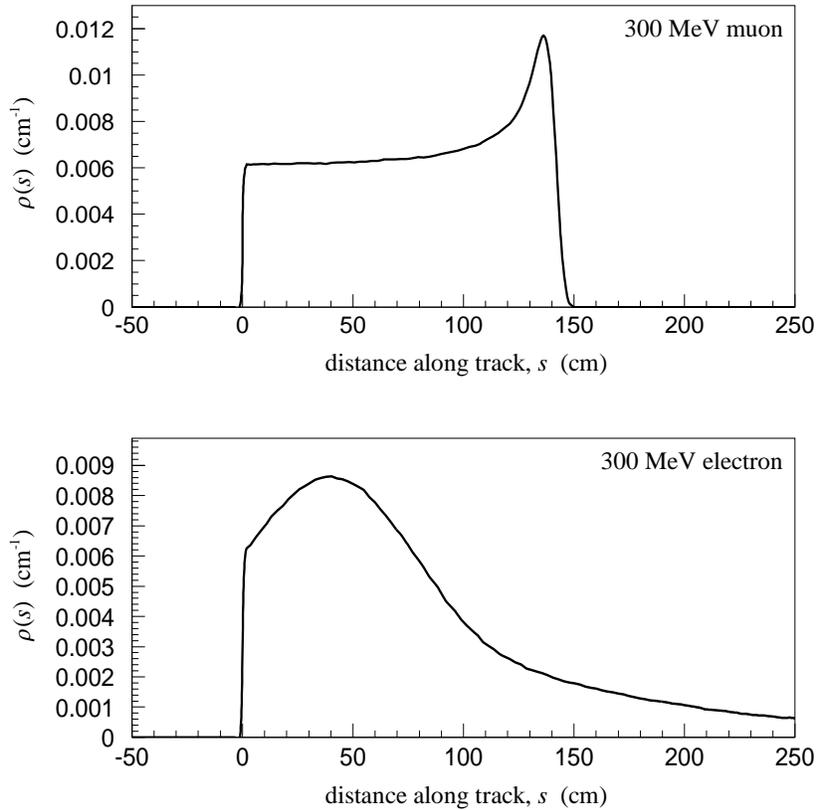


Figure 9.3: Scintillation production profiles $\rho(s)$ for 300 MeV muons (top) and electrons (bottom). See also §C.4.

we are considering. Figure 9.5 shows examples of $g(\cos \theta; s)$ and $\rho_{\text{Ch}}(s)$.

We skipped a few steps in writing the above expression. A more complete treatment would involve $\frac{d^2 P(\cos \theta, \phi)}{d(\cos \theta) d\phi}$, the differential probability of sending an emitted Cherenkov photon in the direction $(\cos \theta, \phi)$.⁶ This quantity would appear inside an integral over the solid angle subtended by the PMT. However, by assuming that $\frac{d^2 P(\cos \theta, \phi)}{d(\cos \theta) d\phi}$ is constant across the PMT face, we effect the integration simply by multiplying the differential probability by the solid angle factor Ω . Further, the azimuthal symmetry of the track allows us to reduce the two-dimensional expression $\frac{d^2 P(\cos \theta, \phi)}{d(\cos \theta) d\phi}$ to a one-dimensional one. We take this one-dimensional expression to be the angular p.d.f. for Cherenkov light emission, $g(\cos \theta)$. Doing so hides a factor of $\frac{1}{4\pi}$ in the definition of Φ_{Ch} since $g(\cos \theta)$ satisfies

$$\int_{-1}^1 g(\cos \theta; s) d(\cos \theta) = 1 \quad (9.11)$$

for all s .

⁶ ϕ here is the azimuth around the track.

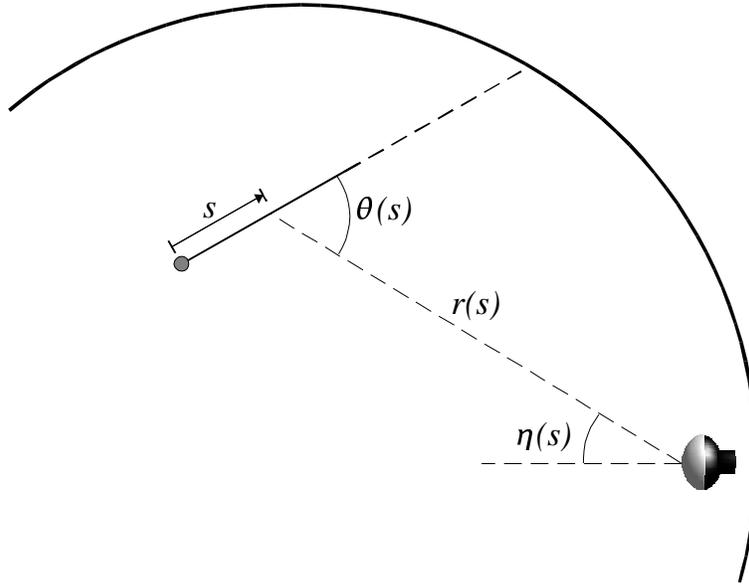


Figure 9.4: The angle θ defined.

Indirect light

If it were not for scattering, fluorescence, and reflections,⁷ we'd be finished. We would add μ_{sci} to μ_{Ch} to get the total predicted charge for the PMT. But our detector does have scattering, fluorescence, and reflections (together: *indirect light*), which we introduce now.

We begin with indirect light that originated as scintillation light. Take an infinitesimal element ds along the track. We view this element as a point source situated in the tank at radius R and at angle Θ relative to the PMT's position. Figure 9.6 shows this configuration.

The direct light from such a source has already been calculated. It is the right hand side of Eq. (9.9) without the integration, namely

$$d\mu_{\text{sci}}^{\text{direct}} = ds \Phi_{\text{sci}} \rho_{\text{sci}}(s) \Omega(s) T_{\text{sci}}(s) \epsilon(s) . \quad (9.12)$$

Writing an analytic expression for the *indirect* light from this point source would involve an elaborate integral over emission angles and scattering points throughout the tank. To avoid this, we observe that the value of such an integral must be proportional to the source strength and must otherwise depend only on the topological variables R and Θ . We can eliminate the source strength by forming a ratio of the indirect and direct light predictions:

$$A_{\text{sci}}(R, \cos \Theta) \equiv \frac{d\mu_{\text{sci}}^{\text{indirect}}}{d\mu_{\text{sci}}^{\text{direct}}} . \quad (9.13)$$

We call the ratio $A_{\text{sci}}(R, \cos \Theta)$ the scintillation “scattering table” (although it includes all

⁷and the computational impracticality of the expressions

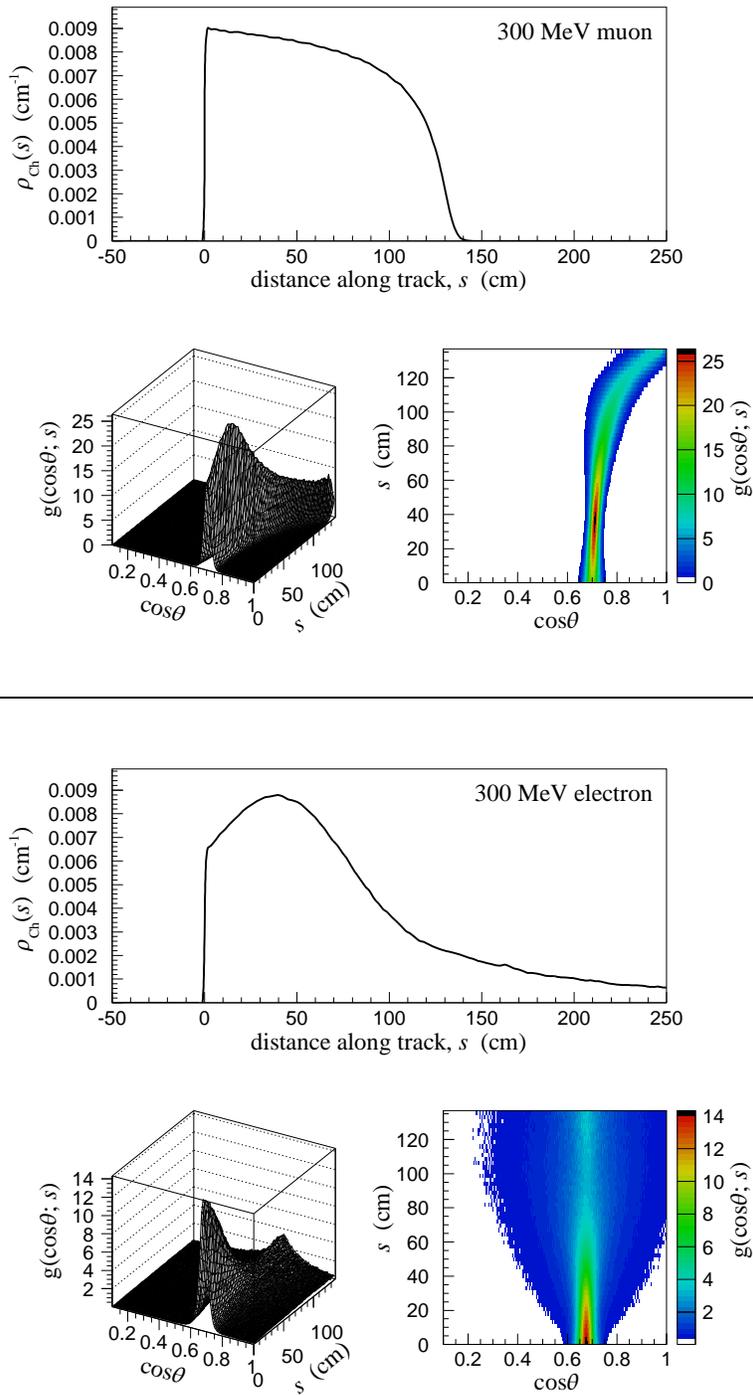


Figure 9.5: Cherenkov emission profiles for 300 MeV muons (top) and electrons (bottom). The wide panels show $\rho_{\text{Ch}}(s)$, the p.d.f. for Cherenkov light production along the track. The square panels provide views of $g(\cos\theta; s)$. The closing of the muon Cherenkov cone can be seen at large s . §C.4 has more.

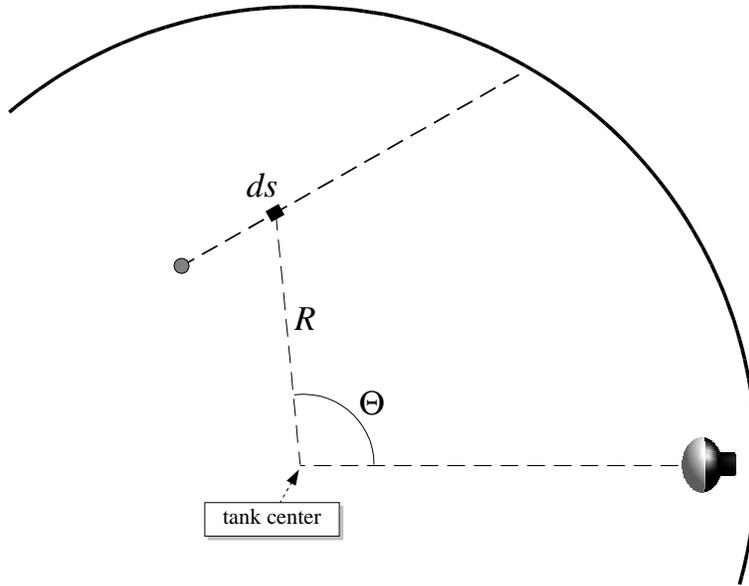


Figure 9.6: The geometry for indirect light. An infinitesimal line element ds of the track is treated as a point source situated as shown. R is the distance from the tank center to the source and Θ is the angle about the tank center from the point source to the PMT.

indirect light, not just scattered light.) This table is a property of the detector. It depends not on the particulars of the track but only on the $(R, \cos \Theta)$ geometry of the source element and PMT. §C.3 describes the creation of $A_{\text{sci}}(R, \cos \Theta)$.

With the scattering table defined, we can immediately incorporate indirect light into the scintillation prediction:

$$\mu_{\text{sci}} = \Phi_{\text{sci}} \int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s) \Omega(s) T_{\text{sci}}(s) \epsilon(s) [1 + A_{\text{sci}}(R(s), \cos \Theta(s))] , \quad (9.14)$$

where we have made explicit the fact that the arguments R and Θ change as s varies.

For Cherenkov light, the situation is more complex since the light emission is anisotropic. We need two additional variables that specify the direction, relative to the PMT position and the tank center, of a *vector* source. We make the following non-unique choices for these two variables:

- θ , the angle between the source direction and the source-to-PMT ray. This is the same θ used elsewhere.
- ϕ , the angle between (a) the plane containing the tank center, the PMT, and the source; and (b) the plane containing the track and the tank center.

Figure 9.7 shows these variables graphically.

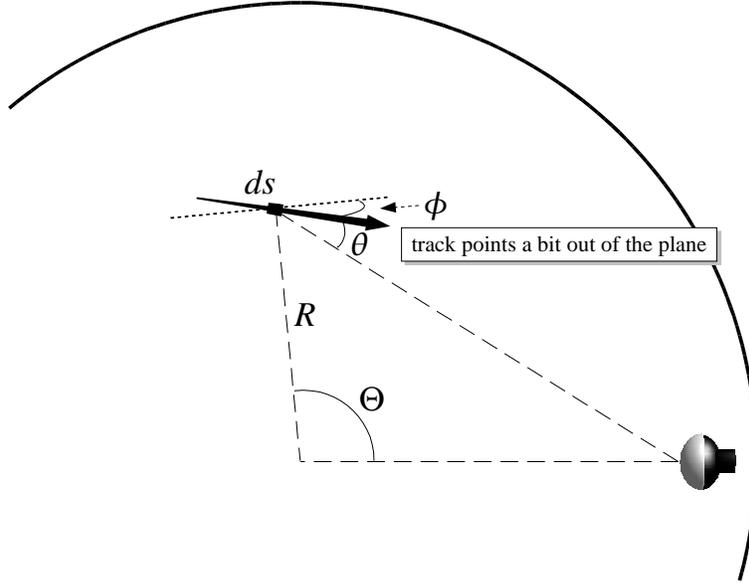


Figure 9.7: Geometry for the Cherenkov scattering table. All broken lines lie in the plane of the page. The track comes out of the plane. θ is the angle between the track and a vector pointing from the source to the PMT. ϕ is the track's angle of rotation about the line labeled R .

With the topology defined, we could write the Cherenkov analog of Eq. (9.13) as

$$A'_{\text{Ch}}(R, \cos \Theta, \cos \theta, \phi) \equiv \frac{d\mu_{\text{Ch}}^{\text{indirect}}}{d\mu_{\text{Ch}}^{\text{direct}}} . \quad (9.15)$$

However, it will be convenient to normalize the indirect Cherenkov light to a fictitious *isotropic* Cherenkov source with predicted charge $d\mu_{\text{Ch}}^{\text{direct,iso}}$, so we actually write

$$A_{\text{Ch}}(R, \cos \Theta, \cos \theta, \phi) \equiv \frac{d\mu_{\text{Ch}}^{\text{indirect}}}{d\mu_{\text{Ch}}^{\text{direct,iso}}} . \quad (9.16)$$

We now construct the total Cherenkov prediction

$$\mu_{\text{Ch}} = \mu_{\text{Ch}}^{\text{direct}} + \mu_{\text{Ch}}^{\text{indirect}} , \quad (9.17)$$

with $\mu_{\text{Ch}}^{\text{direct}}$ given by Eq. (9.10) and with

$$\mu_{\text{Ch}}^{\text{indirect}} = \Phi_{\text{Ch}} \int_{-\infty}^{\infty} ds \rho_{\text{Ch}}(s) \Omega(s) T_{\text{Ch}}(s) \epsilon(s) A_{\text{Ch}}(R(s), \cos \Theta(s), \cos \theta(s), \phi(s)) , \quad (9.18)$$

where we have used

$$d\mu_{\text{Ch}}^{\text{direct,iso}} = ds \Phi_{\text{Ch}} \rho_{\text{Ch}}(s) \Omega(s) T_{\text{Ch}}(s) \epsilon(s) \quad (9.19)$$

in analogy with Eq. (9.12). See §C.3 for more.

We are now finished with our charge predictions. Summing the Cherenkov and scintillation predictions yields the total predicted charge for the PMT,

$$\mu = \mu_{\text{Ch}} + \mu_{\text{sci}} . \quad (9.20)$$

9.2.2 Computational realities

The above expressions for predicted charge all involve integrals along the track. Numerically evaluating these integrals with sufficient spatial granularity for every PMT in every minimization step would result in unusably slow code. We avoid this by performing all integrations beforehand, as described below.

Parabolic approximation (scintillation version)

Eq. (9.9) gives the direct scintillation predicted charge. Its integrand,

$$\rho_{\text{sci}}(s) \Omega(s) T_{\text{sci}}(s) \epsilon(s) \quad (9.21)$$

can be thought of as having two factors, the production profile $\rho_{\text{sci}}(s)$ and the effective acceptance, $J(s)$:

$$J(s) \equiv \Omega(s) T_{\text{sci}}(s) \epsilon(s) . \quad (9.22)$$

If the components of $J(s)$ vary gradually enough, we can take it to have a parabolic form:

$$J(s) = j_0 + j_1 s + j_2 s^2 . \quad (9.23)$$

Doing so allows us to write the predicted charge as

$$\begin{aligned} \mu_{\text{sci}}^{\text{direct}} &= \Phi_{\text{sci}} \int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s) (j_0 + j_1 s + j_2 s^2) \\ &= \Phi_{\text{sci}} \left(j_0 \int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s) + j_1 \int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s) s + j_2 \int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s) s^2 \right) . \end{aligned} \quad (9.24)$$

The first integral in Eq. (9.24) is identically unity, and the last two integrals we label I_1^{sci} and I_2^{sci} , giving

$$\mu_{\text{sci}}^{\text{direct}} = \Phi_{\text{sci}} (j_0 + j_1 I_1^{\text{sci}} + j_2 I_2^{\text{sci}}) . \quad (9.25)$$

Since the shape of $\rho_{\text{sci}}(s)$ depends on the energy E_0 of the particle, these integrals depend on E_0 . However, they depend on no other track parameter. It is thus feasible to tabulate these integrals beforehand, eliminating the integrations from the minimization code. We

obtain the coefficients $\{j_i\}$ by evaluating Eq. (9.22) at three points along the track,

$$\begin{aligned} J_0 &\equiv J(0) \\ J_1 &\equiv J(\Delta s_{\text{mid}}) \\ J_2 &\equiv J(2\Delta s_{\text{mid}}) , \end{aligned}$$

where Δs_{mid} is the mean of the Cherenkov emission profile.⁸ The parabola that goes through these three points has coefficients

$$\begin{aligned} j_0 &\equiv J_0 \\ j_1 &\equiv \frac{-3J_0 + 4J_1 - J_2}{2\Delta s_{\text{mid}}} \\ j_2 &\equiv \frac{J_0 - 2J_1 + J_2}{2\Delta s_{\text{mid}}^2} , \end{aligned}$$

completing the expression for $\mu_{\text{sci}}^{\text{direct}}$ in Eq. (9.25). Figure 9.8 shows an example.

The extension to indirect scintillation light is straightforward. The only difference is the definition of $J(s)$:

$$J(s) \equiv \Omega(s) T_{\text{sci}}(s) \epsilon(s) A_{\text{sci}}(R(s), \cos \Theta(s)) \quad (\text{for indirect scintillation light}) . \quad (9.26)$$

Parabolic approximation (Cherenkov version)

The expression for Cherenkov predicted charge involves the angular profile $g(\cos \theta(s); s)$. Including $g(\cos \theta(s); s)$ in the definition of $J(s)$ would ruin the parabolic treatment, as the function can vary violently near the Cherenkov angle. We therefore keep $J(s)$ as-is,

$$J(s) \equiv \Omega(s) T_{\text{Ch}}(s) \epsilon(s) , \quad (9.27)$$

and we put the angular profile inside the integrals:

$$\begin{aligned} \mu_{\text{Ch}}^{\text{direct}} = \Phi_{\text{Ch}} \left(\right. & j_0 \int_{-\infty}^{\infty} ds \rho_{\text{Ch}}(s) g(\cos \theta(s); s) \\ & + j_1 \int_{-\infty}^{\infty} ds \rho_{\text{Ch}}(s) g(\cos \theta(s); s) s \\ & \left. + j_2 \int_{-\infty}^{\infty} ds \rho_{\text{Ch}}(s) g(\cos \theta(s); s) s^2 \right) . \quad (9.28) \end{aligned}$$

As in the scintillation case, these integrals depend on the energy of the track. They also, however, depend on two parameters that define the PMT-track geometry.⁹ We choose as these two parameters the vertex-PMT distance, $r(0)$, and the cosine of the angle-to-tube as

⁸The choice of where to evaluate $J(s)$ is somewhat arbitrary. Any three points that sample a good bit of the light-producing range of s would do fine.

⁹This latter dependence comes from the function $\theta(s)$. That is, the mapping from s to θ changes with track orientation and tube distance.

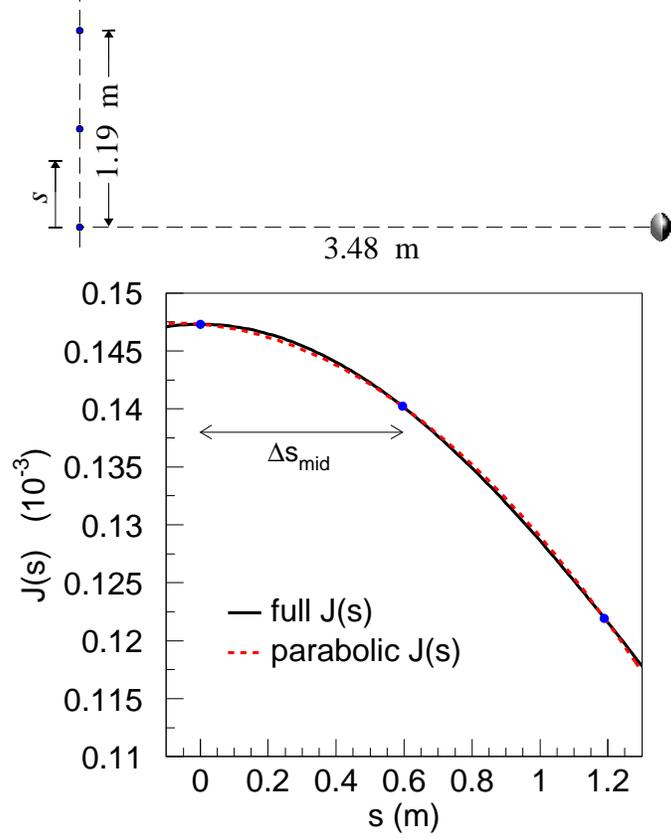


Figure 9.8: An example of $J(s)$. (Top) A possible track-PMT geometry is shown. The value of $J(s)$ is evaluated at three points separated by 0.595 m, the value of Δs_{mid} for a 300 MeV muon. (Bottom) The full $J(s)$ expression (Eq. (9.22)) is shown by the solid curve. The parabolic version, obtained using the three points indicated, is shown by the dashed curve. Using $\rho_{\text{sci}}(s)$ for a 300 MeV muon, the integral $\int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s) J(s)$ with the parabolic $J(s)$ differs from the one with the full $J(s)$ by 0.3%.

viewed from the vertex, $\cos \theta(0)$. Labeling the integrals $\mathcal{I}_i^{\text{Ch}}$ ($i \in \{1, 2, 3\}$) gives

$$\mu_{\text{Ch}}^{\text{direct}} = \Phi_{\text{Ch}} \left(j_0 \mathcal{I}_0^{\text{Ch}} + j_1 \mathcal{I}_1^{\text{Ch}} + j_2 \mathcal{I}_2^{\text{Ch}} \right). \quad (9.29)$$

Although the integrals $\{\mathcal{I}_i^{\text{Ch}}\}$ depend on three parameters, it is still feasible to tabulate their values ahead of time, again keeping integrations out of the minimization code.

For indirect Cherenkov light, recall that the scattering table A_{Ch} is normalized to a fictitious isotropic Cherenkov source. Thus, the parabolic method used for indirect scintillation light works for indirect Cherenkov light. $J(s)$ becomes

$$J(s) \equiv \Omega(s) T_{\text{Ch}}(s) \epsilon(s) A_{\text{Ch}}(R(s), \cos \Theta(s), \cos \theta(s), \phi(s)) \quad (\text{indirect Cher. light}), \quad (9.30)$$

and we tabulate the two energy-dependent integrals $\int ds \rho_{\text{Ch}}(s) s^i$: I_1^{Ch} and I_2^{Ch} .

9.2.3 The functions $P(\text{hit}; \mu)$ and $f_q(q; \mu)$

With the predicted charges in hand, Eq. (9.5) gives us the charge likelihood, provided we know the forms of the functions $P(\text{hit}; \mu)$ and $f_q(q; \mu)$.

$P(\text{hit}; \mu)$:

The actual number of photoelectrons n in a hit is a Poisson variable with mean μ . Assuming that any single photoelectron will fire the discriminator, the probability of a hit $P(\text{hit}; \mu)$ is simply

$$\begin{aligned} P(\text{hit}; \mu) &= P(n > 0; \mu) \\ &= 1 - e^{-\mu} . \end{aligned} \tag{9.31}$$

$f_q(q; \mu)$:

For a hit with a given predicted charge μ , we need the p.d.f. for the hit's measured charge q . This p.d.f. involves

- the underlying Poisson process that determines n ,
- the properties of the PMT,
- the amplification and digitization of the anode signal, and
- the conversion of the digital values into a properly calibrated charge measurement.

In our case, the laser flasks offer a convenient way to vary μ while observing q . Thus, one can simply map out the p.d.f. $f_q(q; \mu)$ empirically [132].¹⁰ Figure 9.9 shows examples of $f_q(q; \mu)$.

9.3 The time likelihood

The time portion of the likelihood requires that we know the p.d.f. $f_t(t; \mathbf{x})$ for a PMT's measured time t given the track parameters \mathbf{x} . Some of the dependence on the track parameters can be eliminated by working with a “corrected time” t^c , which we write as

$$t^c = t - T_0 - \frac{r(\Delta s_{\text{mid}}(E_0))}{c_n} - \frac{\Delta s_{\text{mid}}(E_0)}{c} \tag{9.32}$$

where

- t is the hit's measured time,

¹⁰There are several ways to learn μ for center-flask laser events, the simplest of which is to use Poisson statistics and the observed PMT occupancy. Beyond 15 PE or so, when there are rarely unhit tubes with which to measure PMT occupancy, one must do something fancier. The caption of Figure 8.9 mentions one scheme. Ref. [132], which discusses the creation of $f_q(q; \mu)$, uses another.

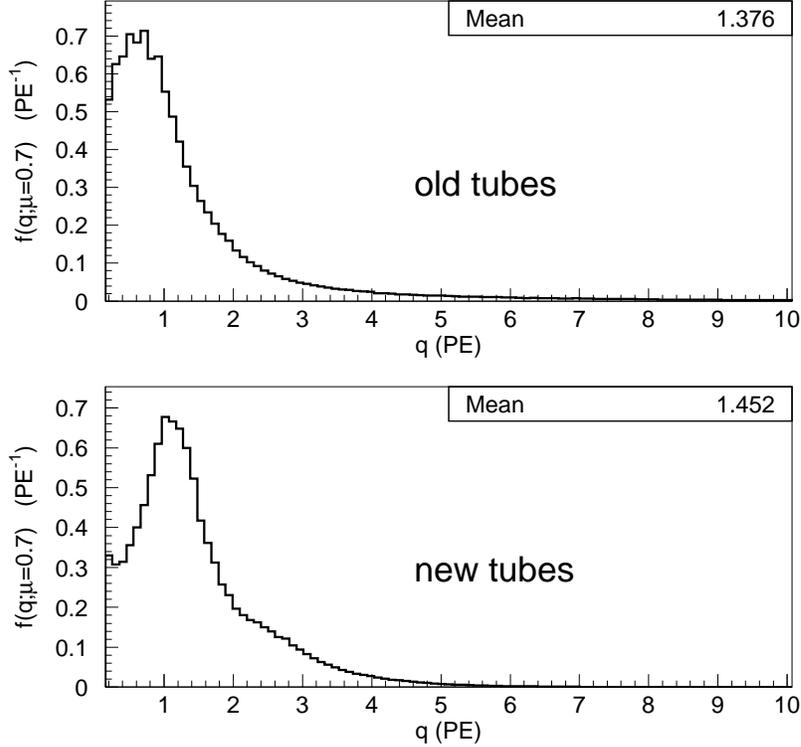


Figure 9.9: Examples of $f(q; \mu)$. These plots show the probability density functions for observing charge q on a *hit* PMT with predicted charge $\mu=0.7$. The top panel shows the p.d.f. for old tubes; the bottom, new tubes. The means are near $\frac{\mu}{1-e^{-\mu}} \approx 1.4$, as expected.

- T_0 is the track’s starting time,
- $\Delta s_{\text{mid}}(E_0)$ is the mean of the Cherenkov emission profile for a track of energy E_0 (the track “midpoint”),
- $r(\Delta s_{\text{mid}}(E_0))$ is the distance to the PMT from the track midpoint,
- c is the speed of light, and
- $c_n = 19.50$ cm/ns is the speed of light in mineral oil used for t^c . (The exact value isn’t critical, as the p.d.f. we create below can be for any quantity.)

The corrected time removes from our observable t the global shift T_0 , the photon time-of-flight from the track’s midpoint, and the particle’s time-of-flight from its origin to its midpoint. We could have chosen to take the corrected time from the track origin rather than the track midpoint, eliminating the last term in Eq. (9.32) and the involvement of Δs_{mid} . However, the above form of t^c improves the validity of the simplification we are about to make.

9.3.1 Reducing the problem

With t^c defined, we seek the p.d.f. $f_{t^c}(t^c)$ for the corrected time t^c given arbitrary PMT-track configurations. The space of these configurations is five dimensional, and producing tables of $f_{t^c}(t^c)$ throughout configuration space is a task orders of magnitude beyond our computing resources. To reduce the task, we make the following assumption: the corrected time p.d.f. depends only on the track energy, the predicted “prompt” charge, and the predicted “late” charge. (We define “prompt” and “late” charge later.) Loosely, this is the assumption that the shape of the corrected time spectrum is dominated by the physical extent of the track (characterized by its energy) and by the raw amounts of prompt and late light reaching the PMT. The extent of the track affects the spread of possible hit times, since there is a spread of photon production times, while the amounts of prompt and late light affect the time distribution’s “peak-to-tail” ratio as well as pile-up and slewing phenomena.

Although this approximation reduces configuration space to three dimensions, we are still limited by computing, so we make one final simplification. For a fixed energy E_0 , we do not make a p.d.f. indexed by the pair $(\mu_{\text{prompt}}, \mu_{\text{late}})$. Rather, we make separate “primitive” distributions, indexed by *either* μ_{prompt} or μ_{late} , and from these primitive distributions we construct the full p.d.f. on-the-fly. For both historical and computational reasons, we actually create Cherenkov and scintillation primitive distributions which we subsequently label as the prompt and late primitive distributions.

9.3.2 Creating the primitive distributions

We first describe the Cherenkov primitive distributions. Everything below is done separately for old and new tubes and for muons and electrons.

We begin with simulated contained events throughout the detector with isotropically chosen directions and with fixed energy E_0 . These events are created with direct Cherenkov light only; all other light is turned off. For each hit, we use the true track parameters to evaluate the predicted direct Cherenkov charge $\mu_{\text{Ch}}^{\text{direct}}$. We then make corrected time histograms in various ranges of this predicted charge. Figure 9.10 shows three such corrected time histograms for 300 MeV muons. Since only Cherenkov light is present, the histograms show no late-time features. Also, the shape of the time spectrum depends on the $\mu_{\text{Ch}}^{\text{direct}}$ range examined, with the spectrum getting narrower and earlier as the predicted charge increases. We parametrize this charge-dependent shape variation by fitting the time spectrum in each charge bin to a Gaussian (see the fits in Figure 9.10) and by subsequently fitting the resulting Gaussian parameters (mean and RMS) to sextics across $\mu_{\text{Ch}}^{\text{direct}}$ values. Figure 9.11 shows an example of these “second-level” fits. We repeat the preceding exercise at several energies, with the two second-level fits providing seven parameters for the Gaussian mean and seven parameters for the Gaussian width at each energy. The energy dependences of these fourteen parameters are then fit in a final third-level parametrization. Figure 9.12 shows the first seven of these. (Note that all this parametrization is required so that the likelihood surface presented to the minimization algorithm is smooth.)

The Cherenkov primitive distributions are now complete. When we need the time

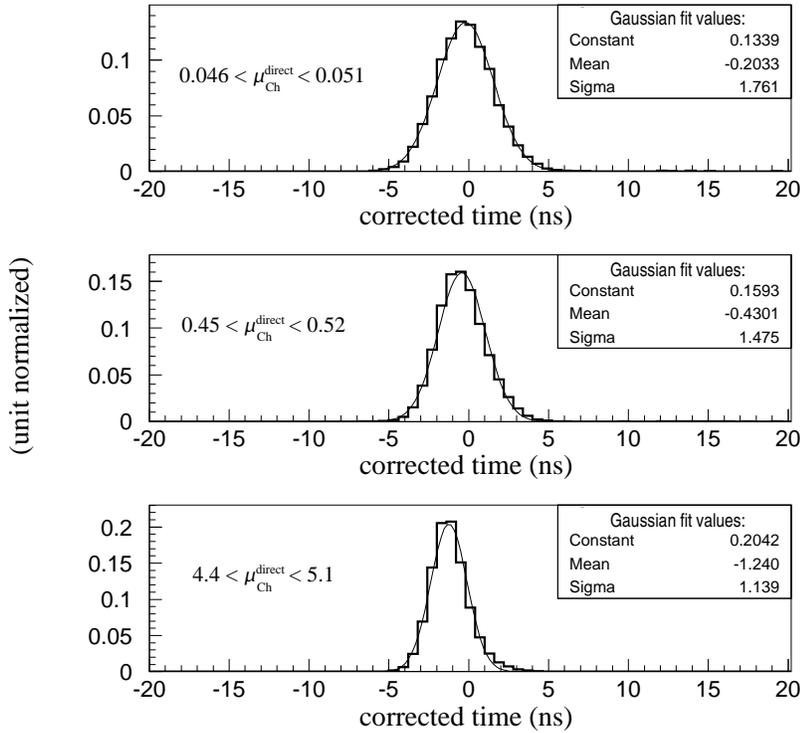


Figure 9.10: Cherenkov-only corrected time histograms for hits with $\mu_{\text{Ch}}^{\text{direct}}$ near 0.05 PE, 0.5 PE, and 5 PE for 300 MeV muons. Gaussian fits like those shown here provide the basis for the $\mu_{\text{Ch}}^{\text{direct}}$ and E_0 parametrizations (see main text.)

distribution for a PMT of predicted charge $\mu_{\text{Ch}}^{\text{direct}}$ in an event with energy E_0 , we do the following. First, we evaluate the fourteen third-level quartic curves at E_0 to obtain the parameters for the two second-level sextic functions. Next, we evaluate the sextic functions at $\log(\mu_{\text{Ch}}^{\text{direct}})$ to obtain the mean and RMS of the prompt time spectrum. This time spectrum is the Cherenkov primitive distribution, and we label it $G_{\text{Ch}}(t^c; E_0, \mu_{\text{prompt}})$.

To check how well all this works, we can make corrected time histograms for hits of various predicted charges from events of various energies, and we can compare these histograms to the shapes we get from our parametrization. Figure 9.13 shows several such comparisons.

For the scintillation primitive distributions, events are generated with scintillation light only, and corrected time histograms are created in bins of $\mu_{\text{sci}}^{\text{direct}}$. Each corrected time histogram is fit to a sum of two exponentials with Gaussian smearing. The exponential decay constants are fixed at $\tau_1=5$ ns and $\tau_2=30$ ns, leaving three free parameters: the time origin, the Gaussian resolution, and the relative weight of the two exponentials. These three parameters are fit across $\log(\mu_{\text{sci}}^{\text{direct}})$ to sextics, and the resulting seven parameters for each sextic are fit across energy to quartics, in analogy with the Cherenkov case. We show the result of the scintillation parametrization for muons in Figure 9.14.

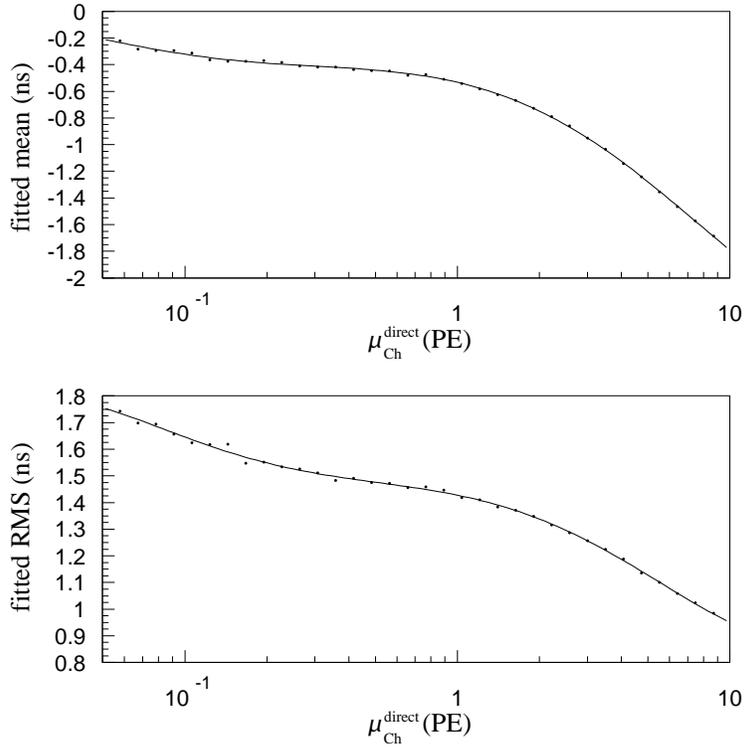


Figure 9.11: Gaussian parameters versus $\mu_{\text{Ch}}^{\text{direct}}$. The mean and RMS returned by the Gaussian fits to the corrected time spectra in each $\mu_{\text{Ch}}^{\text{direct}}$ bin are fitted here to a sextic in $\log(\mu_{\text{Ch}}^{\text{direct}})$. Fits for 300 MeV muons are shown. To connect this figure with Figure 9.10: The middle plot in Figure 9.10 (corresponding to $\mu_{\text{Ch}}^{\text{direct}} \approx 0.5$) has a fitted mean of -0.43 and a fitted RMS (“Sigma” in the plot legend) of 1.48 . These values can be found above at abscissa $\mu_{\text{Ch}}^{\text{direct}} \approx 0.5$.

9.3.3 Using the primitive distributions

To turn the primitive distributions into the sought after p.d.f. $f_{t^c}(t^c)$ for a given hit, we make the following generalizations/approximations:

- Direct Cherenkov light is prompt; everything else is late.
- Hits that contain one or more prompt Cherenkov photoelectrons have corrected times that follow the Cherenkov primitive distribution.
- Hits that contain no prompt Cherenkov photoelectrons have corrected times that follow the scintillation primitive distribution.

The total distribution will be some mixture based on the prompt and late predicted charges. Note that the last item above says that we are using the scintillation-based time spectra for *all* late light. This is reasonable since (a) UV fluorescence, the dominant late light source,

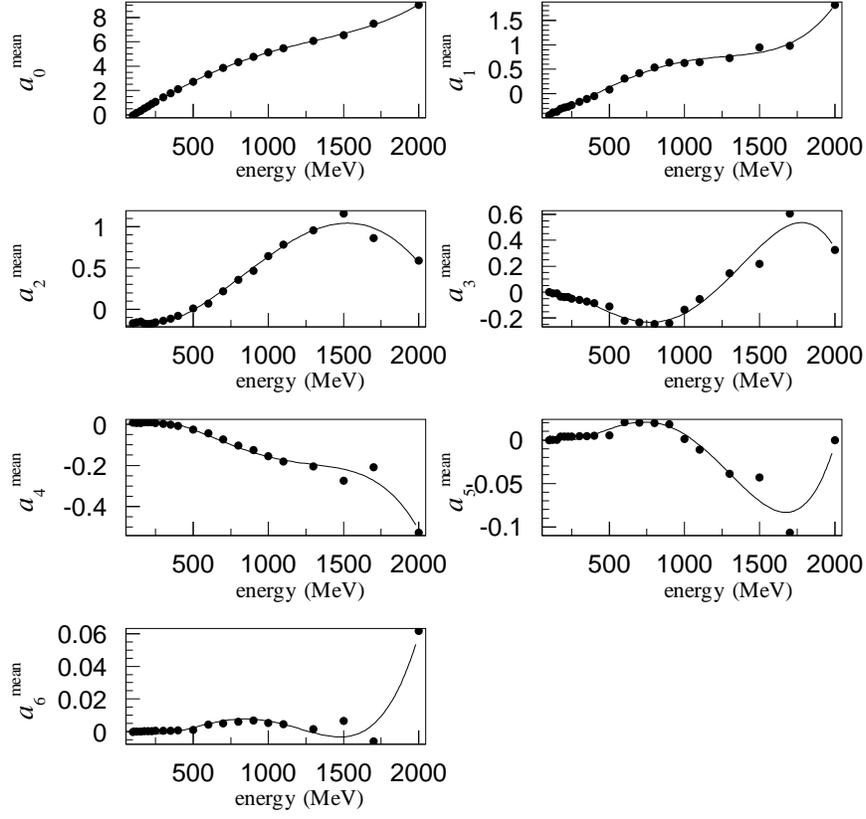


Figure 9.12: Third-level fits. The fits shown in Figure 9.11 for 300 MeV muons are repeated at many energies. The seven fit parameters are then fit versus energy to quartics, as shown here for the Gaussian mean parameters.

shares scintillation's time structure, and (b) prompt light timing is more critical than late light timing.

With the above assertions set, we proceed by defining the prompt and late charge predictions:

$$\mu_{\text{prompt}} \equiv 0.95 \mu_{\text{Ch}}^{\text{direct}} \quad (9.33)$$

$$\mu_{\text{late}} \equiv 0.05 \mu_{\text{Ch}}^{\text{direct}} + \mu_{\text{Ch}}^{\text{indirect}} + \mu_{\text{sci}}^{\text{direct}} + \mu_{\text{sci}}^{\text{indirect}} , \quad (9.34)$$

where 5% of the direct Cherenkov light has been called late to account for PMT late pulsing. The Poisson distribution gives us the probability that the hit has no prompt or late photoelectrons:

$$P(\text{no prompt PEs}) \equiv e^{-\mu_{\text{prompt}}}$$

$$P(\text{no late PEs}) \equiv e^{-\mu_{\text{late}}} .$$

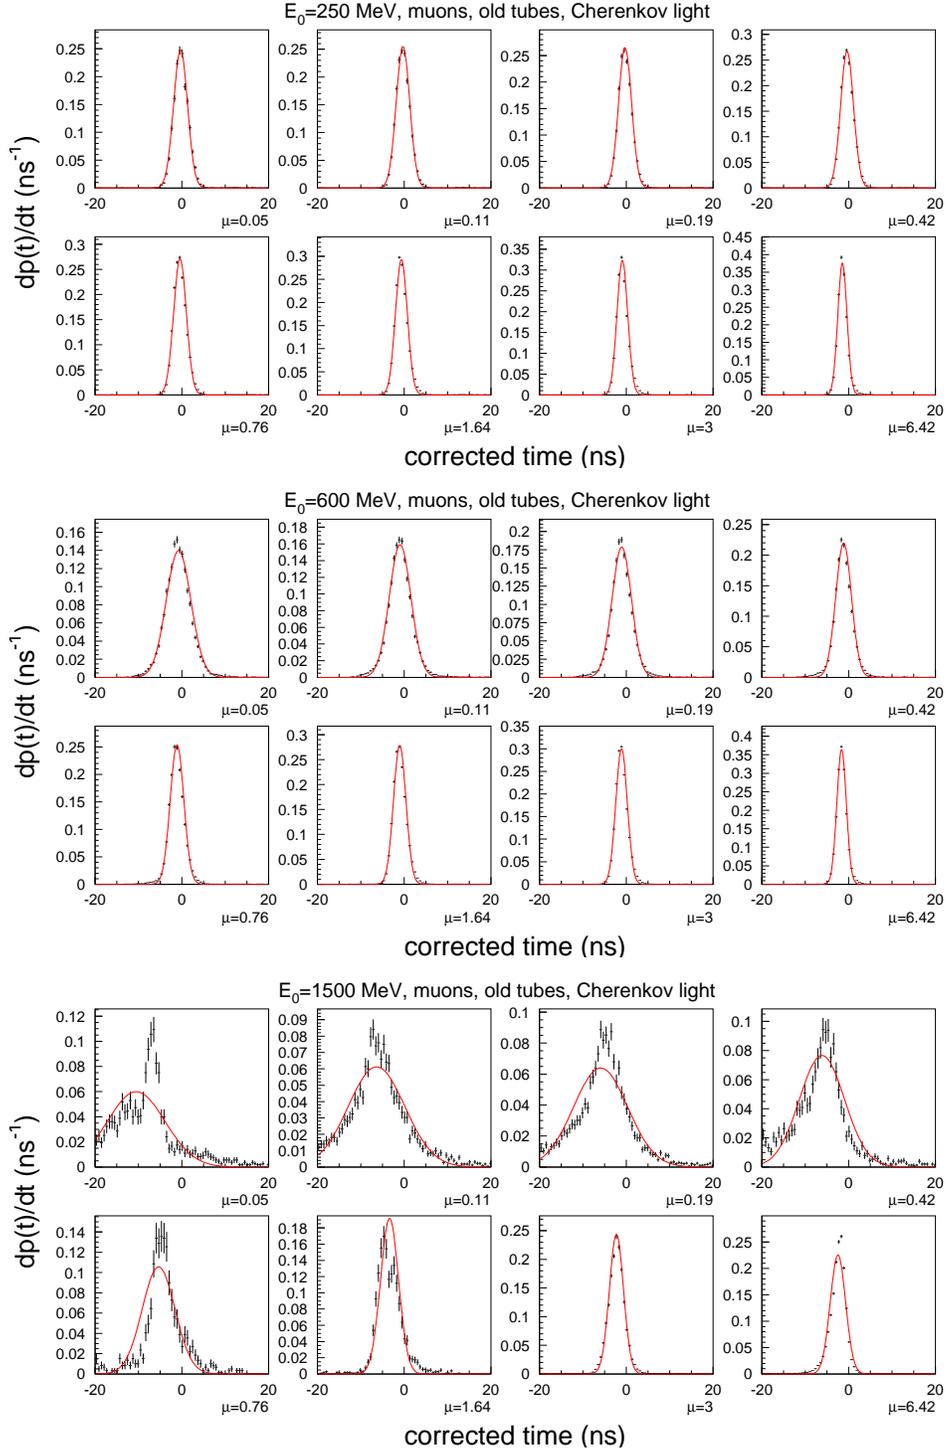


Figure 9.13: Verifying the Cherenkov primitive distributions. For simulated muons of three energies (250 MeV (top), 600 MeV (middle), 1500 MeV (bottom)), the distributions of corrected times for hits in eight $\mu_{\text{Ch}}^{\text{direct}}$ bins are shown as points with statistical error bars. The curve overlaid on each plot is the parametrization discussed in the text. The distributions are well described by the parametrization procedure except for low charge hits in high energy events (rare) where the space-reduction approximations of §9.3.1 begin to break down.

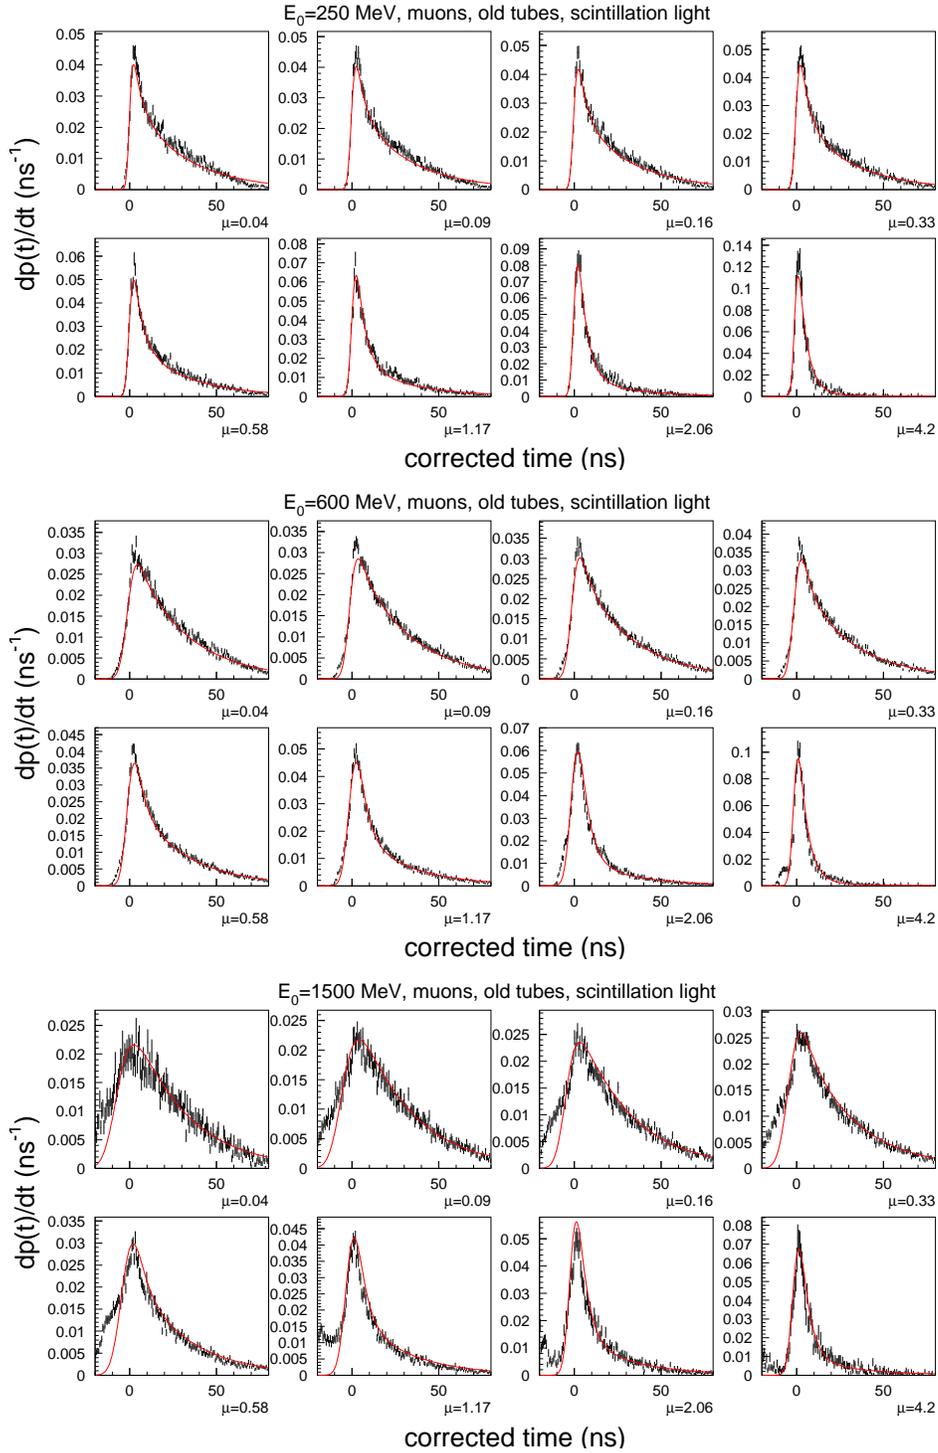


Figure 9.14: Verifying the scintillation primitive distributions. For simulated muons of three energies (250 MeV (top), 600 MeV (middle), 1500 MeV (bottom)), the distributions of corrected times for hits in eight $\mu_{\text{sci}}^{\text{direct}}$ bins are shown as points with statistical error bars. The curve overlaid on each plot is the parametrization discussed in the text.

With these probabilities, we construct the probability that a given hit contains at least one prompt photoelectron:

$$P(\text{prompt PE present} \mid \text{hit}) = \frac{1 - P(\text{no prompt PEs})}{1 - P(\text{no prompt PEs})P(\text{no late PEs})} . \quad (9.35)$$

Using this quantity as the weight w_p for the Cherenkov (prompt) primitive distribution and using $w_l \equiv 1 - w_p$ as the weight for the scintillation (late) primitive distribution, we write, finally, our full p.d.f.,

$$f_{t^c}(t^c; E_0, \mu_{\text{prompt}}, \mu_{\text{late}}) = w_p G_{\text{Ch}}(t^c; E_0, \mu_{\text{prompt}}) + w_l G_{\text{sci}}(t^c; E_0, \mu_{\text{late}}) , \quad (9.36)$$

where we've used G_{Ch} and G_{sci} to represent the Cherenkov and scintillation primitive distributions, respectively.

It is important to note that Eq. (9.35) gives the probability that a prompt photoelectron exists *given* a hit exists. This condition removes the overall probability of a hit occurring. Thus, even if the absolute probability of a prompt photoelectron is low, w_p can be large if $\mu_{\text{late}} \ll \mu_{\text{prompt}}$.

9.4 μ and e hypotheses

The single-track reconstruction algorithm can assume either a muon or an electron track. We list here the components of the algorithm that depend on the track hypothesis *and* that are direct inputs to the fitter. We do not list any quantities that are encapsulated by these, even if they are hypothesis dependent. The creation of a new track hypothesis (*e.g.*, proton) would require the items listed.

Quantities that depend on the track hypothesis:

- $\Phi_{\text{sci}}(E_0), \Phi_{\text{Ch}}(E_0)$
- $I_1^{\text{sci}}(E_0), I_2^{\text{sci}}(E_0), I_1^{\text{Ch}}(E_0), I_2^{\text{Ch}}(E_0)$
- $\mathcal{I}_0^{\text{Ch}}(E_0, r(0), \cos \theta(0)), \mathcal{I}_1^{\text{Ch}}(E_0, r(0), \cos \theta(0)), \mathcal{I}_2^{\text{Ch}}(E_0, r(0), \cos \theta(0))$
- $G_{\text{Ch}}(t^c; E_0, \mu_{\text{prompt}}), G_{\text{sci}}(t^c; E_0, \mu_{\text{late}})$

9.5 Two-track fits

A single track with seven parameters is sufficient for reconstructing cosmic rays, Michel electrons, and charged current quasi-elastic events (ignoring the minor perturbation from the recoil nucleon). Neutral current π^0 events, however, require a two-track likelihood. Figure 9.15 shows the twelve parameters needed to describe two γ tracks originating from

a common vertex.¹¹ The electron tables serve double duty as the γ tables since the Mini-BooNE detector cannot distinguish an e from a γ . However, a γ track picks up a conversion distance s .

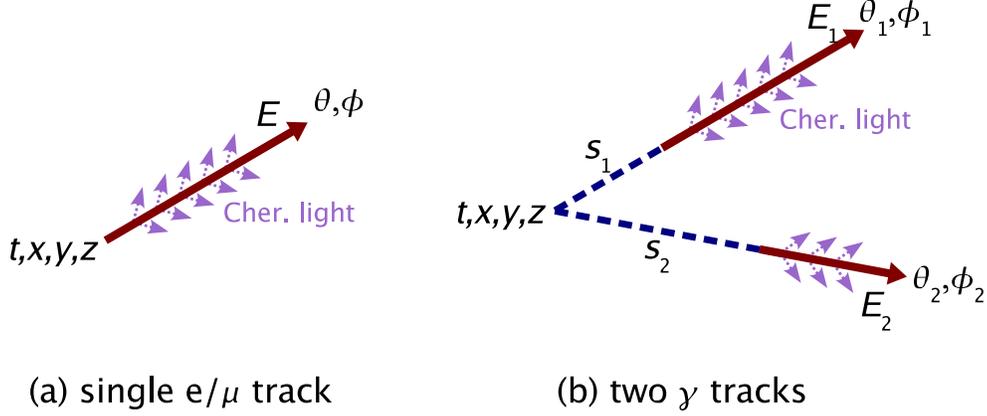


Figure 9.15: Internal fitter parameters for (a) a single muon or electron track and (b) two γ tracks. Each γ track includes a conversion distance s . One can constrain the parameters such that the invariant mass of the $\gamma\gamma$ system is always M_{π^0} .

Two-track charge likelihood

Recall that the charge likelihood F_q depends only on the measured charge q and the total predicted charge μ for each PMT. To form the two track version, one simply needs to add together the predicted charges from the two tracks to form the total predicted charge. That is,

$$\mu = \mu_{\text{track 1}} + \mu_{\text{track 2}} \cdot \quad (9.37)$$

Two-track time likelihood

The time reported by a PMT is the time of the earliest threshold-crossing pulse. This fact motivates the weighting of the primitive distributions described in Eq. (9.36), where the prompt primitive distribution is weighted according to the probability that a prompt photoelectron is present given any photoelectrons, and where the balance of the weight is given to the late primitive distribution. This scheme is extended to handle the case of two tracks as follows.

The two single-track primitive distributions are formed for the i -th track: G_{Ch}^i and G_{sci}^i . In anticipation of the clumping together of all late light below, the two late (“sci”) primitive

¹¹The code is equipped to handle one, two, or three tracks with arbitrary relationships, including distinct 4-vertices. However, this oscillation analysis needs only the cases described in the text.

distributions are averaged to form

$$\hat{G}_{\text{sci}} = \frac{1}{2}(G_{\text{sci}}^1 + G_{\text{sci}}^2) . \quad (9.38)$$

Of the two tracks present, one has a midpoint that is nearer to the target PMT than the other. We label that track's quantities with "n" (near) and the other's with "f" (far). We then define three Poisson probabilities in analogy with Eq. (9.33):

$$\mu_{\text{prompt,n}} \equiv 0.95 \mu_{\text{Ch,n}}^{\text{direct}} \quad (9.39)$$

$$\mu_{\text{prompt,f}} \equiv 0.95 \mu_{\text{Ch,f}}^{\text{direct}} \quad (9.40)$$

$$\mu_{\text{late}} \equiv \mu_{\text{tot}} - \mu_{\text{prompt,n}} - \mu_{\text{prompt,f}} . \quad (9.41)$$

Continuing as before,

$$\begin{aligned} \bar{P}_{\text{n}} &\equiv P(\text{no prompt PEs from near track}) &&\equiv e^{-\mu_{\text{prompt,n}}} \\ \bar{P}_{\text{f}} &\equiv P(\text{no prompt PEs from far track}) &&\equiv e^{-\mu_{\text{prompt,f}}} \\ \bar{P}_{\text{l}} &\equiv P(\text{no late PEs}) &&\equiv e^{-\mu_{\text{late}}} . \end{aligned}$$

This leads to the weights

$$\begin{aligned} w_{\text{n}} &= \frac{1 - \bar{P}_{\text{n}}}{1 - \bar{P}_{\text{n}}\bar{P}_{\text{f}}\bar{P}_{\text{l}}} \\ w_{\text{f}} &= \frac{1 - \bar{P}_{\text{l}}}{1 - \bar{P}_{\text{f}}\bar{P}_{\text{l}}}(1 - w_{\text{n}}) \\ w_{\text{l}} &= 1 - w_{\text{n}} - w_{\text{f}} , \end{aligned}$$

where w_{n} is the probability that a prompt photoelectron from the near track exists *given* that any photoelectron exists, and so forth. We use these weights to combine the two prompt primitive distributions and the averaged late primitive distribution:

$$f_{t^c}(t^c) = w_{\text{n}} G_{\text{Ch,n}}(t^c) + w_{\text{f}} G_{\text{Ch,f}}(t^c) + w_{\text{l}} \hat{G}_{\text{sci}}(t^c) , \quad (9.42)$$

thus forming the complete two-track corrected time p.d.f.

9.6 Minimizing $F = -\log(\mathcal{L})$

With the likelihood \mathcal{L} defined, we must find the parameter set \mathbf{x} that minimizes its negative logarithm F .

9.6.1 e and μ hypotheses (one track)

There are two complexities in the single track minimization.

1. The energy parameter E_0 is tied to the geometry of event via the track profiles $\rho(s; E_0)$

and $g(\cos \theta(s); s, E_0)$. If the spatial parameters (x_0 , etc.) are varied together with E_0 , the minimization algorithm occasionally gets confused by correlations between these. The solution is iteration.

2. Parametrization of input tables results in a smooth likelihood surface $\mathcal{L}(\mathbf{x})$. However, the discrete PMT lattice imprints a small fluctuating signal on $\mathcal{L}(\mathbf{x})$. As a result, minimization algorithms that rely on gradients or that provide little user control over step sizes are poorly suited to the problem. Thus, the MIGRAD method of Minuit [133] has difficulty, but the SIMPLEX method works well. We use the latter.

Seed parameters are derived from a fast fitter [132]. The energy E_0 is held at its seeded value and the six remaining parameters ($x_0, y_0, z_0, T_0, \theta_0, \phi_0$) are varied via Minuit/SIMPLEX to find a temporary minimum of F . The energy is then freed while the other parameters are fixed, and another minimization is performed. E_0 is fixed once more and the others freed to find the final minimum of F . No further iteration is needed, thanks in part to the good energy seed available from the fast fitter. The parameters from the final SIMPLEX call are returned as the answer.

9.6.2 π^0 and $\gamma\gamma$ hypotheses (two tracks)

Consider the two situations shown in the cartoon of Figure 9.16. The event has two γ tracks whose Cherenkov rings are represented in black. In case (a), the twelve fitter parameters are near the correct answer. In case (b), the parameters are such that both tracks are directed toward the dominant ring. Case (b) represents a local minimum which may offer a worse likelihood than case (a), yet the fitter is trapped. All small parameter changes result in an increased F . For example, sweeping one track over to the smaller ring involves passing through a region with little detected Cherenkov light. The intermediate states are disfavored, and the minimization code will have difficulty finding a lower F on the other side of those states, especially if three or four additional parameters (E_1, E_2 , etc.) must be adjusted simultaneously to realize the improvement.

The two-track fits require a minimization approach that avoids these traps in the likelihood surface. Monte Carlo π^0 events are used to identify trapping scenarios, and the minimization algorithm addresses all of these by brute force. Rather than enumerate the individual scenarios, we instead just describe the minimization algorithm. Its complexity is a reflection of the difficulty and importance¹² of finding the global minimum of $F(\mathbf{x})$.

Collecting seeds

The procedure begins with the collecting of seed parameters for the eventual SIMPLEX call.

- The conversion lengths s_1 and s_2 are seeded with either 50 cm or 250 cm, leading to four possible (s_1, s_2) pairs.

¹²since the fit results are used to reject π^0 events in the ν_e selection

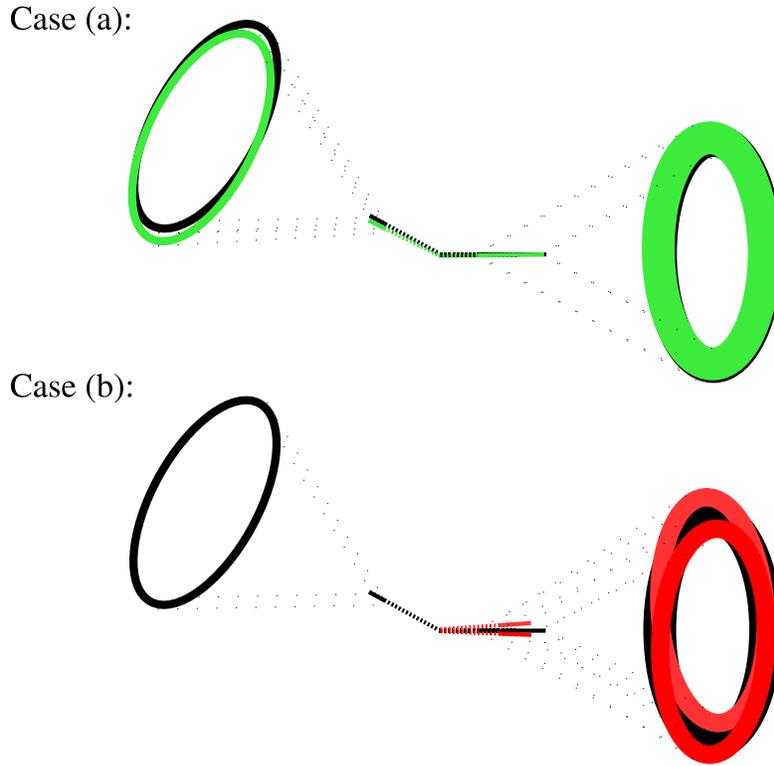


Figure 9.16: The two fitter configurations discussed in the text. The black rings represent the desired configuration. The overlaid colored rings show (a) a near-correct configuration and (b) a trapped configuration.

- θ_1 and ϕ_1 are seeded with the best-fit e -hypothesis (“ e -fit”) direction or with one of eight perturbations (Note 1 below).
- θ_2 and ϕ_2 are seeded with the best (Note 3 below) directions from a full grid of tested directions (Note 2).
- The 4-vertex of the event is seeded with the e -fit 4-vertex, shifted according to s_1 , θ_1 , and ϕ_1 .
- The seed for E_1 is approximately the e -fit energy, while the seed for E_2 is, when possible, the energy needed to give an invariant mass M_{π^0} . The energy seeds are actually based on M_{π^0} and a simulation-established empirical expression that gives the second ring’s fractional energy contribution to the e -fit energy as a function of the angle between the tracks.

Notes:

1. The nine possible seed directions for track 1 are created as follows. The spatial distribution of PMT charge is projected onto a plane perpendicular to the e -fit direction.

The covariance ellipse of the resulting 2D distribution, and in particular the direction of the ellipse's semimajor axis, is found. The nine possible track 1 seed directions are rotations of the best e -fit direction by $0, \pm 0.159, \pm 0.450, \text{ and } \pm 0.644$ (radians) parallel to the semimajor axis, and ± 0.159 perpendicular to the semimajor axis.

2. The track 2 grid search can be coarse (24 ϕ steps and 12 θ steps) or fine (50 and 25).
3. Three types of track 2 directions are saved in the grid search: the best total likelihood, the best *charge* likelihood when the rings have similar energies, and the best charge likelihood when the rings have dissimilar energies. Similarity is defined by $E_{\text{bigger}}/E_{\text{smaller}} < 2$. Additionally, the second best direction from each category is saved, leaving six directions in all. However, see Note 4.
4. When one of the eight alternate track 1 directions is being used, the track 2 grid search skips those directions which imply $E_{\text{bigger}}/E_{\text{smaller}} > 5$. In these highly asymmetric cases, the e -fit direction would be a good (unperturbed) description of the track 1 direction, so there is no reason to test these asymmetric cases while exploring alternate track 1 directions.

Each permutation of (s_1, s_2) gets its own set of these six best/second-best directions, leaving twenty-four seeding parameter sets in total. Each set goes through two Minuit sequences:

- The first sequence has two steps. (1) The energy and angle parameters are fixed and SIMPLEX is used to minimize F with respect to the other parameters. (2) All parameters are then freed, and a final SIMPLEX minimization is run.
- The second sequence differs in that only the energy parameters are held fixed in the first step.

In all forty-eight fit sequences, the final step is a fully free SIMPLEX call. That is, all of the above complexity serves only to provide forty-eight different seeds to this fully free fit. The best parameter set (lowest F) from the forty-eight final SIMPLEX calls is reported as the answer.

Fixed-mass option

The above procedure can be run with a constraint on the invariant mass. This is accomplished by removing E_2 as a free parameter, instead setting it in the likelihood function with

$$E_2 = \frac{M_{\pi^0}^2}{2E_1(1 - \cos \theta_{\text{tracks}})}, \quad (9.43)$$

where θ_{tracks} is the angle between the tracks. This fixed-mass mode is the actual π^0 hypothesis, whereas the free-mass mode allows for mass reconstruction. Both are used for π^0 identification, with the former lending its maximum likelihood \mathcal{L}_{π^0} and the latter providing the mass $M_{\gamma\gamma}$.

Speed

The algorithm was trimmed as much as possible while fulfilling the needs of the oscillation search. Indeed, each element of the minimization procedure is tackling a known trapping scenario. For uses other than π^0 rejection (*e.g.*, π^0 cross section studies), one may be willing to tolerate some loss in performance if CPU time is greatly decreased. The code provides user switching of many of the branches above, and such speedups are used for the π^0 rate measurement of Chapter 11.

Conversion constraint

The γ conversion points (but not the event vertex) are constrained to be within the main detector volume. This prevents the fit from removing a track's influence by inflating the conversion length. (If this were allowed, the π^0 hypothesis could always give a better likelihood than the electron hypothesis.)

9.7 Performance

We examine briefly the performance of the reconstruction algorithm by fitting simulated neutrino events. We apply the μ and e one-track fits to ν_μ and ν_e CC QE events, respectively, and we apply the π^0 and $\gamma\gamma$ two-track fits to NC π^0 events. The simulated events include the appropriate nuclear debris (p , n , γ) in addition to the μ , e , and π^0 tracks. Minimal selection cuts are applied to ensure that events are contained within the main detector region and, for plots not directly related to vertex reconstruction, within the fiducial volume of the analysis.

Figures 9.17 through 9.23 follow, with commentary included in the captions. Some typical performance numbers: vertex resolution, 20 cm; energy resolution, 12%; angular resolution, 3° ; π^0 mass resolution, 13%. Reconstruction biases are generally much smaller than these resolutions, although large biases are not intrinsically harmful to the analysis if they appear in data and simulation alike. (The important question of data/simulation agreement is addressed in Chapter 13.) The many performance metrics we have omitted in the interest of page count reveal nothing unexpected.

Supplement to the caption of Figure 9.19.

- Prompt Cherenkov light provides the fitter with most of its energy, angle, and position information. Thus, muon resolutions get considerably worse at lower energies where Cherenkov production is dropping.
- At low energy, it is difficult to learn anything about the π^0 momentum vector since the symmetric $\gamma\gamma$ final state is randomly oriented with respect to it. However, the Cherenkov light still provides a good π^0 vertex even at low energy.

- A π^0 vertex is harder to reconstruct than a ν_e CC QE vertex since the former is spatially displaced from the light production, owing to the ~ 50 cm γ conversion length.

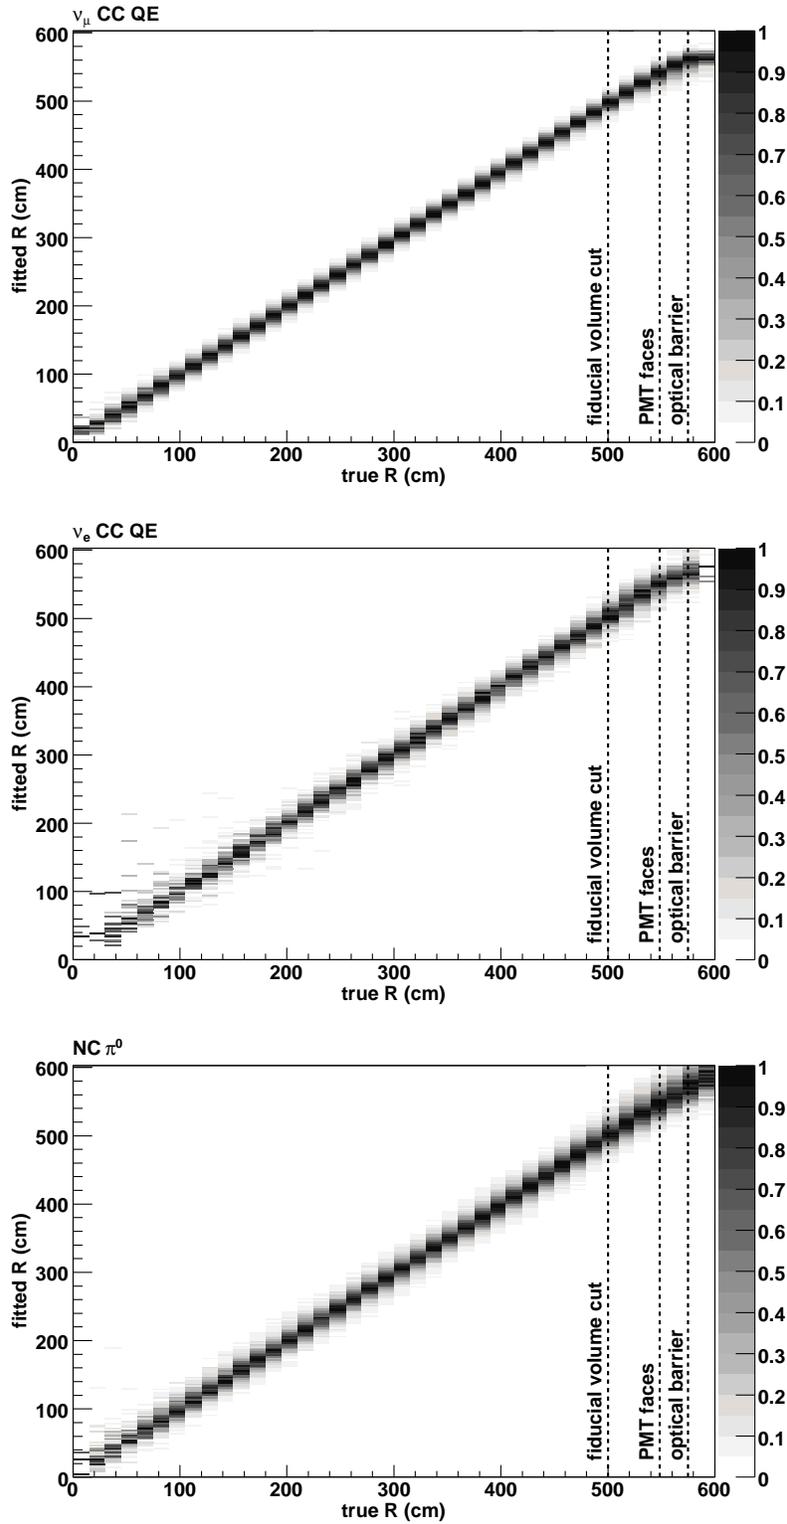


Figure 9.17: Reconstructed radial event position plotted against truth. The top, middle, and bottom panels show ν_μ CC QE, ν_e CC QE, and NC π^0 events. As expected, μ and e events that originate beyond the optical barrier reconstruct as if they begin just at the optical barrier. The fiducial volume cut $R < 500$ cm is used throughout the analysis (along with others).

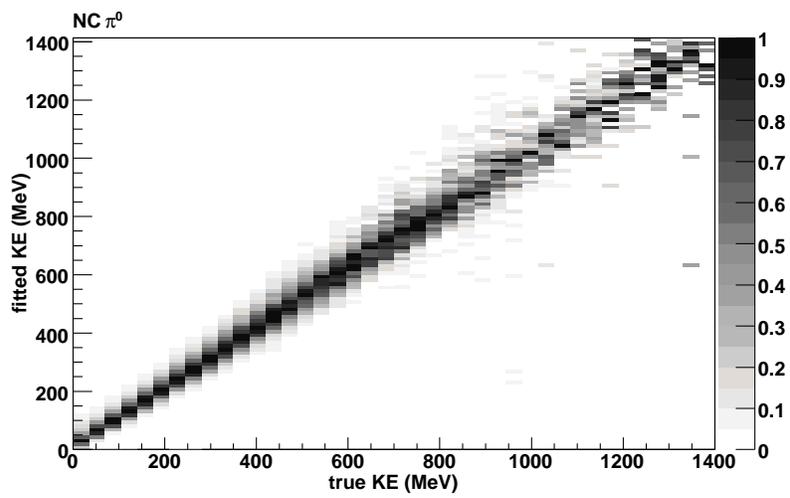
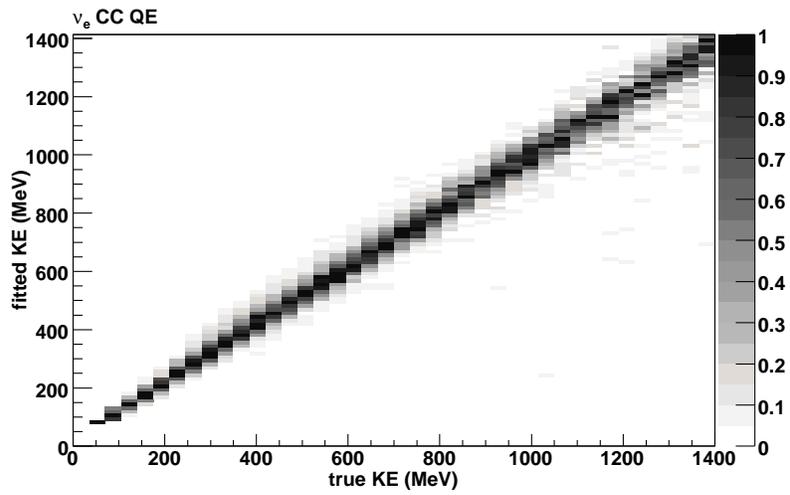
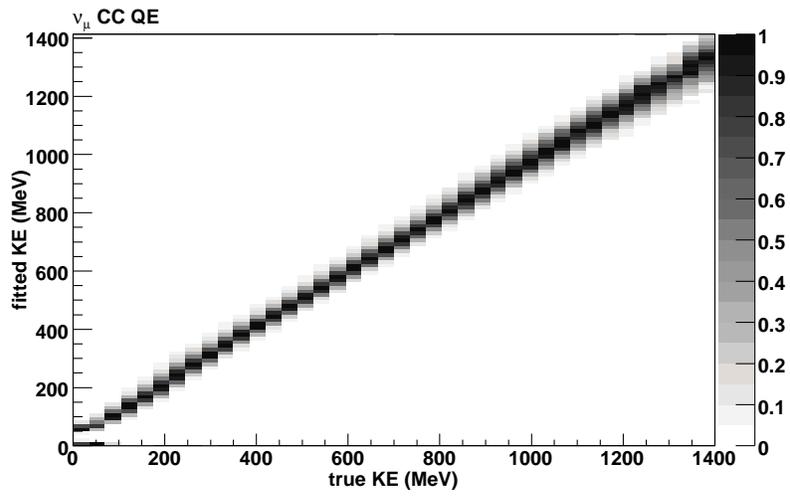


Figure 9.18: μ , e , and π^0 reconstructed kinetic energy plotted against truth.

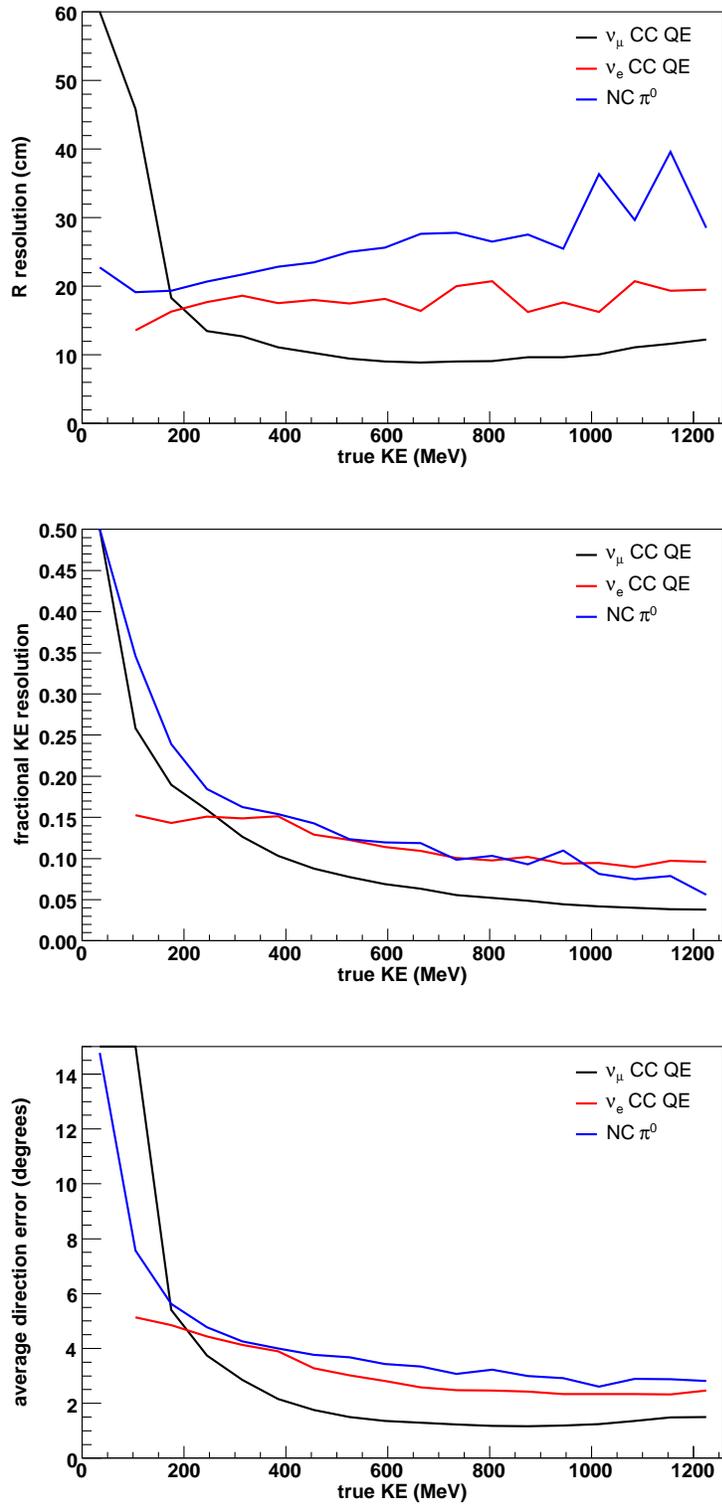


Figure 9.19: Radius, kinetic energy, and direction resolutions. The jitter is due to limited test sample statistics in some regions. The ν_e curves cut off at 100 MeV due to an upstream $N_{\text{main}} > 200$ cut. Also, see the supplement to this caption four pages back.

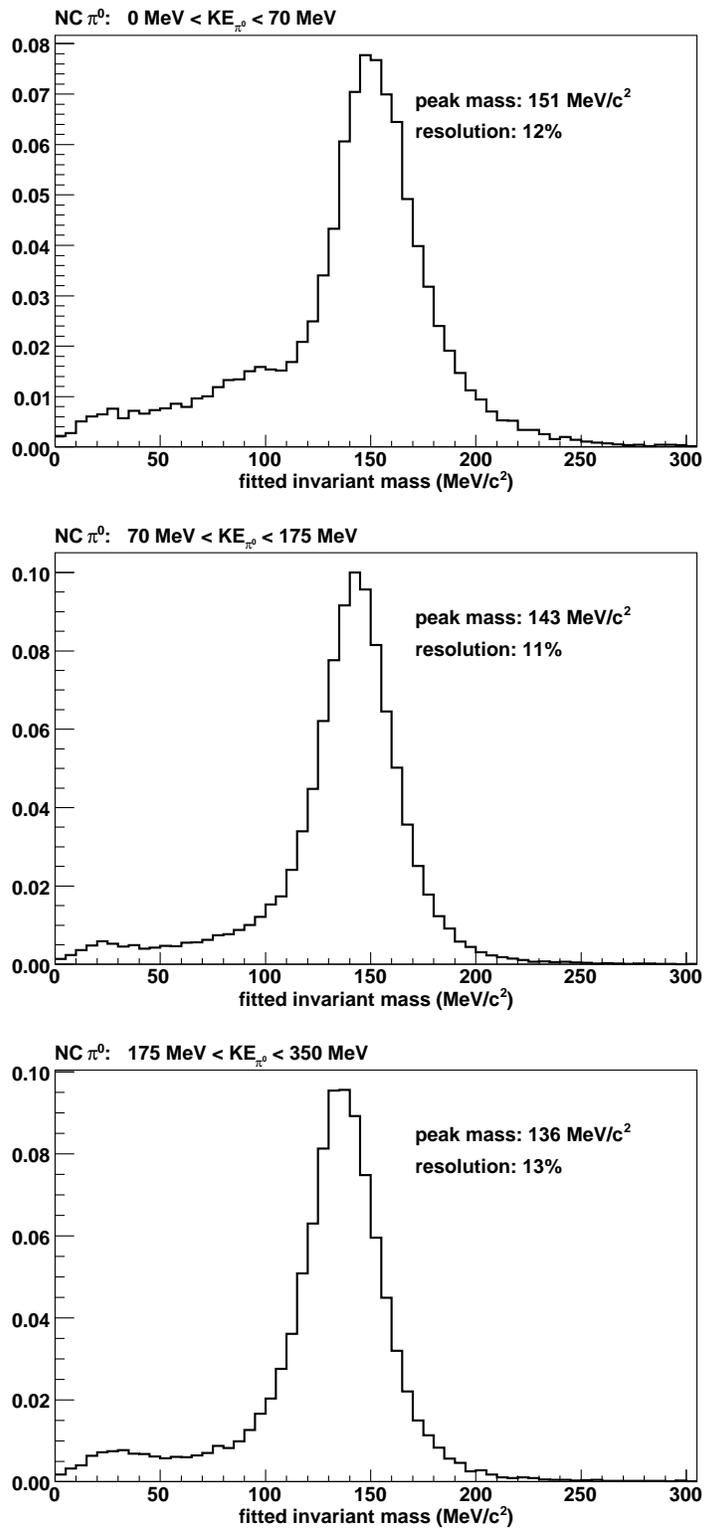


Figure 9.20: Fitted invariant mass $M_{\gamma\gamma}$ for NC π^0 events in three low energy regions. Continued in Figure 9.21.

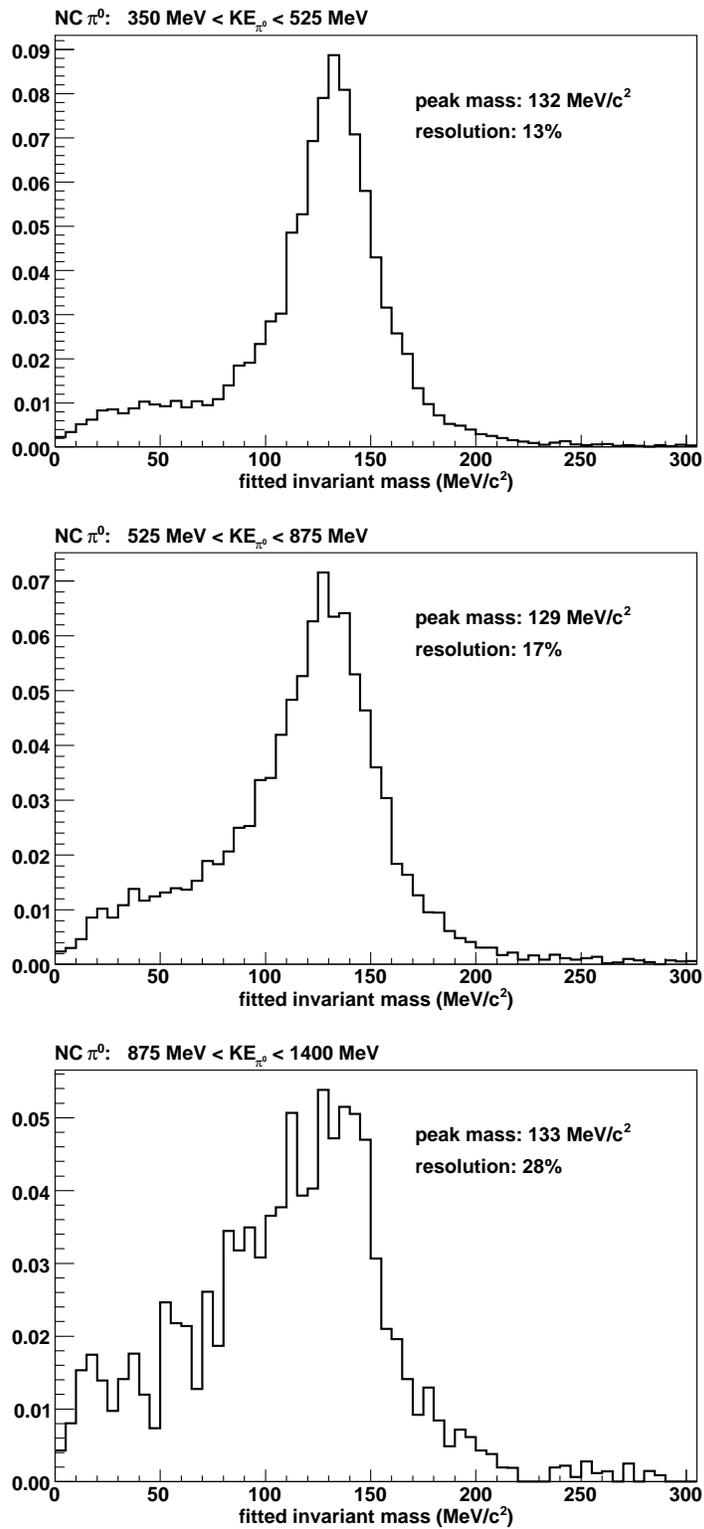


Figure 9.21: Fitted invariant mass $M_{\gamma\gamma}$ for NC π^0 events in three high energy regions. Continued from Figure 9.20.

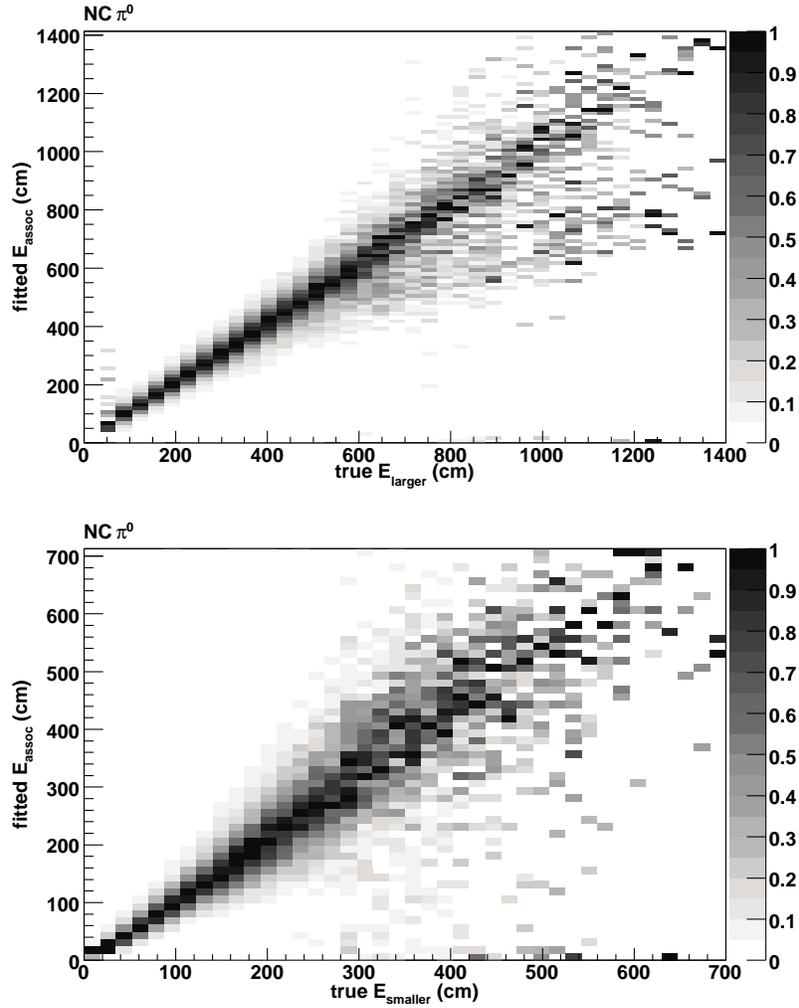


Figure 9.22: Reconstructed γ energies in NC π^0 events plotted against truth. The top (bottom) panel shows the higher (lower) energy γ from each event. The association of each fitted track to the underlying true γ 's is, in general, ambiguous. For these plots, we choose the assignment that gives the smaller combined energy and direction discrepancy (using a sort of least-squares figure of merit).

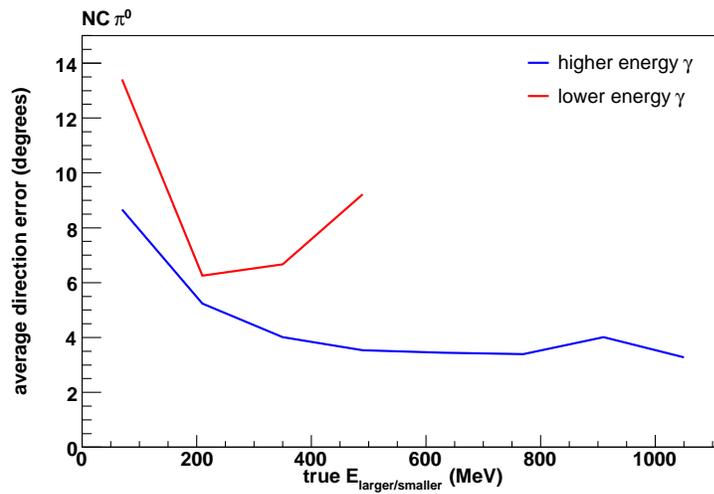
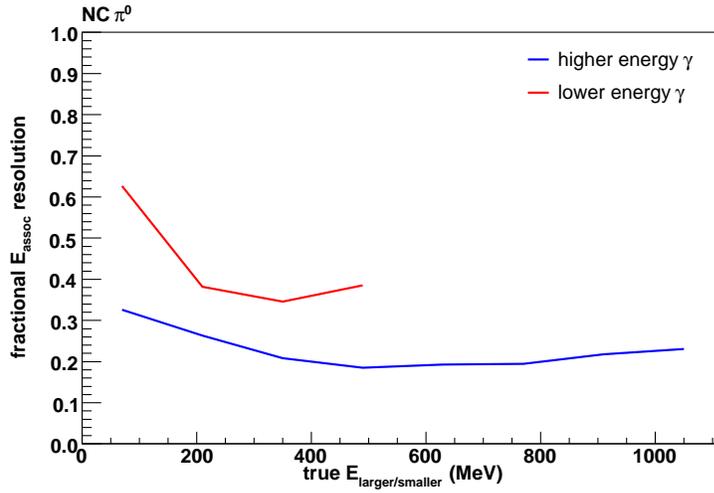


Figure 9.23: Energy (top) and direction (bottom) resolutions for γ 's in NC π^0 events. For a given true γ energy, a dominant γ has expectedly better resolution than does a subordinate γ . See the caption of Figure 9.22 for more.

Chapter 10

ν_μ charged current quasi-elastic scattering

ν_μ charged current quasi-elastic (CC QE) scattering, $\nu_\mu n \rightarrow \mu^- p$, accounts for 40% of MiniBooNE events. Its simple final state allows straightforward event reconstruction. Its ν_μ flux is directly related to the possible $\nu_\mu \rightarrow \nu_e$ signal flux. Its CC QE cross section is shared by the ν_e CC QE events being selected in the oscillation analysis. And, the largest component of the intrinsic ν_e flux ($\pi^+ \rightarrow \mu^+ \rightarrow \nu_e$) comes from the same π^+ 's that produce the ν_μ 's. All of these facts make the ν_μ CC QE sample an essential tool for understanding and constraining components of the oscillation analysis.

10.1 Selecting ν_μ CC QE events

If the μ^- produced in a ν_μ CC QE interaction comes to rest in the detector and decays, one typically sees two subevents, with the first subevent having too much light to be a Michel electron and with the second having an appropriate amount light for a Michel electron. This basic signature is identified with the following cuts (the numeric subscript specifies the subevent number):

- Exactly 2 subevents
- $(N_{\text{veto}})_1 < 6$
- $(N_{\text{veto}})_2 < 6$
- $(N_{\text{main}})_1 > 200$
- $(N_{\text{main}})_2 < 200$

Events near the edge of the detector can be influenced by PMT angular efficiencies, low wavelength light (which has short extinction length), reflections, veto modeling, and sundry other effects that are tough to simulate completely. We apply two cuts to the first subevent to keep events away from the tank edge:

- $R_\mu < 500$ cm
- (endpoint R_μ) < 500 cm .

The first cut limits the radial location of the reconstructed event vertex, while the second limits the location of the track endpoint as estimated from the reconstructed energy:

$$(\text{endpoint } R_\mu) = |\mathbf{R}_\mu + 2\Delta_{\text{mid}}(E_\mu)\mathbf{U}_\mu| , \quad (10.1)$$

where the subscript “ μ ” indicates that the muon hypothesis is used in the event reconstruction and where $\Delta_{\text{mid}}(E_\mu)$ is a range function built into the fitter that gives the mean distance from the vertex to the track midpoint. \mathbf{R}_μ and \mathbf{U}_μ are, respectively, the reconstructed 3-vertex and 3-direction of the muon track.

The PMT electronics take several hundred nanoseconds to recover normal charge response after a hit. This perturbs the identification and reconstruction of Michel electrons from short-lived muons. While this phenomenon is included in the simulation, a holdoff is enforced to minimize its importance:

- $t_2 - t_1 > 1000$ ns ,

where t_i is the time of the i -th subevent.

One final cut is applied to eliminate events with unexplained energy (usually from pions). The cut exploits the near-constant $\frac{dE}{dx}$ of muons in our energy range. Figure 10.1 shows the longitudinal separation

$$\Delta_{\parallel} = (\mathbf{R}_{2,e} - \mathbf{R}_{1,\mu}) \cdot \mathbf{U}_{1,\mu} \quad (10.2)$$

between the muon and Michel vertices plotted against the fitted muon energy $E_{1,\mu}$ for simulated ν_μ CC QE events (no backgrounds). The subscripts here indicate the subevent number and the reconstruction hypothesis. Figure 10.2 shows similar plots for all Monte Carlo and data events passing the pre-selection above, revealing the bit of background that can be removed. We apply the cut

$$|a_0 + a_1 E_{1,\mu} - \Delta_{\parallel}| < 50 \text{ cm} , \quad (10.3)$$

called the “line cut” below, which asks whether the reconstructed energy is consistent with the separation Δ_{\parallel} . A nice feature of the line cut is that it is self-calibrating: the coefficients a_0 and a_1 which describe a typical muon are determined empirically from the CC QE sample itself. The procedure:

1. Apply the pre-selection along with $\log(\mathcal{L}_e/\mathcal{L}_\mu) < 0$. (See §5.7.)
2. Create a profile version of the range *vs.* energy plot and fit a line to it.
3. Use the resulting fit parameters in a first-guess line cut like Eq. (10.3), except cutting at 70 cm rather than 50 cm.
4. With this temporary line cut in place, repeat (2) to establish the final coefficients a_0 and a_1 .

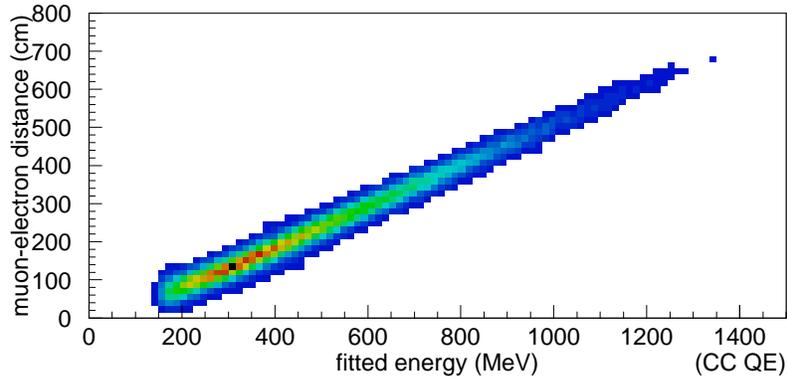


Figure 10.1: Range *vs.* energy for a pure sample of simulated ν_μ CC QE events. The vertical axis gives the longitudinal separation between the muon and Michel vertices. The horizontal axis gives the fitted energy of the muon. The linear behavior is the basis for the final selection cut.

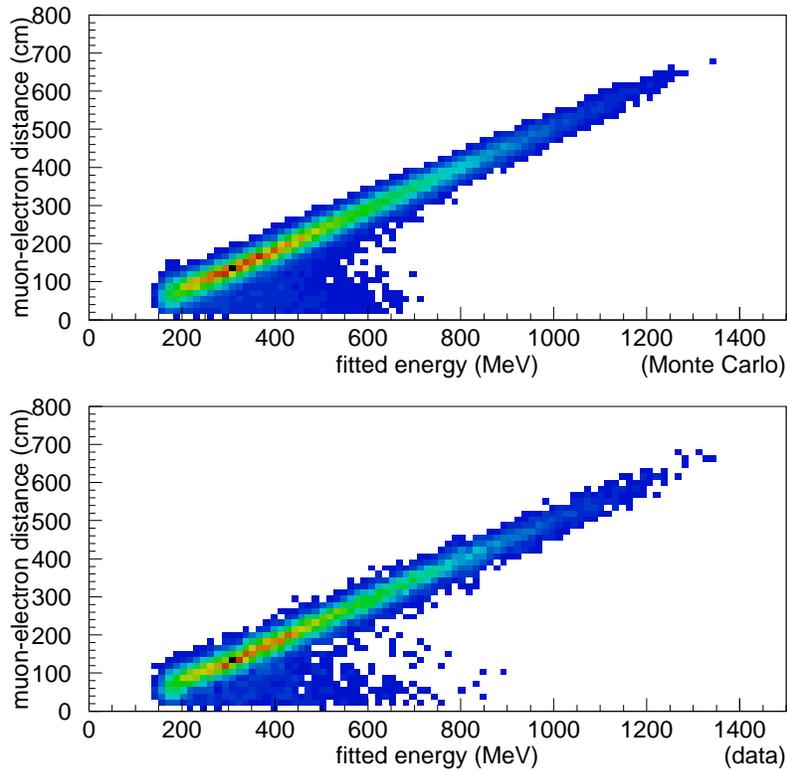


Figure 10.2: Range *vs.* energy for events passing the pre-selection in Monte Carlo (top) and data (bottom). CC π events typically have more energy than is expected given the measured muon range, so they lie to the right of the CC QE line.

A single iteration is sufficient. The linear fit is performed over the kinetic energy range 250–1200 MeV and is demonstrated for data in Figure 10.3. Table 10.1 shows the coefficients obtained from (and used for) data, the default Monte Carlo, and several Monte Carlo variants. Figure 10.4 compares the line cut discriminant in data and the simulation.

This selection yields a 71% pure ν_μ CC QE sample. 55% of the background events have QE-like final states, consisting of a muon, nucleonic debris, possible low energy photons, and (in particular) no pions.

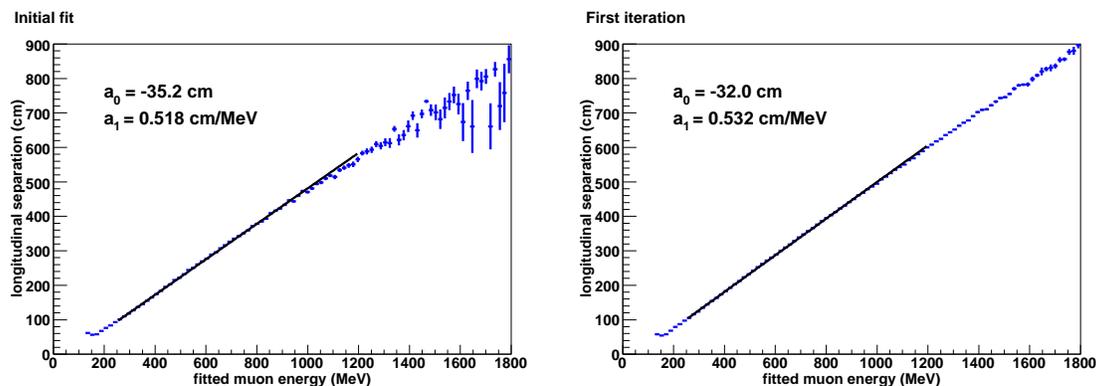


Figure 10.3: Establishing the cut coefficients a_0 and a_1 in data. The initial fit (left) and the iteration (right) are shown.

sample	a_0 (cm)	a_1 (cm/MeV)
data	-32.0	0.532
standard MC	-34.8	0.539
MC variant 35	-28.1	0.541
MC variant 36	-28.7	0.548
MC variant 37	-32.9	0.543
MC variant 38	-30.9	0.538
MC variant 39	-32.9	0.535

Table 10.1: Cut coefficients for different realities. In the variants, detector optical properties and nuclear final state interactions are adjusted within their assigned uncertainties.

10.2 Using the sample

10.2.1 Q^2 dependence of the CC QE cross section

As described in Chapter 7, the CC QE cross section in `nuance` has a Q^2 dependence governed by nucleon form factors and nuclear phenomena. We calculate the approximate Q^2 for an event by assuming a stationary target nucleon:

$$-Q^2 = q^2 = M_l^2 - 2E_\nu^{\text{rec}} (E_l - \mathbf{p}_l \cdot \mathbf{u}_\nu) , \quad (10.4)$$

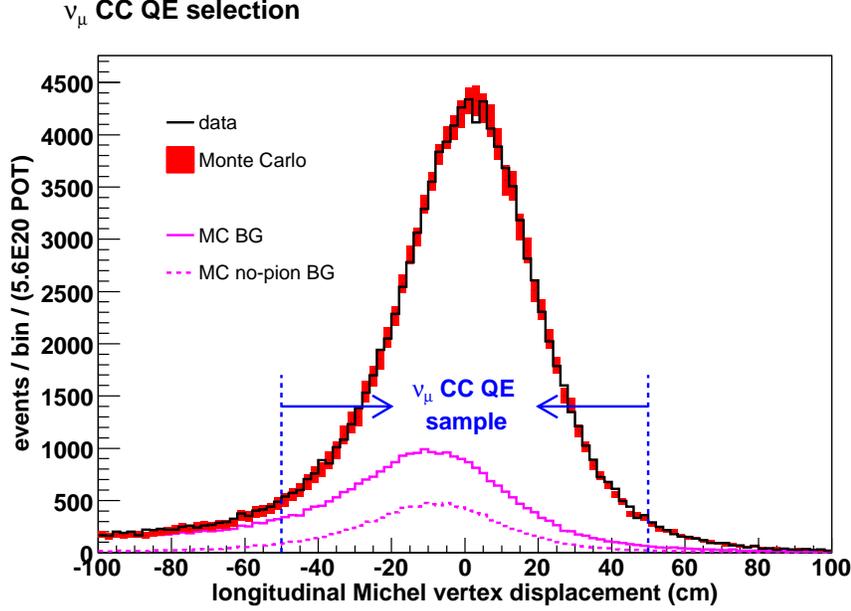


Figure 10.4: Data (solid black curve) and Monte Carlo (red band) distributions for the line cut residual of Eq. (10.3). The red band indicates the systematic error in the simulation. The solid magenta histogram shows the background (non-CC QE events), and the dashed magenta histogram shows the background events that contain no pions (mostly CC π events in which the π gets absorbed before leaving the nucleus.) The latter is nearly indistinguishable from CC QE in the detector.

where (E_l, \mathbf{p}_l) is the reconstructed lepton 4-momentum, M_l is the lepton mass, \mathbf{u}_ν is the incident neutrino direction, and E_ν^{rec} is the reconstructed neutrino energy from Eq. (7.9). Figure 10.5, discussed further below, shows the observed Q^2 distribution.

The nuance cross section has a single form factor parameter – M_A , the axial vector mass. M_A has been measured ($M_A \approx 1.0$ GeV) with an error that is likely no smaller than 10%, although the situation is murky. (See Ref. [134] for a summary of past measurements.) Further, the interpretation of M_A depends on the particular Q^2 parametrizations used for the nucleon form factors (§7.1). We bypass these issues by tuning M_A to obtain Q^2 agreement in our own Monte Carlo and data ν_μ CC QE samples. For $Q^2 \lesssim 0.2$ GeV², nuclear effects also influence the Q^2 distribution. Deficiencies in nuance’s RFG nuclear model [106] can be mitigated by adjusting the Pauli blocking parameter κ introduced in Eq. (7.8).

Since the effect of κ is restricted to low Q^2 values, M_A is determined first via a least-squares fit to the $Q^2 > 0.3$ GeV² portion of the distribution. The covariance matrix used in the fit accounts for all non- M_A sources of error, with detector model uncertainties dominating. The fit result: $M_A = 1.15 \pm 0.07$.

With M_A frozen at this value¹, a fit for κ is performed to the $Q^2 < 0.3$ GeV² region. This

¹though later freed to confirm that M_A - κ correlations are negligible

is not a least-squares fit, though. While κ provides a mechanism for adding or removing low- Q^2 events, we cannot expect it to perfect the shape of the low- Q^2 part of the distribution. Thus, a χ^2 fit can give underestimated errors, as κ may provide insufficient freedom to respond to shape uncertainties in the covariance matrix. However, the eventual $\nu_\mu \rightarrow \nu_e$ signal extraction procedure depends only on reconstructed neutrino energy, which in turn is fairly insensitive to the fine structure of the Q^2 distribution, as Figure 10.6 demonstrates. Thus, κ and its uncertainty are set by examining the total number of low- Q^2 events

$$I = \int_0^{0.25 \text{ GeV}^2} f(Q^2) dQ^2, \quad (10.5)$$

where $f(Q^2)$ is the Q^2 distribution. κ is chosen to give $I_{\text{data}}=I_{\text{MC}}$ (requiring $\kappa=1.024$), and its uncertainty ($\sigma_\kappa=0.02$) is chosen to produce a spread in I equal to the spread caused by all other sources of error (detector model, background levels, etc.)

Figure 10.5 shows the Q^2 distribution for data, the tuned simulation ($M_A=1.15 \pm 0.07 \text{ GeV}$, $\kappa=1.024 \pm 0.02$), and the standard nuance simulation ($M_A=1.03 \text{ GeV}$, $\kappa=1$). The post-tuning agreement is good. An independent extraction of the two CC QE parameters arrives at consistent results with a fairly different fit procedure ($M_A=1.23 \pm 0.20 \text{ GeV}$, $\kappa=1.019 \pm 0.011$) [135].

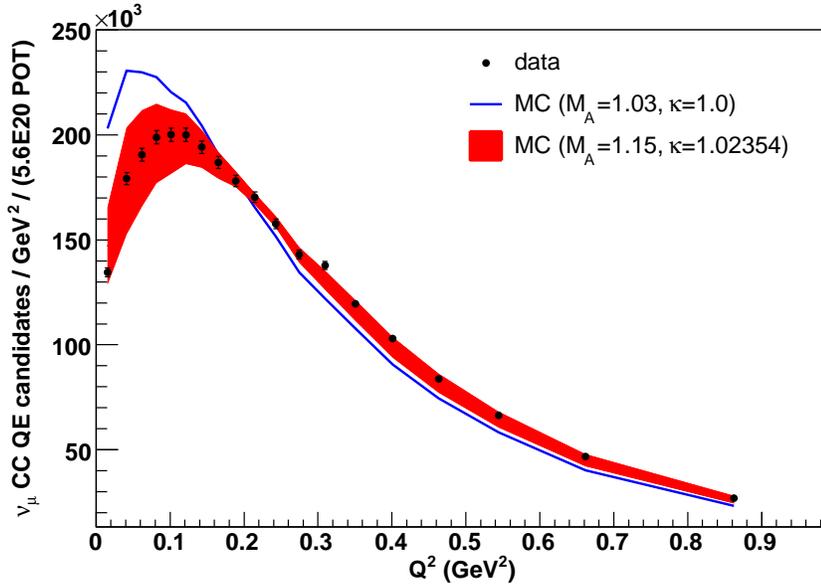


Figure 10.5: Distribution of Q^2 for ν_μ CC QE candidate events. The data (black points) and tuned simulation (red band) are shown along with the distribution obtained using default nuance parameters (blue curve). The red band indicates the range of systematic variation from all cross section uncertainties.

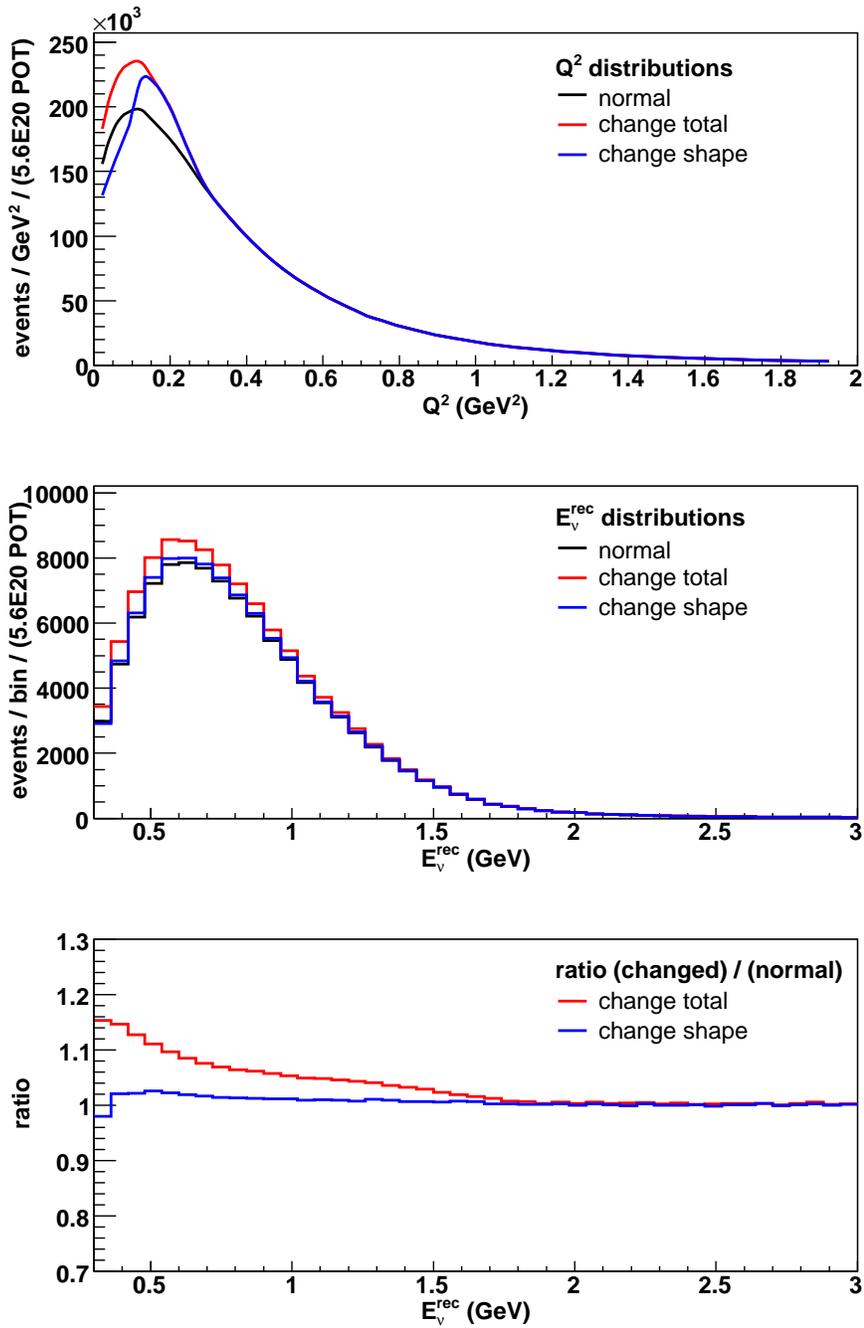


Figure 10.6: Insensitivity of the E_{ν}^{rec} distribution to the low- Q^2 shape. (Top panel) Three reconstructed Q^2 distributions are created by weighting (true) low- Q^2 events in various ways. The “change total” version has a 20% excess of low- Q^2 events over the normal case, while the “change shape” version keeps Eq. (10.5) constant (within 3%) but has a different Q^2 shape. (Middle panel) The corresponding E_{ν}^{rec} distributions. (Bottom panel) Ratios of these. The low- Q^2 shape change has negligible effect on the E_{ν}^{rec} distribution.

10.2.2 E_ν dependent corrections

We now build generalized, energy-based corrections that improve our Monte Carlo predictions in other samples by taking advantage of correlations between those samples and this ν_μ CC QE one. To begin, we create a map connecting the reconstructed neutrino energy E_ν^{rec} in ν_μ CC QE events with the underlying true neutrino energies E_ν^{true} . To build this map, we populate a two-dimensional E_ν^{rec} -vs.- E_ν^{true} histogram with Monte Carlo ν_μ CC QE events. Each E_ν^{rec} row of the resulting histogram is then unit normalized. Figure 10.7 shows the result. At the same time, we form one-dimensional E_ν^{rec} histograms for data and for Monte Carlo, as shown in Figure 10.8.

Let R_i represent the height of the i -th bin of the E_ν^{rec} histogram. Also, let \mathbf{M} be the matrix incarnation of the row-normalized 2D map. For the Monte Carlo sample, the quantities

$$T_j = \sum_i M_{ij} R_i \quad (10.6)$$

are simply the bin heights of the true energy (E_ν^{true}) histogram that we could have formed directly. For data, $\{T_j\}$ gives an *inferred* E_ν^{true} histogram which we can be compared to the Monte Carlo version. Figure 10.9 shows the inferred true energy histograms for data and Monte Carlo along with their ratio. This ratio forms our E_ν^{true} -dependent correction

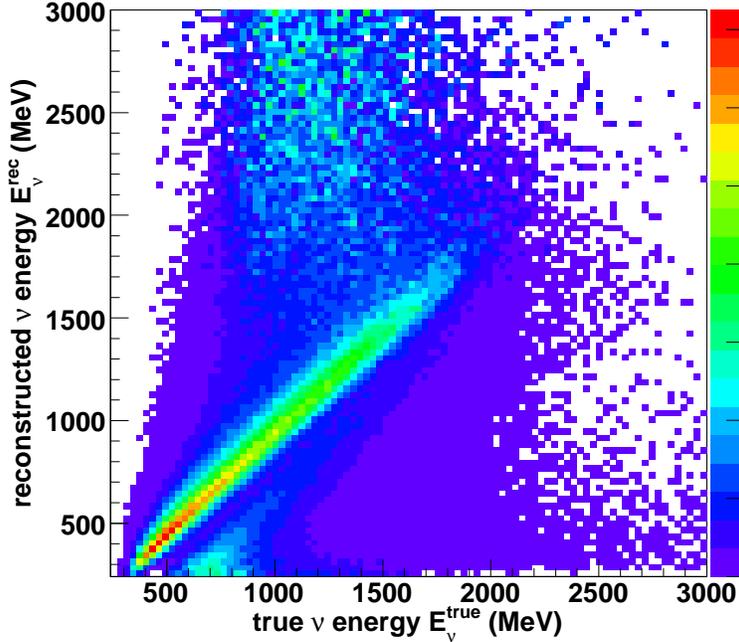


Figure 10.7: Row-normalized 2D histogram of E_ν^{rec} vs. E_ν^{true} for ν_μ CC QE events in the Monte Carlo simulation.

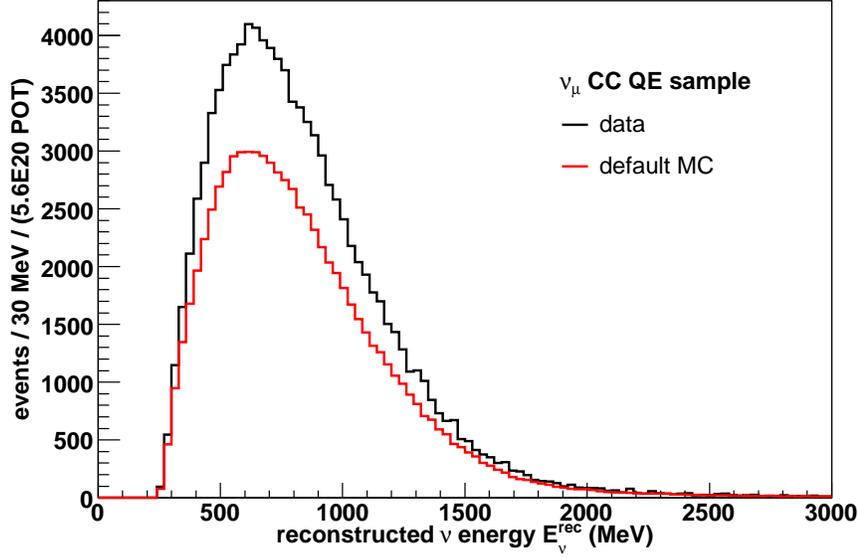


Figure 10.8: Data and (uncorrected) Monte Carlo histograms of E_ν^{rec} . The $\sim 30\%$ discrepancy, which is within uncertainties, is what we wish to measure and correct for the benefit of other event samples.

function

$$f(E_j) = \frac{T_j^{\text{data}}}{T_j^{\text{MC}}}, \quad (10.7)$$

or, removing the discrete notation, $f(E_\nu^{\text{true}})$.

We note that this is not an unsmearing procedure, which would involve inverting a response matrix similar to \mathbf{M} . Such an approach has significant limitations in problems with statistical noise (such as this one). Rather, the matrix \mathbf{M} acts as a sort of discrete Green’s function, providing a series of E_ν^{true} probability density functions which are weighted and summed according to the observed E_ν^{rec} distribution. The power this procedure has for correcting the Monte Carlo is governed by the energy resolution apparent in \mathbf{M} .

Using the correction function

The simplest application of $f(E_\nu^{\text{true}})$ is to weight each ν_μ CC QE event in the simulation according to the function. We actually use $f(E_\nu^{\text{true}})$ more generally, weighting

- all ν_μ CC events,
- all ν_μ NC events *except* those containing π^0 ’s or radiative Δ decays (next chapter), and
- $\nu_\mu \rightarrow \nu_e$ oscillation events,

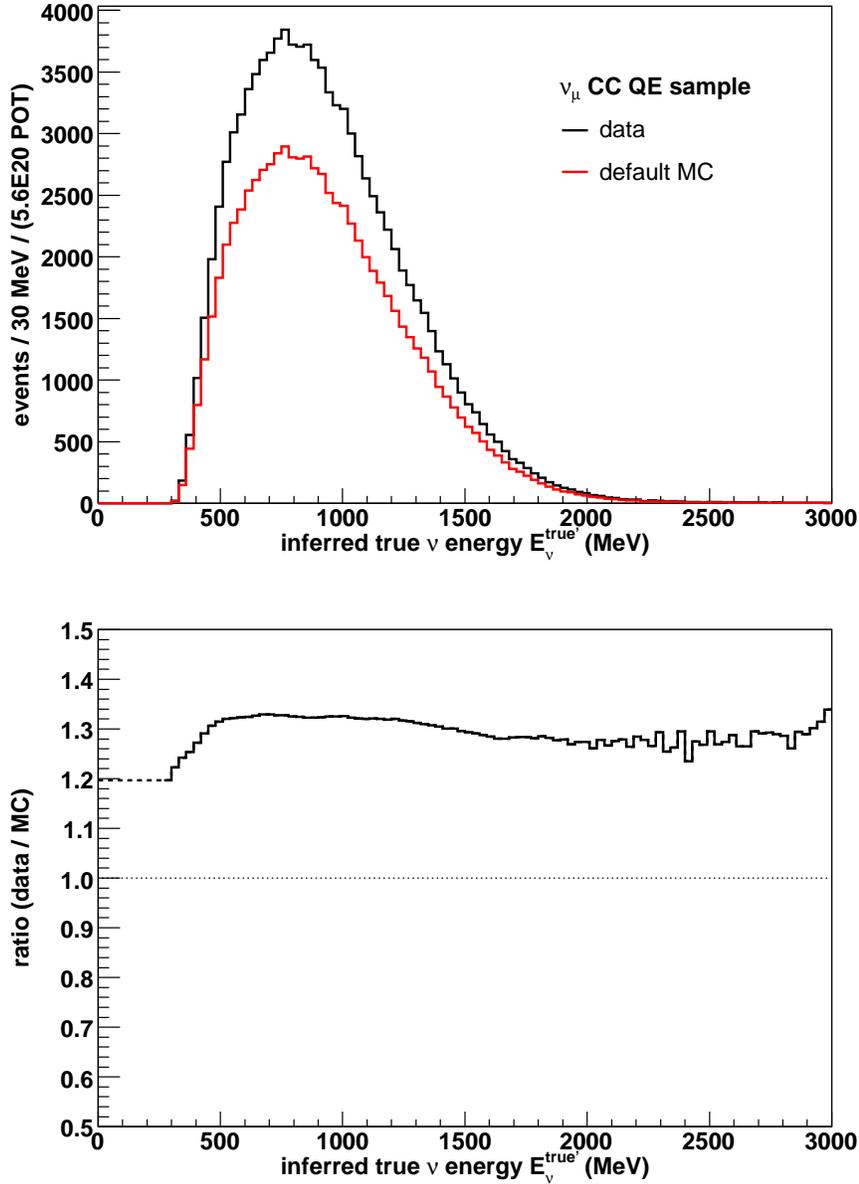


Figure 10.9: (Top) Inferred true neutrino energy distributions for data and Monte Carlo. (Bottom) Their ratio $\Leftrightarrow f(E_\nu^{\text{true}})$. At low E_ν^{true} where no events pass the selection cuts, $f(E_\nu^{\text{true}})$ is extended to the left as shown by the dotted line. A similar extension occurs to the right of 3000 MeV. (This is a technical matter only, as few neutrinos enter any relevant sample from these extremes.)

with the additional exception that events occurring outside the detector – so-called “dirt” events – are not weighted (next chapter). The desire to correct more than ν_μ CC QE events motivates the translation of E_ν^{rec} into E_ν^{true} . (Otherwise, an E_ν^{rec} -based weighting would suffice.)

The mapping \mathbf{M} is perfect for Monte Carlo (from which it is derived), but it is only approximate for data, as \mathbf{M} is built with particular flux, cross section, and detector model assumptions which are subject to systematic uncertainty. Before describing how we handle this complication, we must first outline the primary error propagation approach used in the oscillation analysis. A more complete discussion comes in Chapter 12.

Uncertainties and $f(E_\nu^{\text{true}})$

To turn uncertainties for some component of the simulation (say, π^+ production in the target) into uncertainties on quantities of interest (say, the ν_e sample’s E_ν^{rec} histogram), we do the following. We first pick a random model variation within the range specified by the relevant covariance matrix. So for π^+ uncertainties, we choose a random Sanford-Wang parameter set according to the parameters’ 7×7 covariance matrix. The Monte Carlo variant so produced gives us a predicted E_ν^{rec} distribution for the ν_e sample (or anything else) that differs from the default prediction. Repeating this many times creates an ensemble of prediction histograms from which we build the desired covariance matrix. Other sources of error yield similar matrices, and all of these are added together.

It is a simple matter to incorporate the correction procedure here: for each Monte Carlo variant, a new ν_μ CC QE analysis is performed, with $f(E_\nu^{\text{true}})$ established from scratch using a new map \mathbf{M} and a new histogram $\{R_i\}$. (Data does not change.) That is, each Monte Carlo variant is corrected according to its own ν_μ CC QE sample. Correlations between the ν_μ CC QE and downstream target samples play out variant-by-variant. That is, modifying the flux parameters (say) induces correlated changes in the two samples, implying that one sample can be used to improve our knowledge of the other.

Figure 10.10 shows the function $f(E_\nu^{\text{true}})$ obtained from and used for a handful of Monte Carlo π^+ cross section variants. Figure 10.11 shows the change in the predicted $\nu_\mu \rightarrow \nu_e$ E_ν^{rec} spectrum and its uncertainty before and after the (ν_μ CC QE)-established correction is applied. Note that this is not the final uncertainty; other covariance matrices get added to this π^+ one.

10.2.3 ν_e from μ decay

The largest background in the oscillation search comes from intrinsic beam ν_e ’s produced in μ^+ decay. The μ^+ ’s themselves come from the decays of the same π^+ ’s that provide the ν_μ flux. Figure 10.12 shows this decay sequence graphically. The resulting ν_e flux is only 0.4% as large as the related ν_μ flux thanks to the long lifetime of the μ^+ .² However, 0.4% is quite significant when compared to the LSND oscillation probability.

²Most of the μ^+ ’s hit the 50 m absorber before decaying. The fraction that decay in-flight will be $\mathcal{O}\left(\frac{50 \text{ m}}{\gamma\beta c\tau}\right)$, which for $E_\mu=1 \text{ GeV}$ is $\mathcal{O}(0.008)$.

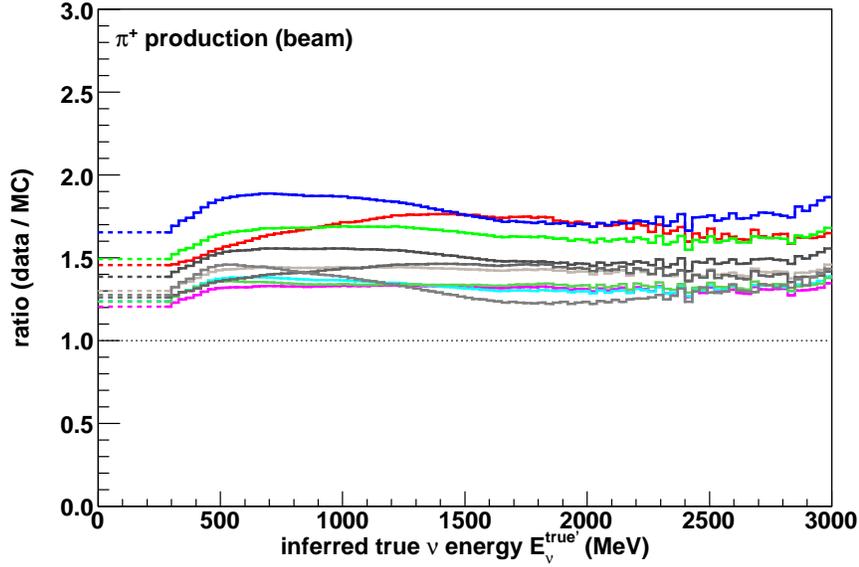


Figure 10.10: $f(E_\nu^{\text{true}})$ for ten Monte Carlo variants in which the π^+ Sanford-Wang parameters were altered as allowed by their uncertainties. Each $f(E_\nu^{\text{true}})$ is used to correct the Monte Carlo from which it was derived.

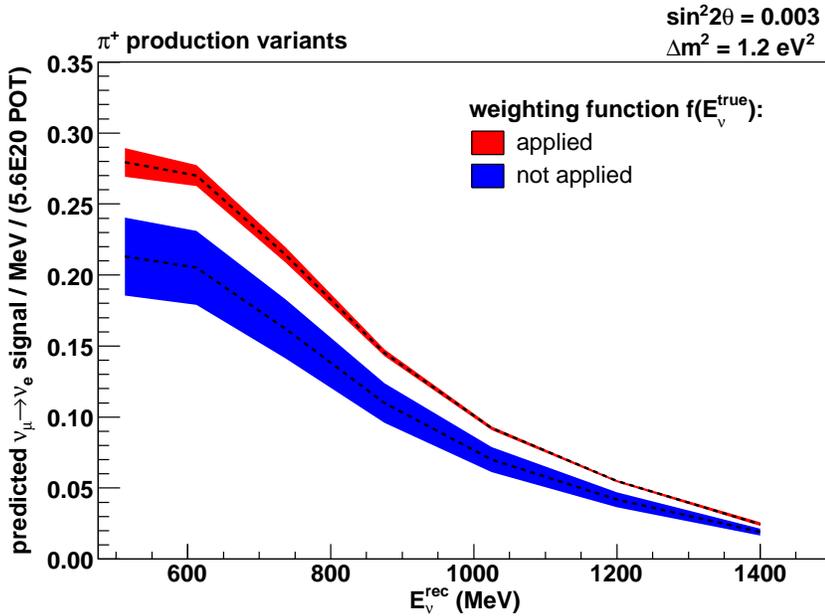


Figure 10.11: Predicted E_ν^{rec} spectrum for $\nu_\mu \rightarrow \nu_e$ events ($\sin^2 2\theta=0.003$, $\Delta m^2=1.2 \text{ eV}^2$) after oscillation cuts and with/without (red/blue) the ν_μ CC QE constraint. The error bands show the uncertainty due to π^+ production.

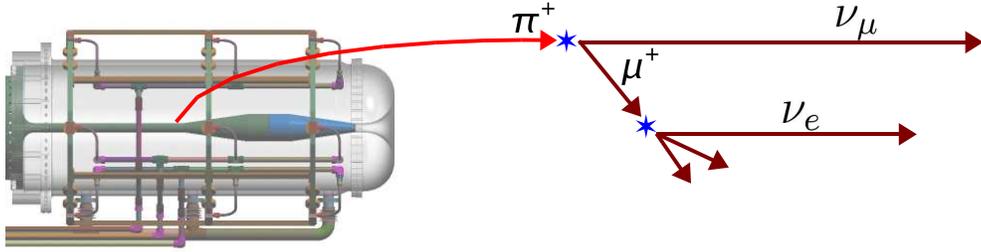


Figure 10.12: Cartoon of the production of ν_e from μ^+ decay. We use the observed ν_μ energy spectrum to improve our knowledge of the π^+ spectrum (and, in turn, the μ^+ and ν_e spectra) as described in the text.

Because the MiniBooNE detector subtends a small angle when viewed from the decay region, a ν_μ that reaches the detector must have come from a forward decay of a typically forward-going π^+ . This geometric constraint introduces a strong correlation between the energy E_{ν_μ} of a detected ν_μ and the energy E_π of its parent π^+ . For perfectly forward trajectories, the ν_μ has energy

$$\begin{aligned} E_{\nu_\mu} &\approx \left(1 - \frac{m_\mu^2}{m_\pi^2}\right) E_\pi \\ &\approx 0.43 E_\pi, \end{aligned} \tag{10.8}$$

where m_μ and m_π are the μ^+ and π^+ masses, and where we have neglected terms of $O\left(\frac{m_\pi^2}{E_\pi^2}\right)$. Figure 10.13 shows the relationship between the *reconstructed* ν_μ energy E_ν^{rec} and the energy E_π^{true} of the parent π^+ in the simulation. The 2D histogram in the figure serves as a new response matrix \mathbf{M}' that we use to create an inferred E_π^{true} spectrum from the observed E_ν^{rec} spectrum, in exact analogy with the $f(E_\nu^{\text{true}})$ formalism. A weighting function $g(E_\pi^{\text{true}})$ is created from the ratio of data and Monte Carlo E_π^{true} distributions. We use $g(E_\pi^{\text{true}})$ to weight ν_e -from- μ^+ events in the simulation according to their parent π^+ energies. Few systematic uncertainties affect the $\pi^+ \rightarrow \mu^+ \rightarrow \nu_e$ portion of the ν_e prediction, so correcting the π^+ energy spectrum goes a long way toward reducing ν_e -from- μ^+ uncertainty. Figure 10.14 shows the resulting shift and error reduction in the predicted E_ν^{rec} spectrum for ν_e -from- μ^+ events in the oscillation sample. As before, this shows only one of several error components.

10.2.4 Closing

The (ν_μ CC QE)-based $f(E_\nu^{\text{true}})$ and $g(E_\pi^{\text{true}})$ corrections are applied any time a Monte Carlo prediction (and corresponding systematic uncertainty) is produced. It is as if we have made flux and/or cross section and/or detector response measurements with this sample and have adjusted the simulation accordingly. Embedding the measurements within the error propagation system as we have accounts for the inherent correlated uncertainties.

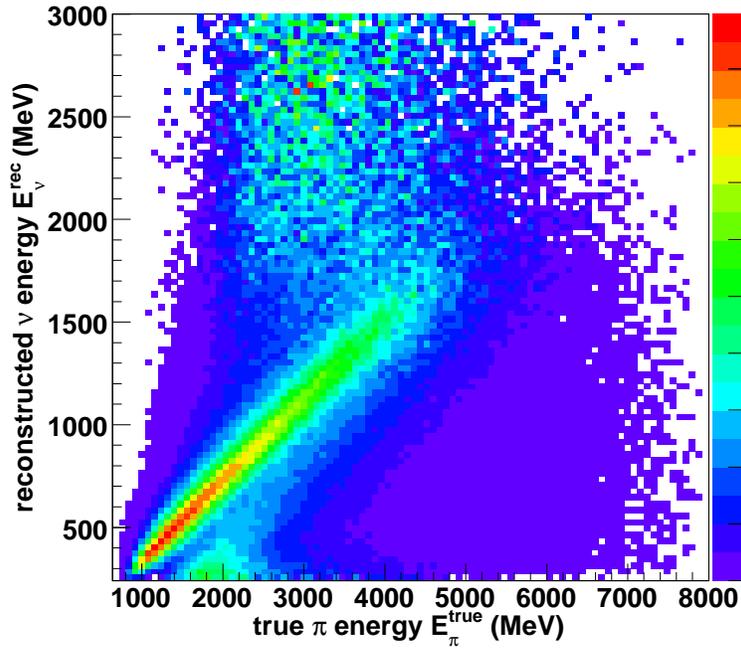


Figure 10.13: Row-normalized 2D histogram of E_ν^{rec} vs. E_π^{true} for ν_μ CC QE events built from the Monte Carlo simulation.

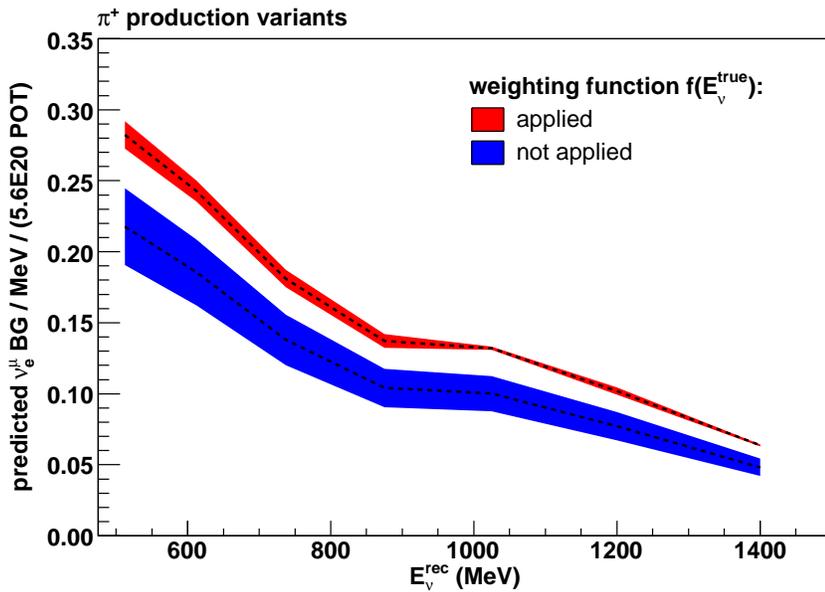


Figure 10.14: Predicted E_ν^{rec} spectrum for ν_e -from- μ^+ events passing oscillation cuts with/without (red/blue) the ν_μ CC QE constraint. The error bands show the uncertainty due to π^+ production.

Chapter 11

Neutral current π^0 events

11.1 π^0 as a $\nu_\mu \rightarrow \nu_e$ background

A π^0 produced in MiniBooNE quickly decays to two photons,¹ each of which, upon conversion² to an e^+e^- pair, produces an electromagnetic shower and corresponding Cherenkov ring. These γ -induced rings are indistinguishable from e -induced ones, so rejection of neutral current π^0 events in the ν_e selection relies on the presence of two distinct rings.

A typical π^0 momentum in MiniBooNE is ~ 0.3 GeV/ c . Thus, the $\gamma\gamma$ final state, symmetric in the π^0 rest frame, is somewhat boosted in the lab frame. If one photon is significantly lower in energy than the other, its Cherenkov ring may go unnoticed, resulting in misidentification. A detailed two-photon likelihood function and a persistent likelihood maximization algorithm (§9.6.2) combat the π^0 background, the largest ν_μ -induced background in the oscillation analysis.

A second avenue for misidentification arises when one of the two photons does not interact within the detector volume. Figure 11.1 shows two dominant scenarios for this: (1) A neutrino interaction in the dirt surrounding the detector produces a π^0 , and one of the decay photons passes invisibly through the veto region and converts within the fiducial volume. (2) A π^0 produced near the tank wall produces an outward going photon which exits the detector unnoticed while the remaining photon converts within the fiducial volume.

The cross section for neutral current π^0 production is not well known at MiniBooNE energies (Figure 7.4). Further, the dominant resonant channel is accompanied by a subdominant coherent channel ($\nu C \rightarrow \nu C \pi^0$) and final state interactions (*e.g.*, charge exchange in the nucleus) which introduce additional uncertainty. We avoid all this by directly measuring the inclusive rate of neutral current π^0 production in the detector using cleanly identified π^0 events.

¹ $c\tau = 25$ nm, B.R. ($\pi^0 \rightarrow \gamma\gamma$) = 98.8%

²radiation length $X_0/\rho = 52.4$ cm

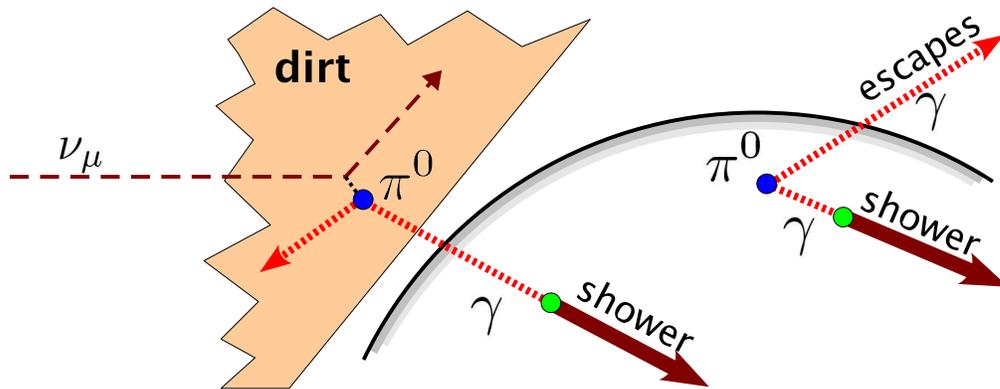


Figure 11.1: Two scenarios for getting a single γ -induced shower from a π^0 . (Left) A neutrino interaction in the dirt can lead to a single photon converting in the tank. (Right) $\pi^0 \rightarrow \gamma\gamma$ near the tank wall can result in one photon exiting the tank before interacting, leaving the other behind.

11.2 π^0 event selection

The neutral current π^0 selection begins with a basic set of cuts akin to those in the ν_μ CC QE selection (§10.1) :

- Exactly 1 subevent (no Michel electron expected)
- $N_{\text{veto}} < 6$
- $N_{\text{main}} > 200$
- $R_e < 500$ cm .

About half the events passing these cuts are ν_μ CC QE interactions in which the Michel decay was not seen. This CC contamination is easily removed by recognizing the muon-like appearance of the event via the ratio of the e - and μ -hypothesis reconstruction likelihoods:

- $\log(\mathcal{L}_e/\mathcal{L}_\mu) > 0.05$.

Figure 11.2 shows the $\log(\mathcal{L}_e/\mathcal{L}_\mu)$ distribution, with the NC π^0 and ν_μ CC QE components indicated. Note that $\log(\mathcal{L}_e/\mathcal{L}_\mu)$ naturally separates μ from e , not μ from π^0 . The fact that π^0 events “look more electron like than muon like” (*i.e.*, that they typically have $\log(\mathcal{L}_e/\mathcal{L}_\mu) > 0$) is not necessarily obvious, although it might be expected since a π^0 event is made up of electron-like rings and since two diffuse rings would not likely mimic a single sharp muon-like ring.³

³Ideally, we would just use $\log(\mathcal{L}_{\pi^0}/\mathcal{L}_\mu)$ to remove muons. However, it is advantageous to eliminate as much background as possible before running any π^0 fits, as they are rather CPU intensive.

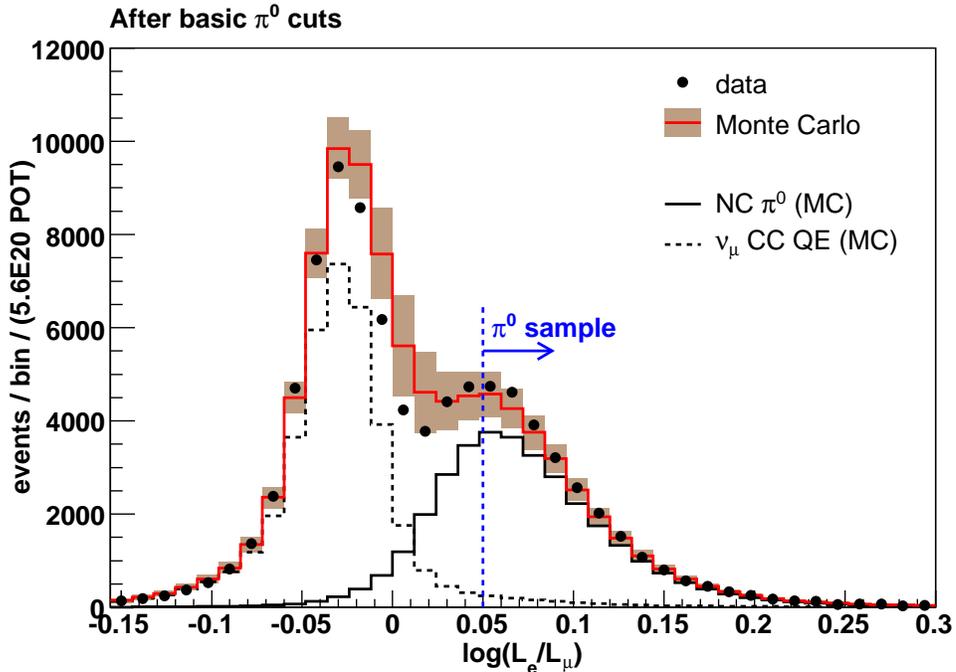


Figure 11.2: $\log(\mathcal{L}_e/\mathcal{L}_\mu)$ distribution for events passing the basic π^0 cuts. The two peaks straddling zero arise from ν_μ CC events (left) and NC π^0 events (right). The uncertainty band represents the square roots of the diagonal elements of the full covariance matrix. The blue line indicates the location of the cut applied to remove ν_μ CC events.

The desire to remain blind to an oscillation signal motivates two additional cuts. As we will see in Chapter 12, ν_e CC QE (signal-like) events reconstruct with lower mass and with higher $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ values than do π^0 events, by design. To reduce potential signal ν_e events to an acceptably small level, we require:

- $M_{\gamma\gamma} > 50 \text{ MeV}/c^2$
- $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0}) < 0$.

Figure 11.3 shows these cuts graphically along with other cuts that we use later on. Figures 11.4 and 11.5 compare data and simulation.

11.2.1 Momentum dependent correction

The events that pass this suite of cuts are grouped into nine bins of reconstructed π^0 momentum, as shown in Figure 11.6. In each momentum bin, the number of events satisfying

$$80 \text{ MeV}/c^2 < M_{\gamma\gamma} < 200 \text{ MeV}/c^2 \quad (11.1)$$

is measured, and the predicted background contribution is subtracted (from both data and Monte Carlo). Here, “background” is anything that is not a neutral current event with a decayed π^0 .

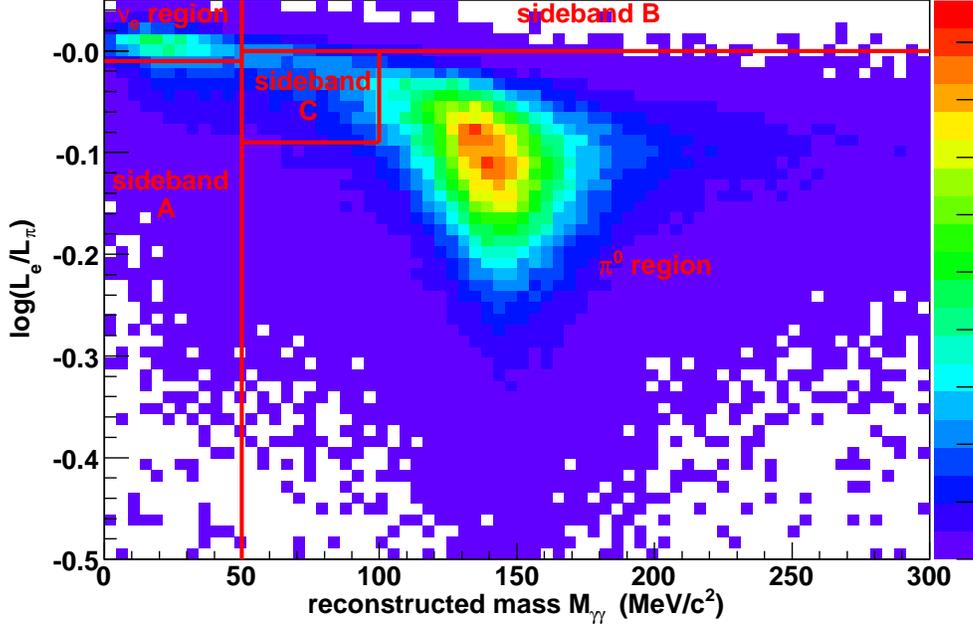


Figure 11.3: Monte Carlo 2D distribution of $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ and $M_{\gamma\gamma}$ for events passing the π^0 pre-selection (all cuts through $\log(\mathcal{L}_e/\mathcal{L}_{\mu}) > 0.05$ in the text). The “ π^0 region” includes the subregion “sideband C”. The “ ν_e region” contains the ν_e -like events which are excluded from the π^0 sample for blindness reasons. The sideband regions are discussed in Chapter 13.

The simulation is used to construct a 9×9 momentum smearing matrix \mathbf{U} , following [136]. The π^0 yields $\mathbf{T} = \{T_i\}$ in bins of true π^0 momentum are calculated with

$$\mathbf{T} = \mathbf{U}^{-1}\mathbf{R} \quad (11.2)$$

where \mathbf{R} is the vector of background-subtracted yields $\{R_i\}$ in bins of reconstructed momentum.⁴ The ratios $T_i^{\text{data}}/T_i^{\text{MC}}$ provide a momentum-dependent weighting that gets applied to all Monte Carlo π^0 events. The previously mentioned Figure 11.6 shows reconstructed mass distributions before and after the correction. The weighting function itself is plotted in Figure 11.7. Systematic errors from the background subtraction procedure and from statistical errors – including effects from the unsmearing process – are indicated in the graph. More details are available in Ref. [137]. Figure 11.8 shows which π^0 ’s are important as oscillation backgrounds.

11.2.2 Coherent π^0 production (angular distribution)

The reconstructed neutrino energy one calculates for misidentified π^0 events depends on the π^0 direction. Uncertainty in the π^0 angular distribution comes in part from the poorly

⁴ \mathbf{M} is built from a Monte Carlo sample large enough to abate noise issues in the unsmearing process [137].

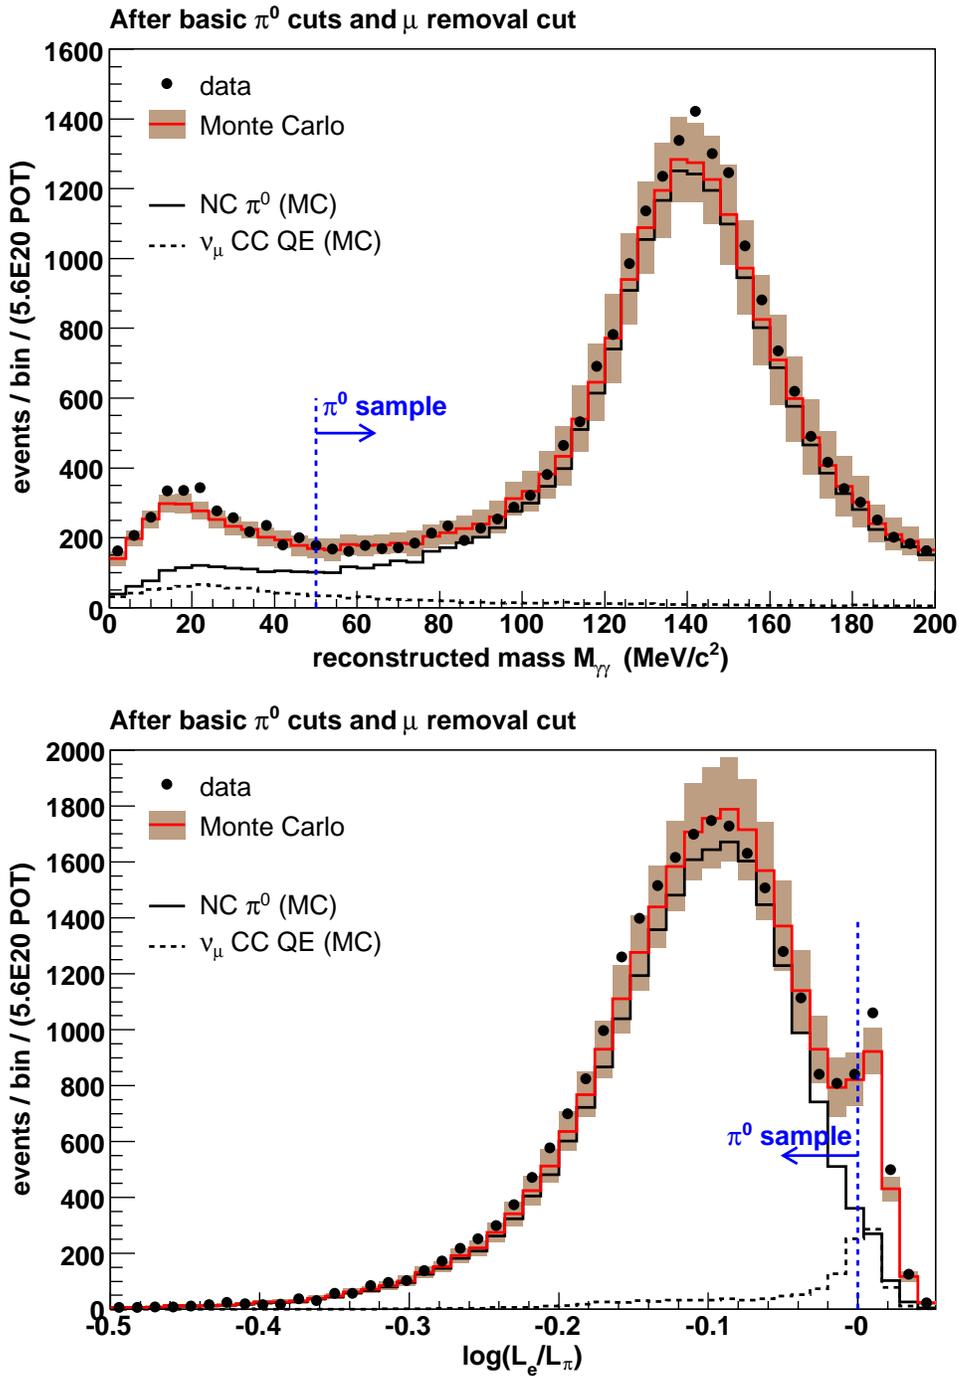


Figure 11.4: (Top) Data [black points] and Monte Carlo [red histogram and band] distributions of $M_{\gamma\gamma}$ for events passing the π^0 pre-selection. The ν_μ CC QE and NC π^0 components are indicated. The blue line gives the location of the blindness cut. (Bottom) The analogous plot for $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$.

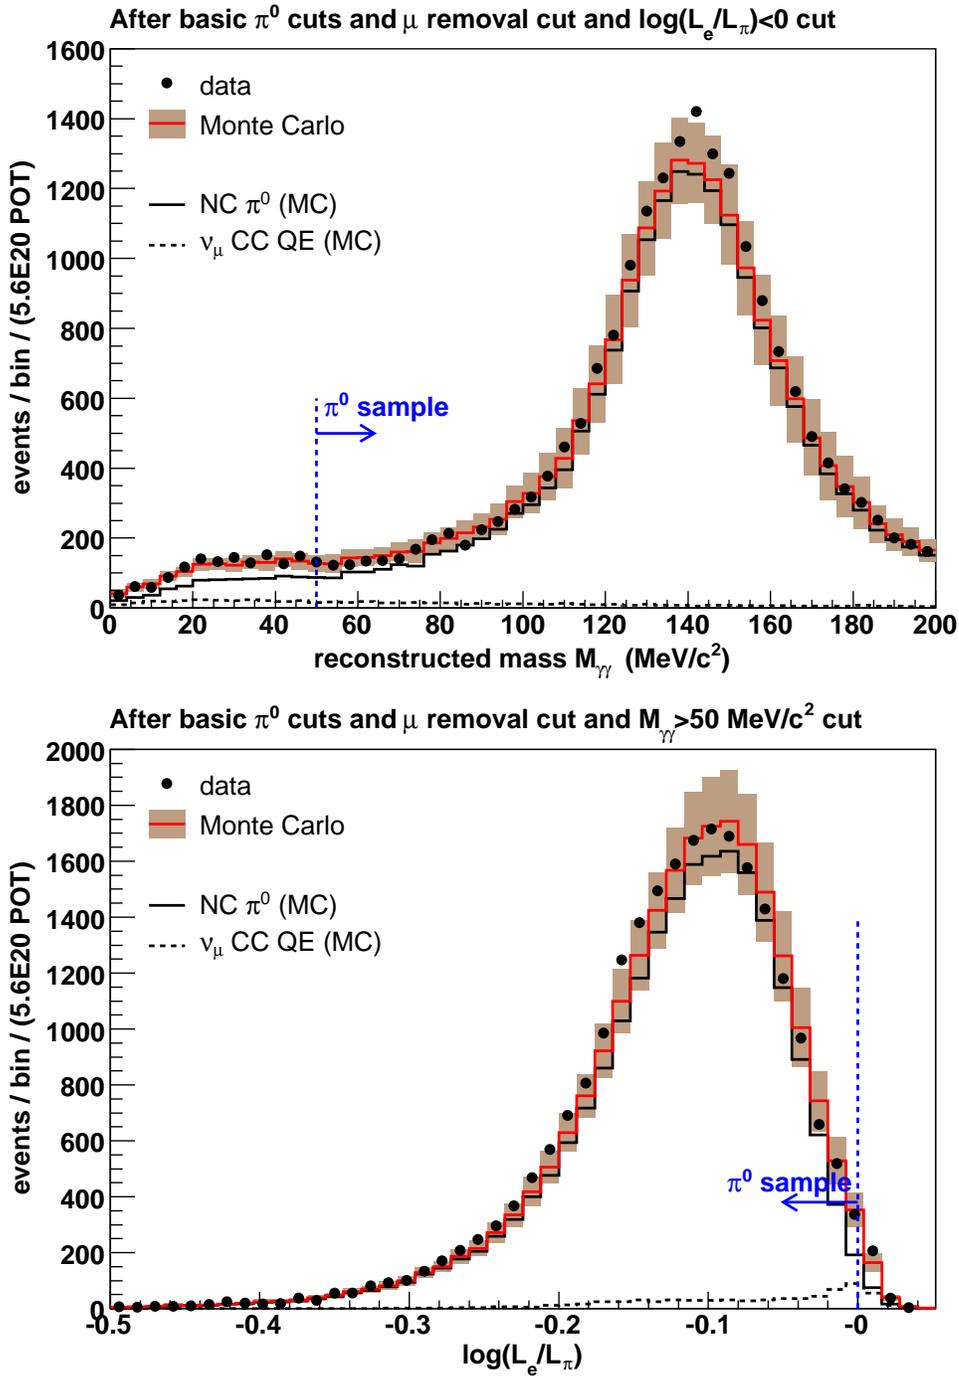


Figure 11.5: This figure differs from Figure 11.4 in that one blindness cut has been applied, either the $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ cut (top) or the $M_{\gamma\gamma}$ cut (bottom).

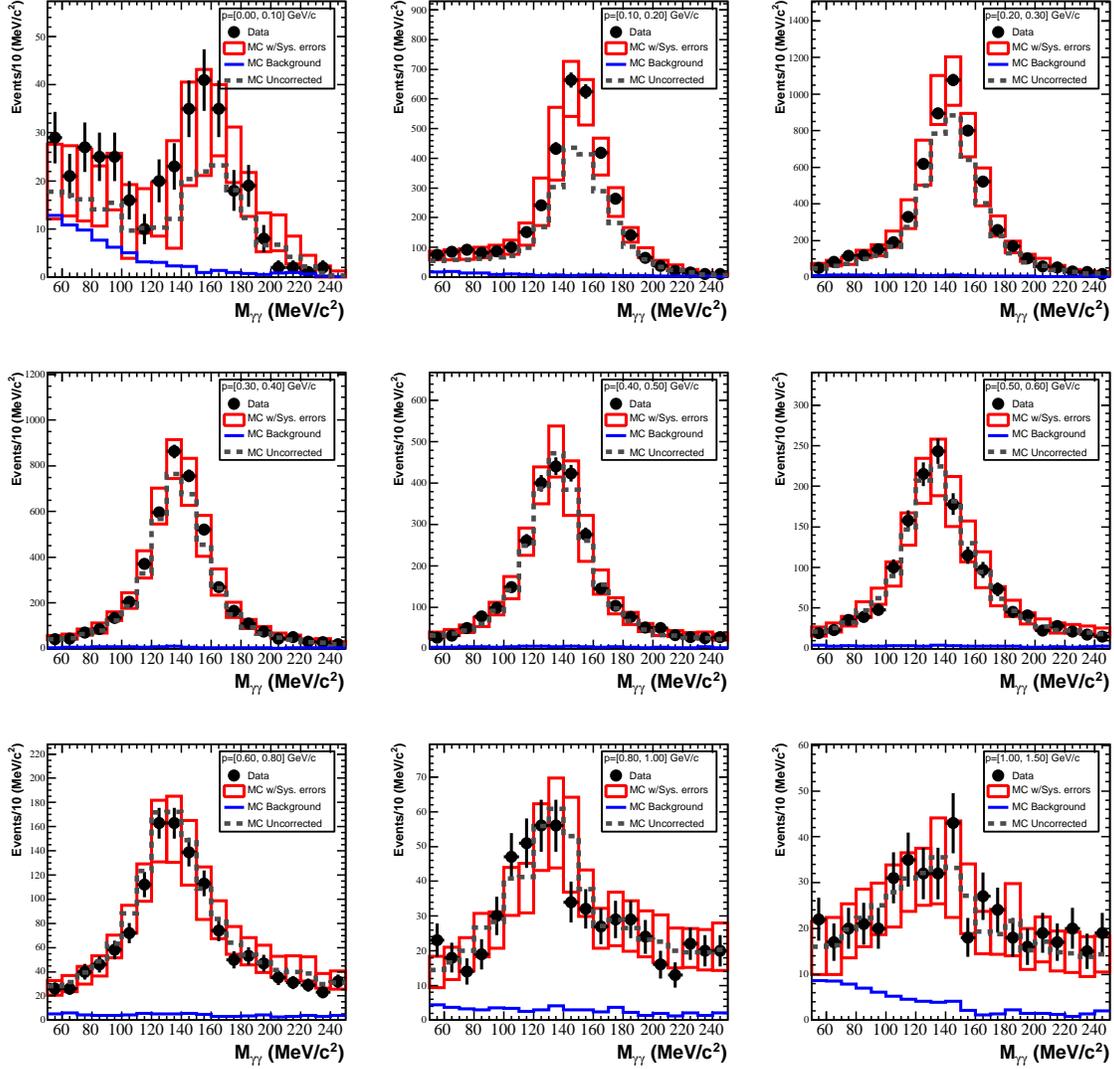


Figure 11.6: Reconstructed $M_{\gamma\gamma}$ distributions for data (black points with statistical error bars), uncorrected Monte Carlo (dashed histogram), and corrected Monte Carlo with systematic errors (red boxes) in bins of reconstructed π^0 momentum, from 0 GeV/c to 1.5 GeV/c. The blue histogram in each panel shows the Monte Carlo estimated backgrounds, which are typically quite small.

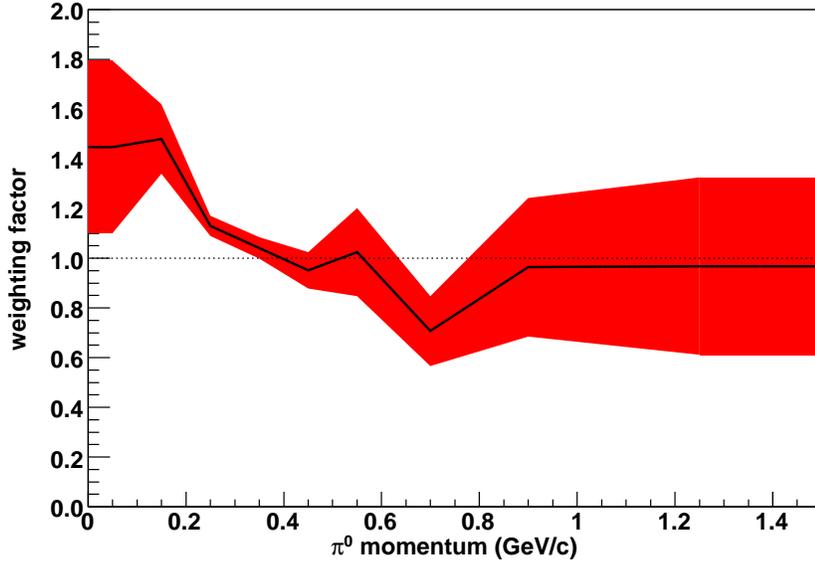


Figure 11.7: The NC π^0 weighting function. The black line passes through the ratios $T_i^{\text{data}}/T_i^{\text{MC}}$, with linear interpolation in between. The red band gives the size of the (correlated) errors.

known cross section for coherent π^0 production:

$$\nu + C \rightarrow \nu + C + \pi^0 . \quad (11.3)$$

(Coherently produced π^0 's have a more forward peaked angular distribution than resonantly produced π^0 's.) A χ^2 fit to the 2D distribution of $M_{\gamma\gamma}$ vs. $E_{\pi^0}(1 - \cos \theta_{\pi^0})$ is carried out to determine the fraction of π^0 's produced coherently, as described in Ref. [137]. Figure 11.9 shows the result.

11.2.3 Radiative Δ decay

While the $\Delta(1232)$ resonance typically produces a pion upon decay ($\Delta \rightarrow N\pi$), it can also de-excite electromagnetically ($\Delta \rightarrow N\gamma$), creating a nearly irreducible oscillation background. Uncertainty in the rate of these radiative decays comes from the branching fraction (listed as $0.0052 - 0.0060$ by the PDG [9]) and from the rate of Δ production. The latter is constrained by our observed π^0 rate, though with an estimated 6% uncertainty due to final state interactions which can eliminate or create a π^0 , and with a 15% uncertainty due to the π^0 rate measurement itself and the imperfectly known resonant π^0 fraction. Resulting correlations between the π^0 and radiative Δ uncertainties are treated in the oscillation analysis.

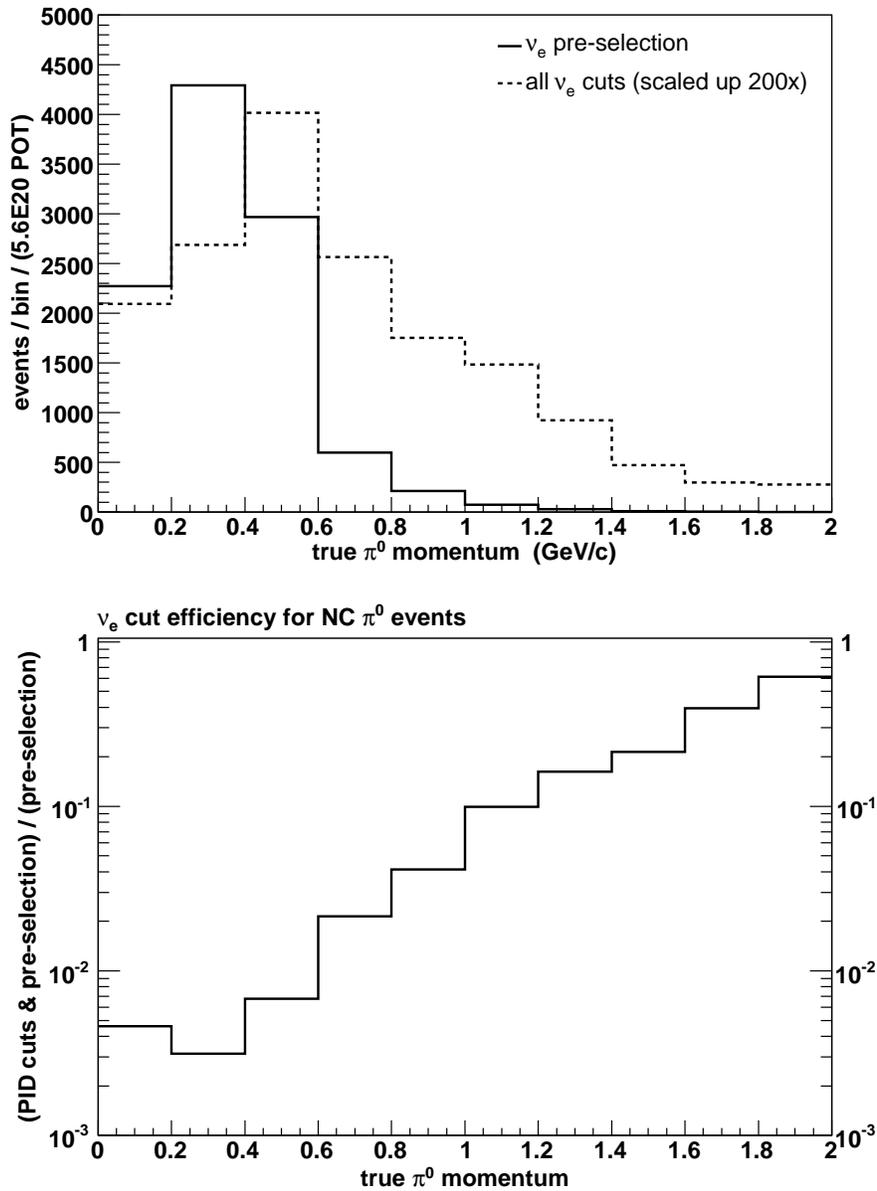


Figure 11.8: (Top) The π^0 momentum distribution after ν_e pre-selection (solid) and full selection (dashed), as defined in the next chapter. The latter histogram is scaled up by a factor of 200 for visibility. (Bottom) The ν_e particle ID cut efficiency for π^0 events. This curve is the ratio of the top two histograms without the $200\times$ scaling.

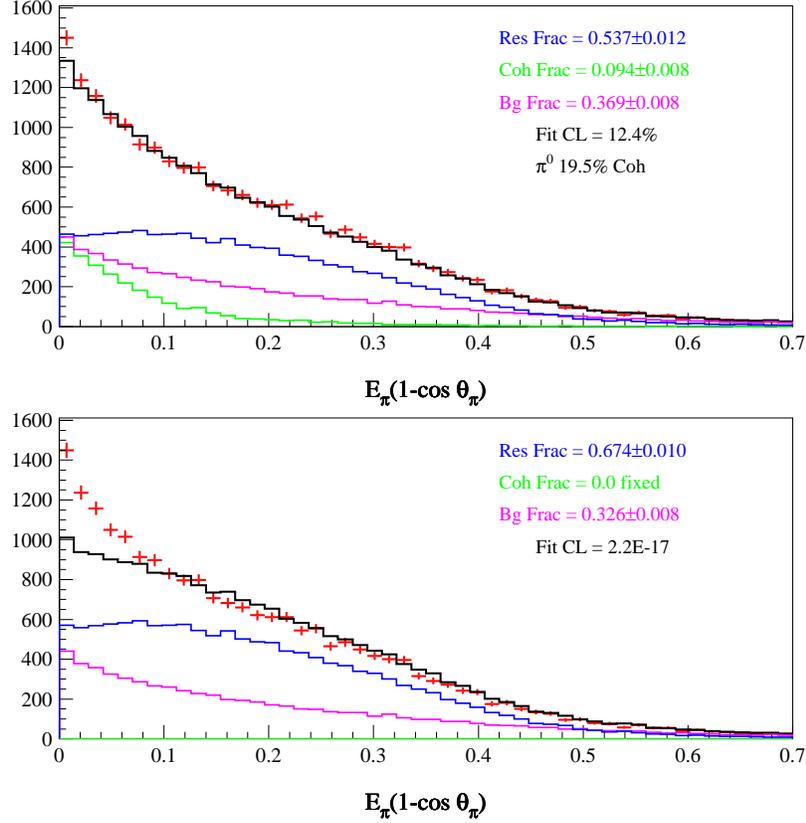


Figure 11.9: Determination of the coherent π^0 fraction. (Top) $E_{\pi^0}(1 - \cos \theta_{\pi^0})$ distribution for the best-fit coherent fraction. (Bottom) The fit result when the coherent fraction is fixed at zero.

11.2.4 Events from outside the tank

The final π^0 -related Monte Carlo correction does not use the π^0 selection at all. Single- γ events from π^0 's produced outside the tank volume (most notably in the dirt) mimic low-energy electrons. Since the radiation length in our oil is only 50 cm or so, these events appear near the tank wall with inward trajectories. We construct a quantity R_{twb} (read: “distance to wall, backward”) which gives the length of active detector that an entering γ would have to traverse (invisibly) to produce the observed track:

$$R_{\text{twb}} = \mathbf{R}_e \cdot \mathbf{U}_e + \left[(\mathbf{R}_e \cdot \mathbf{U}_e)^2 - |\mathbf{R}_e|^2 + R_0^2 \right]^{\frac{1}{2}}, \quad (11.4)$$

where \mathbf{R}_e is the reconstructed e -hypothesis 3-vertex, \mathbf{U}_e is the reconstructed 3-direction, and R_0 is the radius of the “wall”, defined here as the sphere passing through the PMT faces: $R_0 = 550$ cm. Figure 11.10 shows the R_{twb} distribution in data and simulation for low energy ($140 \text{ MeV} < E_e < 280 \text{ MeV}$) and high radius ($480 \text{ cm} < R_e < 510 \text{ cm}$) events.

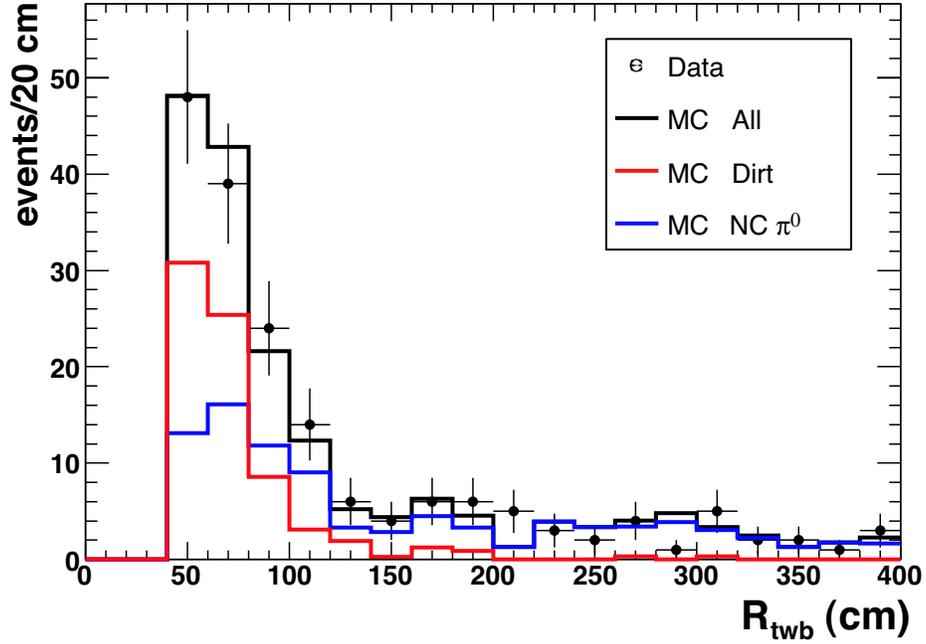


Figure 11.10: The R_{twb} distribution for low-energy, high-radius events. This sample is rich in dirt events (red histogram), and the dirt normalization is consistent with the data. The boundary at $R_{\text{twb}} = 40$ cm is an artifact of the $R_e < 510$ cm in the sample definition.

(These cuts eliminate possible ν_e signal events, thus preserving blindness while keeping plenty of dirt events.) Plots similar to this one are used to establish a scaling factor that gets applied to all Monte Carlo dirt events [138]. The measured factor is 1.00 ± 0.15 .

Chapter 12

ν_e selection and signal fit

12.1 Selecting ν_e CC QE events

We now turn to the ν_e sample itself. The selection begins with these basic cuts:

- Exactly 1 subevent
- $N_{\text{veto}} < 6$
- $N_{\text{main}} > 200$
- $R_e < 500$ cm
- (endpoint R_μ) < 488 cm

The first eliminates events with apparent muon decays; the second keeps cosmic ray muons out and neutrino events in; the third eliminates Michel electrons; and the last two ensure that the entirety of the event is sufficiently far from the tank wall. (See §10.1.) To complete the ν_e pre-selection, we define an E_ν^{rec} analysis range:

- $475 \text{ MeV} < E_\nu^{\text{rec}} < 3000 \text{ MeV}$.

The lower E_ν^{rec} limit was originally 300 MeV¹, but the threshold was increased to 475 MeV for reasons discussed in §13.2.

The pre-selection is followed by three particle identification (PID) cuts which eliminate π^0 and untagged muon backgrounds. The PID cuts are based on the kinematics and maximum likelihoods returned by the event reconstruction algorithm:

- $\log(\mathcal{L}_e/\mathcal{L}_\mu) > a_0 + a_1 E_e + a_2 E_e^2$
- $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0}) > b_0 + b_1 E_e + b_2 E_e^2$
- $M_{\gamma\gamma} < c_0 + c_1 E_e + c_2 E_e^2$.

¹below which hadronic propagation uncertainties become difficult to estimate adequately

The cuts depend quadratically on E_e (the reconstructed lepton energy under an electron hypothesis). This leaves nine cut coefficients whose values are set such that the estimated oscillation sensitivity near $\Delta m^2 = 1.0 \text{ eV}^2$ and $\sin^2 2\theta = 0.004$ is maximized.² Table 12.1 gives the values for the cut coefficients, and Figure 12.1 shows the cuts with representative Monte Carlo event samples. The figure also includes a graph of the cut efficiencies.

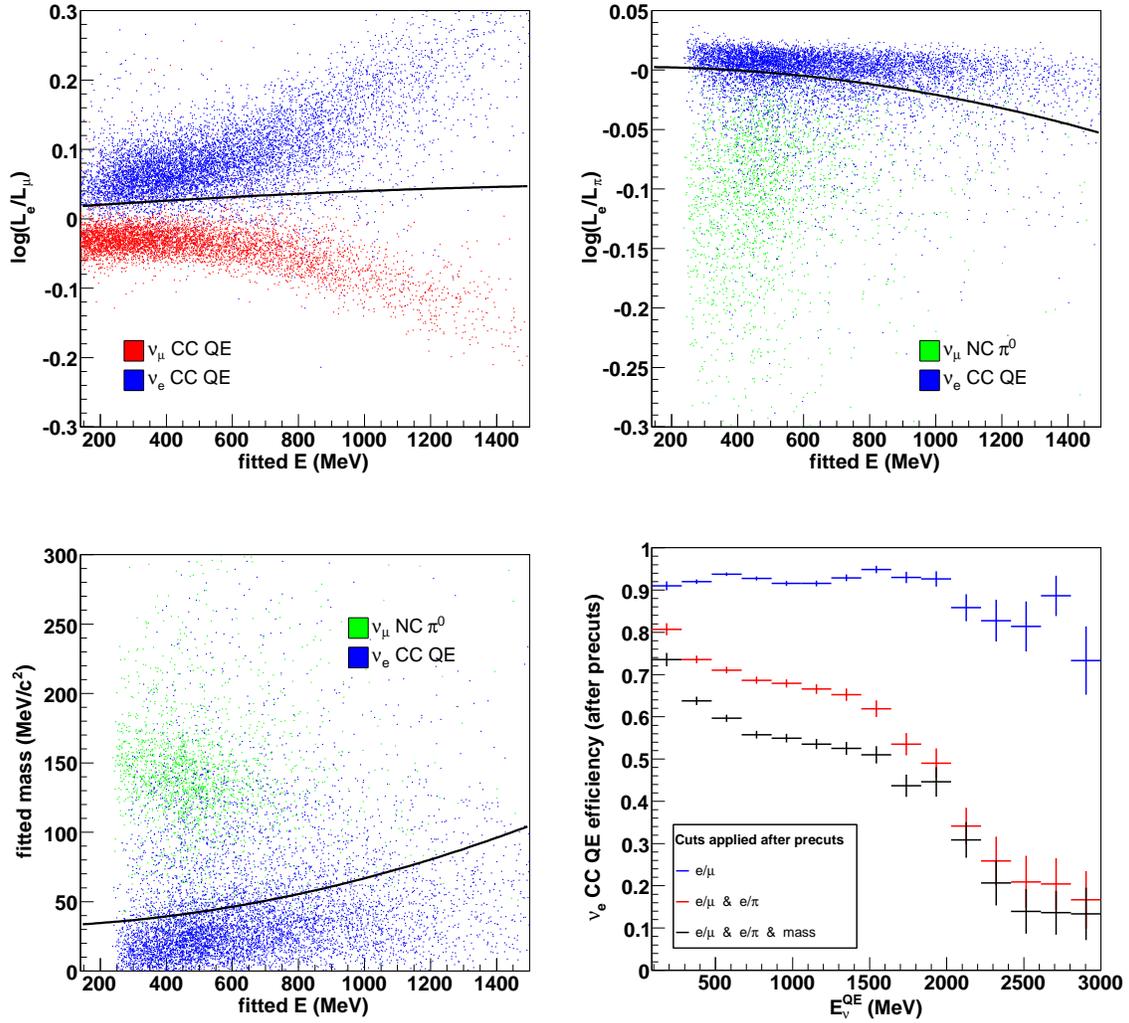


Figure 12.1: The PID cuts versus E_e . The scatter plots in each panel show the signal and most relevant background component: (blue) ν_e CC QE [signal], (red) ν_μ CC QE, (green) NC π^0 . The bottom right panel shows the PID selection efficiency for ν_e CC QE events as a function of E_ν^{rec} . Each subsequently lower curve includes cuts from higher curves.

²Oscillation sensitivity, formally defined later (§12.5), quantifies how strongly we expect to exclude a given $(\sin^2 2\theta, \Delta m^2)$ point if no signal is present. Sensitivity throughout $(\sin^2 2\theta, \Delta m^2)$ space was monitored during cut optimization, but the test point mentioned here was used whenever a single figure of merit was needed.

$\log(\mathcal{L}_e/\mathcal{L}_\mu)$		$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$		$M_{\gamma\gamma}$	
a_0 :	1.355×10^{-2}	b_0 :	2.471×10^{-3}	c_0 :	$3.203 \times 10^{-2} \text{ GeV } c^{-2}$
a_1 :	$3.467 \times 10^{-2} \text{ GeV}^{-1}$	b_1 :	$4.115 \times 10^{-3} \text{ GeV}^{-1}$	c_1 :	$7.417 \times 10^{-3} c^{-2}$
a_2 :	$-8.259 \times 10^{-3} \text{ GeV}^{-2}$	b_2 :	$-2.738 \times 10^{-2} \text{ GeV}^{-2}$	c_2 :	$2.738 \times 10^{-2} \text{ GeV}^{-1} c^{-2}$

Table 12.1: Optimized PID cut coefficients.

The next several plots, unseen prior to our lifting blindness, step through the PID variables. In each, data and expected background distributions are shown along with an example signal distribution (best-fit LSND parameters). The Monte Carlo error band gives the square root of the diagonal elements of the total covariance matrix. Since the PID cuts are energy dependent, there is no single cut value to indicate on each plot. Instead, a range is shown by marking the $E_e = 200 \text{ MeV}$ and $E_e = 1200 \text{ MeV}$ cut values. First, Figure 12.2 shows the $\log(\mathcal{L}_e/\mathcal{L}_\mu)$ distribution for all events passing the ν_e pre-selection in data and Monte Carlo. The example oscillation signal is inflated for visibility in this plot only. Figure 12.3 shows the $M_{\gamma\gamma}$ and $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ distributions after the $\log(\mathcal{L}_e/\mathcal{L}_\mu)$ has been applied. Finally, Figure 12.4 shows the same two distributions after the application of each others' cuts. In all cases, the simulation agrees quite well with data, a statement we quantify later.

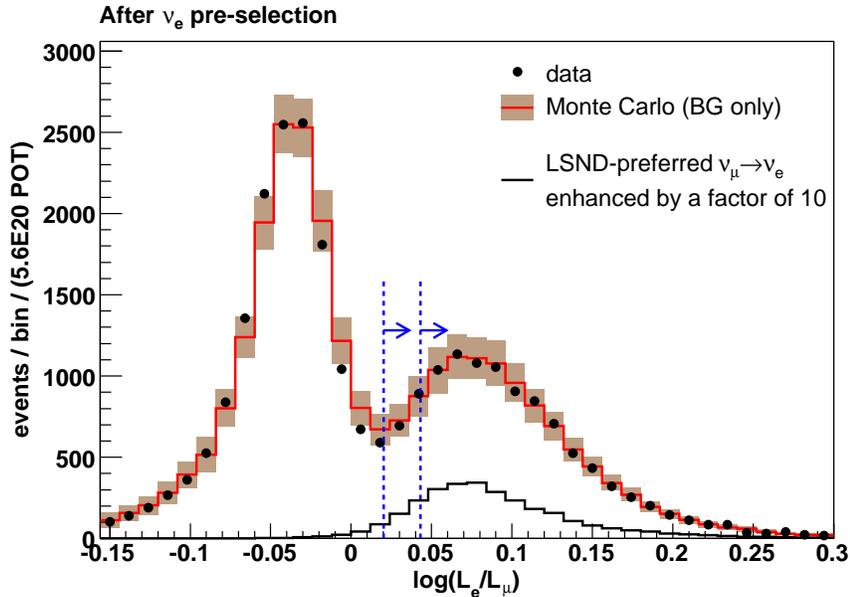


Figure 12.2: $\log(\mathcal{L}_e/\mathcal{L}_\mu)$ distribution for events passing the ν_e CC QE pre-selection. Data (black points) and Monte Carlo simulation (red histogram with error band) are shown. An LSND-like oscillation histogram is also shown, scaled up by a factor of 10 for visibility. $\log(\mathcal{L}_e/\mathcal{L}_\mu)$ cut values for $E_e = 200 \text{ MeV}$ and $E_e = 1200 \text{ MeV}$ are indicated by the blue lines.

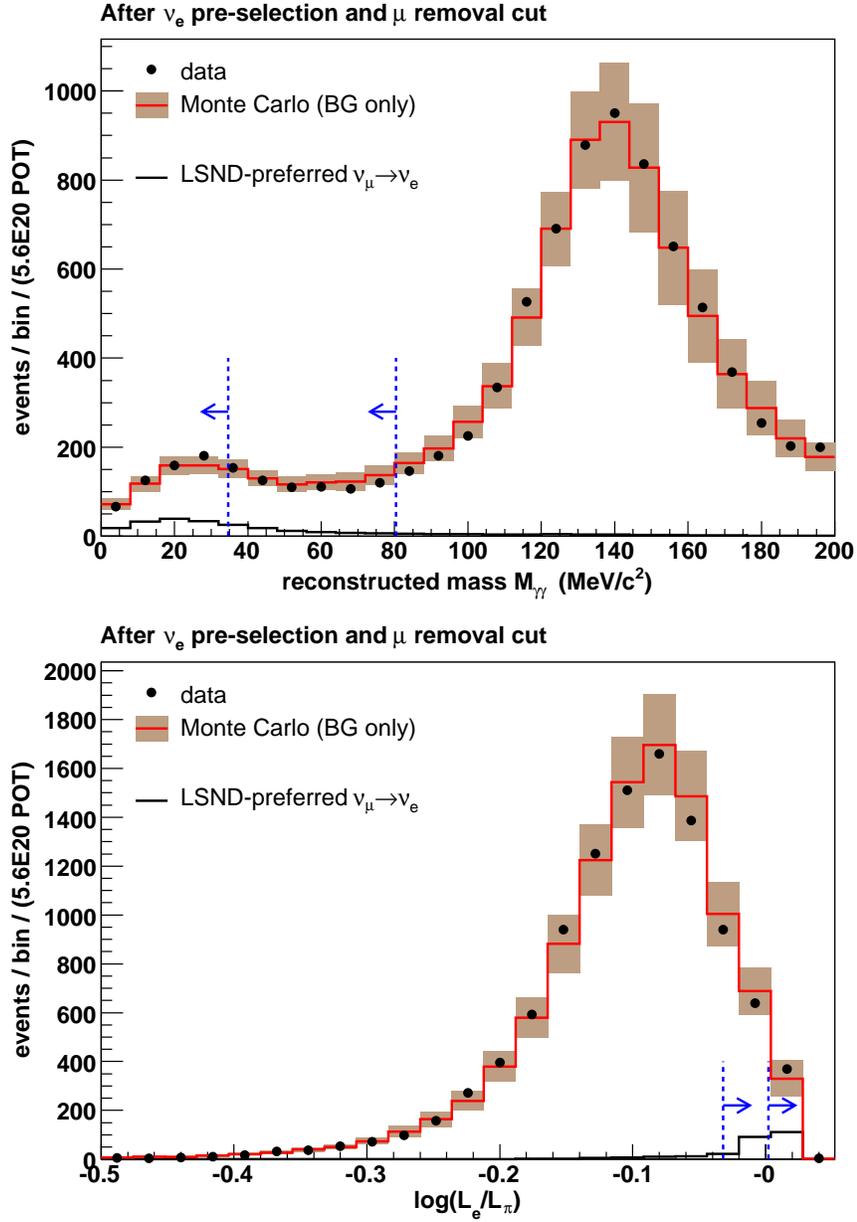


Figure 12.3: $M_{\gamma\gamma}$ and $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ distributions after the ν_e pre-selection and the muon removal cut. The example oscillation signal has not been enhanced. Otherwise, the plot conventions are as in Figure 12.2.

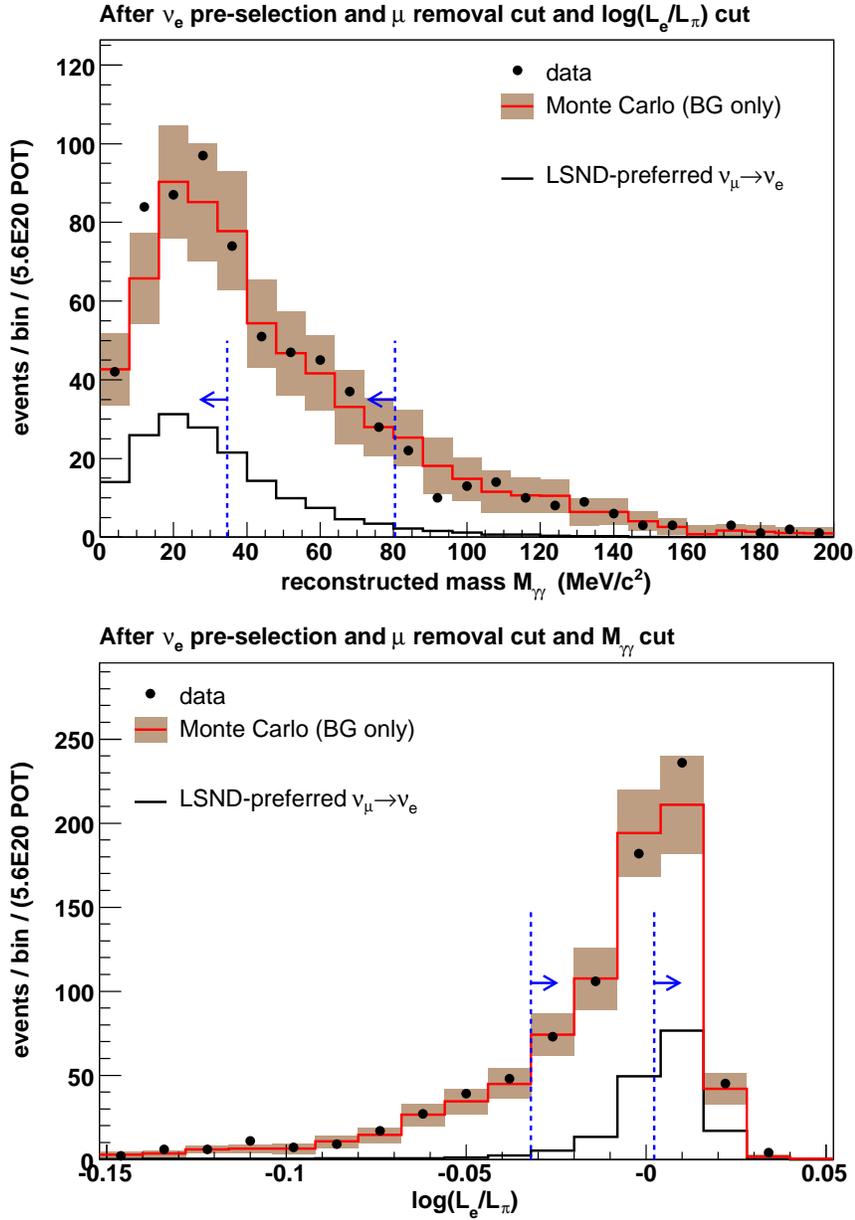


Figure 12.4: $M_{\gamma\gamma}$ and $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ distributions after the ν_e pre-selection, the muon removal cut, and the opposing PID cut. (That is, the $M_{\gamma\gamma}$ cut is applied on the bottom plot; $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ on the top.) The example oscillation signal has not been enhanced. Otherwise, the plot conventions are as in Figure 12.2.

Figure 12.5 shows the backgrounds expected after the above cuts along with an example signal. Without the constraints of the preceding chapters, a $\nu_\mu \rightarrow \nu_e$ signal of the size shown would be swamped by 30 – 50% flux and cross section uncertainties on the backgrounds beneath it. Table 12.2 gives a more detailed breakdown of the sample’s components.

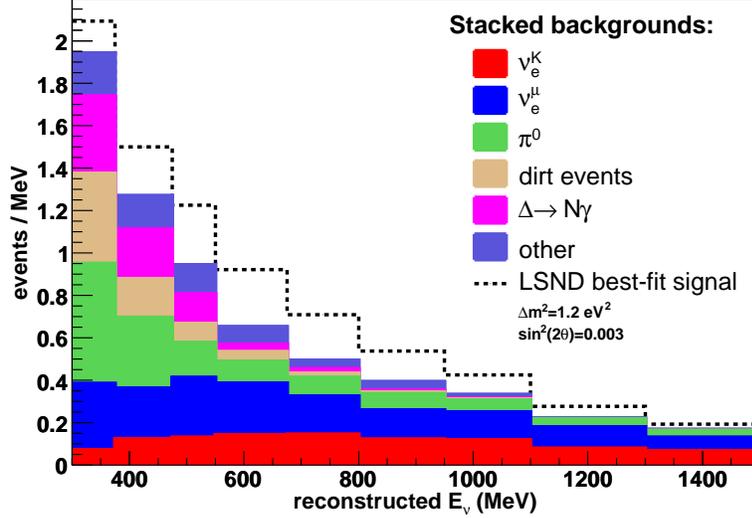


Figure 12.5: A histogram of E_ν^{rec} showing the backgrounds (stacked upon one another) expected after all ν_e selection cuts, except for the cut on E_ν^{rec} itself. The neutral current misidentification backgrounds pile up at low E_ν^{rec} since the outgoing neutrino removes energy from the event.

ν_μ CC QE	10 ± 2
$\nu_\mu e \rightarrow \nu_\mu e$	7 ± 2
other ν_μ events	13 ± 5
NC π^0	62 ± 10
NC $\Delta \rightarrow N\gamma$	20 ± 4
dirt events	17 ± 3
$\mu \rightarrow \nu_e$	132 ± 10
$K^+ \rightarrow \nu_e$	71 ± 26
$K_L^0 \rightarrow \nu_e$	23 ± 7
$\pi^+ \rightarrow \nu_e$	3 ± 1
total background	358 ± 35
0.26% $\nu_\mu \rightarrow \nu_e$	163 ± 21
LSND best-fit $\nu_\mu \rightarrow \nu_e$	126 ± 16

Table 12.2: Expected numbers of background events in the reconstructed neutrino energy range $475 \text{ MeV} < E_\nu^{\text{rec}} < 1250 \text{ MeV}$. (This range is used for the “counting” results in Chapter 14.) The systematic errors listed are significantly correlated with one another. Example numbers of $\nu_\mu \rightarrow \nu_e$ oscillation events are shown either assuming $P_{\text{osc}} = 0.26\%$ with no spectral distortion or putting $\sin^2 2\theta = 0.003$, $\Delta m^2 = 1.2 \text{ eV}^2$ (LSND).

12.2 Signal χ^2

The events that pass the ν_e cuts are used to fill an E_ν^{rec} histogram like the one above. The histogram from the simulation (*i.e.*, the prediction histogram) can have some amount of $\nu_\mu \rightarrow \nu_e$ signal included. We form a standard χ^2 expression:

$$\chi^2(\boldsymbol{\alpha}) = \sum_i \sum_j (d_i - p_i(\boldsymbol{\alpha})) V_{ij}^{-1} (d_j - p_j(\boldsymbol{\alpha})) , \quad (12.1)$$

where d_i is the height of i -th bin of the data histogram, p_i is the height of the i -th bin of the prediction histogram, $\boldsymbol{\alpha} \equiv (\sin^2 2\theta, \Delta m^2)$ controls the amount of oscillation signal present in the $\{p_i\}$, and V_{ij}^{-1} is the (i, j) -th element of the inverse of the covariance matrix \mathbf{V} that describes the uncertainties in the $\{p_i\}$. We use $\chi^2(\boldsymbol{\alpha})$ to extract best-fit parameters and to set confidence intervals in the $(\sin^2 2\theta, \Delta m^2)$ plane.

12.3 Forming the covariance matrix \mathbf{V}

Uncertainties in the prediction come from several unrelated sources. For example, our knowledge of K^+ production in the target comes from fits to available K^+ production data while our quite separate knowledge of optical photon Rayleigh scattering in the detector comes from laboratory measurements with mineral oil. Uncertainties in these are completely decoupled. However, lack of knowledge about Rayleigh scattering affects (for example) our ability to measure scintillation production using Michel electrons. Thus, our scintillation and scattering uncertainties are correlated.

The total covariance matrix \mathbf{V} is formed by summing eleven individual covariance matrices, each one having no correlations with the others. The total matrix is

$$\mathbf{V} = \mathbf{V}^{\pi^+} + \mathbf{V}^{K^+} + \mathbf{V}^{K^0} + \mathbf{V}^{\pi^-} + \mathbf{V}^{\text{beam}} + \mathbf{V}^\sigma + \mathbf{V}^{\pi^0} + \mathbf{V}^{\text{dirt}} + \mathbf{V}^{\text{det}} + \mathbf{V}^{\text{qt}} + \mathbf{V}^{\text{MC}} .$$

The component matrices are summarized in Table 12.3.

12.3.1 The first eight component matrices

Consider the matrix \mathbf{V}^{π^+} , which represents the uncertainty in the N bin heights $\mathbf{p} = \{p_i\}$ due to uncertainties in the π^+ production cross section. One can consider \mathbf{p} to be the signal E_ν^{rec} histogram, but it may be any histogram (or, more generally still, any N quantities) whose prediction depends on our knowledge of π^+ production. As described in §6.1.2, the π^+ production cross section is parametrized in the beam simulation with a seven-parameter Sanford-Wang function

$$f(p, \theta; \mathbf{s}) \equiv \frac{d^2 \sigma}{dp d\Omega} , \quad (12.2)$$

matrix	uncertainty due to...
\mathbf{V}^{π^+}	π^+ production (Chapter 6)
\mathbf{V}^{K^+}	K^+ production (Chapter 6)
\mathbf{V}^{K^0}	K^0 production (Chapter 6)
\mathbf{V}^{π^-}	π^- production (Chapter 6)
\mathbf{V}^{beam}	beamline modeling and hadron cross sections (Chapter 6)
\mathbf{V}^σ	neutrino cross sections (Chapter 7)
\mathbf{V}^{π^0}	neutral current π^0 rate measurement (Chapter 11)
\mathbf{V}^{dirt}	neutrino events in the dirt (Chapter 11)
\mathbf{V}^{det}	detector modeling <i>and</i> statistical errors (Chapter 8)
\mathbf{V}^{qt}	PMT charge-time correlation (Chapter 8)
\mathbf{V}^{MC}	finite Monte Carlo statistics

Table 12.3: The component covariance matrices. Note that statistical errors are included in \mathbf{V}^{det} for reasons described in §12.3.2.

where \mathbf{s} is the vector of seven imperfectly known Sanford-Wang parameters. Let $\hat{\mathbf{s}}$ be the parameters used in the default simulation, and let \mathbf{S} be the error matrix representing our ignorance of the parameter values.

What we don't do

The full Monte Carlo simulation and analysis chain, which includes the Sanford-Wang function $f(p, \theta; \mathbf{s})$ and which produces our prediction histogram \mathbf{p} , can be thought of as a set of functions $\{p_i(\mathbf{s})\}$ for the histogram bin heights $\{p_i\}$. Commonly, one calculates the desired matrix elements $V_{ij}^{\pi^+}$ with

$$V_{ij}^{\pi^+} \approx \sum_m \sum_n \left. \frac{dp_i}{ds_m} \right|_{\hat{\mathbf{s}}} \left. \frac{dp_j}{ds_n} \right|_{\hat{\mathbf{s}}} S_{mn} , \quad (12.3)$$

where the derivatives $\left\{ \frac{dp_i}{ds_m} \right\}$ are estimated using the simulation. The approximation in Eq. (12.3) breaks down if the Taylor expansion of $p_i(\mathbf{s})$ about $\hat{\mathbf{s}}$ contains nonlinear terms that are significant in a region about $\hat{\mathbf{s}}$ of size comparable to the standard deviations of \mathbf{s} .

What we do do

To avoid this and other issues, we instead propagate the parameter errors \mathbf{S} through to \mathbf{V}^{π^+} using a Monte Carlo technique.³ Recall that the covariance matrix \mathbf{S} corresponds to a p.d.f. $g(\mathbf{s})$ for the parameters \mathbf{s} , where

$$g(\mathbf{s}) = \frac{1}{((2\pi)^M |\mathbf{S}|)^{\frac{1}{2}}} \exp \left(-\frac{1}{2} (\mathbf{s} - \hat{\mathbf{s}})^\top \mathbf{S}^{-1} (\mathbf{s} - \hat{\mathbf{s}}) \right) \quad (12.4)$$

³The fact that the function $\mathbf{p}(\mathbf{s})$ happens also to involve Monte Carlo is immaterial to the error propagation approach described here.

and where M is the number of parameters. ($M=7$ in this example.) We draw K random parameter sets $\{\mathbf{s}^{(1)}, \mathbf{s}^{(2)}, \dots, \mathbf{s}^{(K)}\}$ according to this p.d.f., and we calculate the bin heights $\{\mathbf{p}^{(1)}, \mathbf{p}^{(2)}, \dots, \mathbf{p}^{(K)}\}$ for each parameter set using our function $\mathbf{p}(\mathbf{s})$. We can estimate the covariances of \mathbf{p} due to \mathbf{S} from the resulting ensemble:

$$V_{ij}^{\pi^+} \approx \frac{1}{K-1} \sum_{k=1}^K \left(p_i^{(k)} - \langle p_i \rangle \right) \left(p_j^{(k)} - \langle p_j \rangle \right) , \quad (12.5)$$

where $\langle p_i \rangle = \frac{1}{K} \sum_k p_i^{(k)}$ and where $K=100$ for this analysis.

This is an attractive technique in part because we can weight events from a single high-statistics Monte Carlo production at $\hat{\mathbf{s}}$ to evaluate the many $\{\mathbf{p}^{(k)}\}$ quickly. To obtain the k -th π^+ Sanford-Wang variation, we weight each event that has a (p -Be)-produced π^+ in its history by

$$w = \frac{f(p, \theta; \mathbf{s}^{(k)})}{f(p, \theta; \hat{\mathbf{s}})} , \quad (12.6)$$

where p and θ are the initial momentum and angle of the event's parent π^+ . The variation so induced in the downstream histogram $\mathbf{p}^{(k)}$ relative to the unweighted version \mathbf{p} (or relative to the average $\langle \mathbf{p} \rangle$) reflects how much this histogram is affected by the change in π^+ cross section parameters.

The first eight component matrices in Table 12.3 are handled this way. Each matrix has its own function analogous to $f(p, \theta; \mathbf{s})$ whose ratios provide the weights of Eq. (12.6).

12.3.2 \mathbf{V}^{det}

The parameters in the detector simulation affect the nature of events rather than the likelihood of their occurring. Thus, \mathbf{V}^{det} cannot be calculated through the weighting of events in a single large Monte Carlo run. Fresh events must be generated for each random parameter set $\mathbf{s}^{(k)}$. With the weighting method, finite Monte Carlo statistics play a small role due to the cancellation of fluctuations.⁴ For \mathbf{V}^{det} , though, each parameter set gets a new statistically independent event sample, and one must run an impractically large number of Monte Carlo events to have sufficiently low noise in the differences of Eq. (12.5).

To get around this, each of our $K=64$ detector model variants is generated with a number of events equal to that expected in data. The k -th parameter set, then, yields a histogram $\mathbf{p}^{(k)}$ that differs from the default histogram not only because the detector model parameters have changed ($\hat{\mathbf{s}} \mapsto \mathbf{s}^{(k)}$) but also because the histogram has statistical fluctuations. However, these fluctuations are the same size as those expected in data. If we calculate covariances from the K -fold ensemble using

$$V_{ij}^{\text{det}} \approx \frac{1}{K} \sum_{k=1}^K (p_i^{(k)} - p_i(\hat{\mathbf{s}}))(p_j^{(k)} - p_j(\hat{\mathbf{s}})) , \quad (12.7)$$

⁴The Monte Carlo predictions enter Eq. (12.5) only in differences: $(p_i^{(k)} - \langle p_i \rangle)$. Statistical fluctuations in p_i tend to cancel in these differences since the same underlying events provide both terms.

we get a matrix \mathbf{V}^{det} that represents the combination of the detector systematic error matrix and the statistical error matrix. The statistical contribution should be equal to $V_{ij}^{\text{stat}} = \hat{p}_i \delta_{ij}$, a fact we can use for diagnostics on \mathbf{V}^{det} . Since \mathbf{V}^{det} comes with the statistical error built in, we must not add a separate \mathbf{V}^{stat} to the total matrix \mathbf{V} .

Note that the covariances calculated in Eq. (12.7) are taken about the histogram $\mathbf{p}(\hat{\mathbf{s}})$ rather than about the ensemble average as it is in Eq. (12.5). As described in §8.7, the p.d.f. for the detector model parameters was created without reference to the default parameter set $\hat{\mathbf{s}}$. Thus, the ensemble $\{\mathbf{p}^{(k)}\}$ could, in principle, be distributed about a point somewhat removed from the prediction $\mathbf{p}(\hat{\mathbf{s}})$ which actually appears in the χ^2 of Eq. (12.1). The form of Eq. (12.7) ensures that any such displacement enlarges the estimation of the error correspondingly. (In practice, no large displacement is seen, mostly because $\mathbf{V}^{\text{det}} \approx \mathbf{V}^{\text{stat}}$. That is, detector systematics have a small effect on the analysis.)

12.3.3 \mathbf{V}^{qt}

The correlation between a PMT's charge and its time is modeled (§8.4), but its implementation does not lend itself to the above error propagation methods. We thus form a separate matrix \mathbf{V}^{qt} to handle our uncertainty in this phenomenon. We begin with a conservative variant of the charge-time correlation model to produce a prediction histogram \mathbf{p}^{qt} . We could calculate differences between this and the default prediction $\hat{\mathbf{p}}$ to estimate the covariance V_{ij}^{qt} :

$$V_{ij}^{\text{qt}} \approx \left(p_i^{\text{qt}} - \hat{p}_i \right) \left(p_j^{\text{qt}} - \hat{p}_j \right) . \quad (12.8)$$

However, as with \mathbf{V}^{det} , CPU time makes it difficult to overcome the effect of Monte Carlo statistical fluctuations. So, we rewrite the above as

$$\begin{aligned} V_{ij}^{\text{qt}} &= \left(\frac{p_i^{\text{qt}}}{\hat{p}_i} \hat{p}_i - \hat{p}_i \right) \left(\frac{p_j^{\text{qt}}}{\hat{p}_j} \hat{p}_j - \hat{p}_j \right) \\ &= [\eta(E_i) - 1] [\eta(E_j) - 1] \hat{p}_i \hat{p}_j , \end{aligned} \quad (12.9)$$

where we have introduced $\eta(E_i) \equiv \frac{p_i^{\text{qt}}}{\hat{p}_i}$ whose notation reminds us that our observables $\{p_i\}$ are, for the signal extraction, consecutive bins of an energy histogram. If we could run an infinite amount of Monte Carlo, $\eta(E)$ would be a smooth and slowly varying function of E differing from $\eta=1$ as the model variant dictated. Since we have finite Monte Carlo statistics, $\eta(E)$ actually exhibits bin-to-bin fluctuations which we can remove by acknowledging the underlying smoothness. A polynomial, whose degree is sufficiently large to encapsulate the systematic trends in $\eta(E)$ while still being much smaller than the number of elements (bins) of \mathbf{p} , is fitted to $\eta(E)$ to form $\eta'(E)$, giving us our final expression

$$V_{ij}^{\text{qt}} = [\eta'(E_i) - 1] [\eta'(E_j) - 1] \hat{p}_i \hat{p}_j . \quad (12.10)$$

Eq. (12.10) enjoys the same reduction in sensitivity to jitter as does the weighting case above. \mathbf{V}^{qt} is the only matrix that uses this smoothing technique.

12.3.4 V^{MC}

This diagonal matrix provides the statistical uncertainty due to the finite number of events used to form the default prediction $\hat{\mathbf{p}}$. For non-signal events originating within the detector, the Monte Carlo sample is $7.9\times$ larger than the data sample. For events from the dirt: $4.1\times$. For signal events: $50\times$ (if $P_{\text{osc}}\approx 0.26\%$).

12.4 Signal extraction

12.4.1 Best-fit oscillation parameters

The oscillation parameters that best describe the data can be obtained by minimizing the χ^2 in Eq. (12.1). We use the MIGRAD method of Minuit [133] to find the minimum, but we must defend against settling into a non-global minimum.⁵ This possibility stems from the oscillatory nature of $P(\nu_\mu \rightarrow \nu_e)$; recalling Eq. (1.25) :

$$P(\nu_\mu \rightarrow \nu_e) = \sin^2 2\theta \sin^2[\kappa \Delta m^2 (L/E)] .$$

The period of oscillation of the χ^2 surface along Δm^2 can be estimated from the argument of the second sinusoid:

$$\begin{aligned} T_{\Delta m^2} &\approx \frac{\pi}{1.27} \left\langle \frac{E}{L} \right\rangle \frac{\text{km eV}^2}{\text{GeV } c^4} \\ &\approx 3 \text{ eV}^2/c^4 , \end{aligned} \tag{12.11}$$

where we've used $\langle \frac{E}{L} \rangle \approx 1.2 \text{ GeV/km}$. This $\sim 3 \text{ eV}^2/c^4$ periodic structure can be seen in the LSND allowed regions of Figure 1.8.

To find the global minimum, we seed Minuit with parameters from each χ^2 valley in turn (up to the limit of our Δm^2 resolving power), reporting in the end the best minimum found from the many Minuit calls. A more precise description:

1. Begin with $\Delta m^2 = 0.1 \text{ eV}^2/c^4$ (which is within the first period).
2. Find the $\sin^2 2\theta$ value that minimizes χ^2 at this fixed Δm^2 .
3. Seed Minuit with the resulting pair of parameters and run MIGRAD to find a minimum. Both parameters are allowed to vary within the physical plane.⁶
4. Increase Δm^2 by $1 \text{ eV}^2/c^4$ and repeat (2) and (3), recording the new χ^2 minimum and its parameters.
5. Repeat (4) until $\Delta m^2 > 30 \text{ eV}^2/c^4$.
6. Return the best parameters (lowest χ^2) recorded throughout the process.

⁵MIGRAD is efficient at finding the nearest local minimum, but it has little ability to recognize the presence of multiple minima.

⁶ $\Delta m^2 \geq 0$ and $0 \leq \sin^2 2\theta \leq 1$

This procedure gives MIGRAD several shots at each valley in the χ^2 surface. The $1 \text{ eV}^2/c^4$ steps are smaller than the $\sim 3 \text{ eV}^2/c^4$ periodicity so that one need not worry about the phase of the valley structure.

The covariance matrix \mathbf{V} is held fixed during the above six steps. Once the best parameters are found, the matrix is updated to account for the influence of any signal, and the full procedure is iterated. (The first pass uses a no-signal matrix.) These matrix iterations end when the minimum χ^2 changes by less than 0.1 from one iteration to the next. This condition is usually satisfied going from the second to the third fit. That is, the second fit is usually sufficient, but a third is required to confirm convergence.

12.4.2 Global scan confidence intervals

Confidence intervals in the $(\sin^2 2\theta, \Delta m^2)$ plane can be established through the $\Delta\chi^2$ surface

$$\Delta\chi^2(\boldsymbol{\alpha}) = \chi^2(\boldsymbol{\alpha}) - \chi_{\text{best}}^2, \quad (12.12)$$

where $\chi_{\text{best}}^2 = \chi^2(\boldsymbol{\alpha}_{\text{best}})$ is the minimum just found and where $\chi^2(\boldsymbol{\alpha})$ is evaluated using Eq. (12.1).⁷ If this surface is parabolic and the parameters are unbounded (a common scenario), a contour of constant $\Delta\chi^2(\boldsymbol{\alpha}) = \Delta\chi_c^2$ defines a confidence interval whose confidence level (C.L.) c depends on the constant $\Delta\chi_c^2$ and on the number of free parameters in the problem. Tables of $\Delta\chi_c^2$ are available in the literature [9]. For two free parameters, the $c=90\%$ C.L. contour corresponds to $\Delta\chi_c^2=4.61$.

However, we have a non-parabolic $\Delta\chi^2(\boldsymbol{\alpha})$ surface and bounded parameters. In general, this means that $\Delta\chi_c^2$ no longer has a fixed value for a given c . Rather, the contour c is formed by the locus of points satisfying

$$\Delta\chi^2(\boldsymbol{\alpha}) = \Delta\chi_c^2(\boldsymbol{\alpha}), \quad (12.13)$$

where the constant above has been replaced by the surface $\Delta\chi_c^2(\boldsymbol{\alpha})$. Often, $\Delta\chi_c^2(\boldsymbol{\alpha})$ is near the standard (constant) tabulated value. In particular, deviations are insignificant when one is interested in confidence contours that are far from parameter boundaries or that are isolated in a locally parabolic region of the $\Delta\chi^2(\boldsymbol{\alpha})$ surface. Generally, though, one needs the complete ‘‘cutoff’’ surface $\Delta\chi_c^2(\boldsymbol{\alpha})$ for all parameters $\boldsymbol{\alpha}$ and for each desired confidence level c . The cutoff at $\boldsymbol{\alpha}$ can be determined by performing many simulated experiments with signal $\boldsymbol{\alpha}_{\text{true}} = \boldsymbol{\alpha}$ and by examining the resulting distribution of $\Delta\chi^2(\boldsymbol{\alpha}_{\text{true}})$ from Eq. (12.12). There exists some ξ for which a fraction c of the obtained $\Delta\chi^2(\boldsymbol{\alpha}_{\text{true}})$ values satisfies $\Delta\chi^2(\boldsymbol{\alpha}_{\text{true}}) < \xi$. It follows, then, that $\Delta\chi_c^2(\boldsymbol{\alpha}_{\text{true}}) = \xi$.

The $\boldsymbol{\alpha} = \boldsymbol{\alpha}_{\text{null}}$ case

As an example of this exercise, we calculate $\Delta\chi_c^2(\boldsymbol{\alpha}_{\text{null}})$. Figure 12.6 shows the $\Delta\chi^2(\boldsymbol{\alpha}_{\text{null}})$ distribution obtained from an ensemble of simulated signal-free experiments ($\boldsymbol{\alpha}_{\text{true}} = \boldsymbol{\alpha}_{\text{null}}$).

⁷We use the covariance matrix from the final iteration.

34400 fake experiments (no signal present)

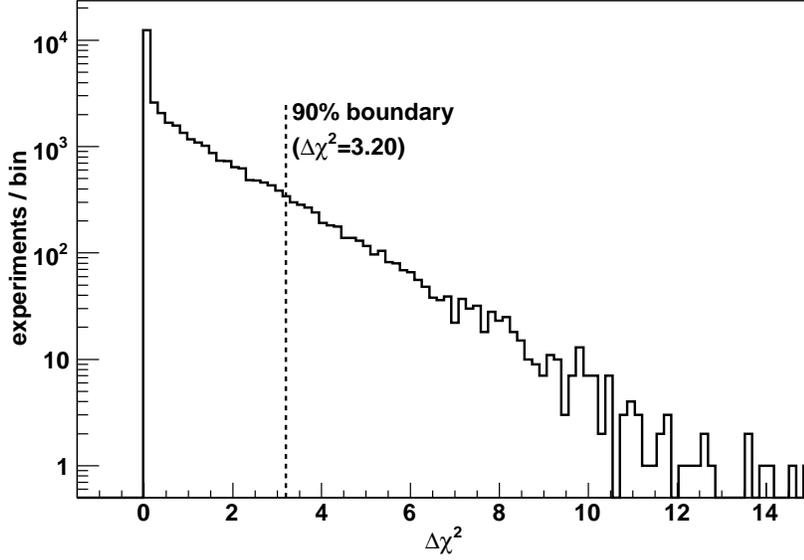


Figure 12.6: Values of $\Delta\chi^2(\alpha_{\text{null}})$ obtained in 34 400 fake signal-free experiments. 90% of experiments have $\Delta\chi^2(\alpha_{\text{null}}) < 3.20$. 95.7% of experiments have $\Delta\chi^2(\alpha_{\text{null}}) < 4.61$. (4.61 is the expected 90% cutoff for a two-parameter unbounded parabolic problem.)

Naively, we would expect 90% of the $\Delta\chi^2$ values to satisfy $\Delta\chi^2 < 4.61$. In fact, 95.7% of simulated experiments satisfy the condition. Table 12.4 presents various integrals of the distribution. We find that $\Delta\chi^2_{c=90\%}(\alpha_{\text{null}}) = 3.20$.

In this α_{null} case, the cutoff deviation is due primarily to the parameter bounds. An experiment whose E_ν^{rec} histogram fits well to an unphysical ($\sin^2 2\theta < 0$) scenario could lead to a lower χ^2_{best} were it allowed, but it gets stopped at the $\sin^2 2\theta = 0$ boundary. $\Delta\chi^2 = 0$ is thus over-represented (since $\alpha_{\text{null}} \Leftrightarrow \sin^2 2\theta = 0$).

The above procedure [139] must be repeated throughout parameter space to establish the

desired C.L.	simple $\Delta\chi^2(\alpha_{\text{null}})$	required $\Delta\chi^2(\alpha_{\text{null}})$	coverage using simple cutoff
68.27%	2.30	1.23	83.1(2)%
90.00%	4.61	3.20	95.7(1)%
99.00%	9.21	7.11	99.75(3)%
99.73%	11.83	9.08	99.96(1)%

Table 12.4: Properties of the $\Delta\chi^2$ distribution (Figure 12.6) obtained from an ensemble of fake signal-free experiments ($\alpha_{\text{true}} = \alpha_{\text{null}}$). For simple problems with two fitted parameters, the $\Delta\chi^2$ cutoffs listed in the 2nd column yield the confidence levels listed in the 1st. For our actual χ^2 function and parameter bounds, the desired confidence levels are reached at α_{null} by using the cutoffs listed in the 3rd column. The last column shows the coverage actually obtained with the simple cutoff values. (Statistical uncertainty is indicated parenthetically.)

complete cutoff surface. Because this is quite CPU intensive, many oscillation experiments forgo it, opting instead to use standard $\Delta\chi_c^2$ values. We do the same, except when quoting acceptance/rejection of the null hypothesis.

The $\alpha \neq \alpha_{\text{null}}$ cases

To verify that the standard $\Delta\chi_c^2$ values give acceptably accurate coverage away from α_{null} , we construct an approximate $\Delta\chi_c^2(\alpha)$ cutoff surface for $c=90\%$. For each of eighty-one signals (nine values of $\sin^2 2\theta$ crossed with nine values of Δm^2), $\mathcal{O}(1000)$ experiments are simulated. Each experiment's N -bin E_ν^{rec} histogram is created by drawing a vector \mathbf{p} of bin heights from the p.d.f.

$$G(\mathbf{p}) = \frac{1}{((2\pi)^N |\mathbf{V}|)^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{p} - \hat{\mathbf{p}})^\top \mathbf{V}^{-1}(\mathbf{p} - \hat{\mathbf{p}})\right), \quad (12.14)$$

where $\hat{\mathbf{p}}$ is the default Monte Carlo prediction for the bin heights (assuming signal α_{true}) and \mathbf{V} is the total covariance matrix for $\hat{\mathbf{p}}$. The resulting histogram is treated as data, and the fit procedure of §12.4.1 is performed. $\Delta\chi^2(\alpha_{\text{true}})$ is recorded.

As before, we set $\Delta\chi_{c=90\%}^2(\alpha_{\text{true}})$ such that 90% of experiments satisfy $\Delta\chi^2(\alpha_{\text{true}}) < \Delta\chi_{c=90\%}^2(\alpha_{\text{true}})$. Doing this for each of the 9×9 values of α_{true} yields the surface $\Delta\chi_{c=90\%}^2(\alpha)$ shown in the left panel of Figure 12.7. Because CPU constraints keep us from producing a sufficiently smooth surface and because we need a continuous function in order to evaluate $\Delta\chi_{c=90\%}^2(\alpha)$ for any α , we fit the 81 points to an analytic function which, frankly, doesn't work all that well at the higher Δm^2 values but which will suffice for our validation purposes. This smoothed surface is shown in the right panel of Figure 12.7. We will see in Chapter 14 that this surface gives similar contours to those obtained with the standard 90% cutoff value.

12.4.3 Raster scan confidence intervals

The global scan procedure we have been discussing is just one of several approaches for presenting a neutrino oscillation result. While it is used by some experiments, still others report confidence intervals using a “raster scan” method. A raster scan contour represents either a two-sided range of allowed $\sin^2 2\theta$ values as a function of Δm^2 or a one-sided upper limit on $\sin^2 2\theta$ as a function of Δm^2 . MiniBooNE has historically calculated sensitivity contours with a one-sided raster scan, and the oscillation results in this work are presented using this method, although we also show global scan contours (Chapter 14). The raster scan procedure is as follows.

At a fixed Δm^2 , the $\sin^2 2\theta$ value that minimizes χ^2 is found iteratively, updating the covariance matrix \mathbf{V} to account for any fitted signal after each iteration. Using the best-fit mixing angle $\sin^2 2\theta_{\text{best}}$ and the final covariance matrix, we find $\sin^2 2\theta_{\text{limit}}$ such that

$$\chi^2(\sin^2 2\theta_{\text{limit}}, \Delta m^2) - \chi^2(\sin^2 2\theta_{\text{best}}, \Delta m^2) = 1.64 \quad (12.15)$$

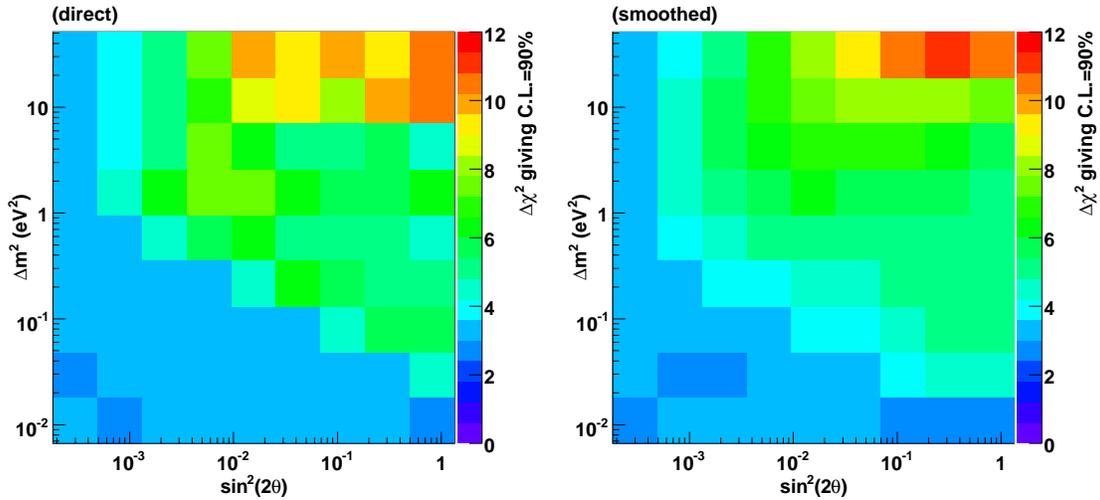


Figure 12.7: (Left) 90% C.L. $\Delta\chi^2$ cutoffs determined at eighty-one points in parameter space. (Right) The cutoffs evaluated using the fitted analytic function.

and $\sin^2 2\theta_{\text{limit}} > \sin^2 2\theta_{\text{best}}$. This is repeated at all Δm^2 values. The union of $(\sin^2 2\theta_{\text{limit}}, \Delta m^2)$ pairs forms the 90% C.L. upper limit contour for $\sin^2 2\theta$.⁸

12.5 Oscillation sensitivity

A natural, if somewhat cumbersome, definition of sensitivity is put forth in Ref. [139]: the $\sin^2 2\theta$ sensitivity at Δm^2 is given by the ensemble average of $\sin^2 2\theta_{\text{limit}}$ values obtained at that Δm^2 from a collection of simulated experiments.⁹ We use an alternate sensitivity definition¹⁰ which results in similar contours with much less computational overhead: the sensitivity contour is that contour obtained from a single fake experiment in which the data exactly matches the null hypothesis prediction.¹¹

Figure 12.8 shows the 90% sensitivity contour for the analysis. By removing individual component matrices (*e.g.*, \mathbf{V}^{π^+}), we can explore how the sensitivity might improve if effort were put into reducing each error source. The four panels of Figures 12.9 and 12.10 show the results of such an exercise. Some comments:

- Eliminating K^+ production uncertainties leads to the largest single improvement, followed by cross section uncertainties.
- No other single systematic component provides noticeable gain when removed.

⁸Since we are working at fixed Δm^2 values, the $\chi^2(\sin^2 2\theta)$ function is parabolic, so we do not face that issue. Effects from the $\sin^2 2\theta$ boundaries are found to be minimal for the contours we eventually form, so the standard cutoff of 1.64 is used.

⁹This is not well-defined for global scan contours which may be zero- or double-valued for some Δm^2 .

¹⁰likely used by other experiments, though this is often impossible to tell from journal articles

¹¹This definition is compatible with a global scan.

- Eliminating all systematic errors together (a rather unrealistic scenario) improves the sensitivity considerably more than any single component does. (Consider, for example, the ν_e -from- K^+ background. Reducing the ν_e flux uncertainty leaves large cross section uncertainties, and *vice versa*. One must reduce both to see significant gains.)
- Removing the statistical error is as beneficial as removing all systematic error.^{12,13} (This does not include the secondary benefit of eliminating statistical errors in the π^0 and ν_μ constraint procedures.)

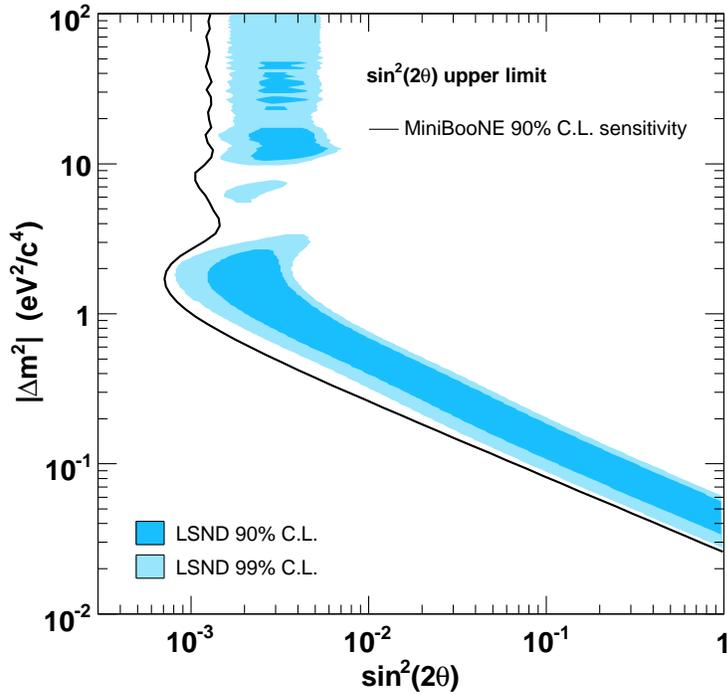


Figure 12.8: The 90% C.L. sensitivity contour. The LSND allowed region is shown for context.

¹²Since we cannot remove statistical errors without also removing detector model errors, the “syst. errors only” case excludes both, via $\mathbf{V}^{\text{det}} = 0$. However, we can approximate the converse situation to show that statistical errors overwhelm the detector errors. (See next note.)

¹³To turn off only detector model errors, we cannot just set $\mathbf{V}^{\text{det}} = 0$ since that also removes statistical errors. Instead, we replace \mathbf{V}^{det} with \mathbf{V}^{stat} , a swap that has the effect of subtracting out and then adding back in a significant contribution to \mathbf{V} . Since the Monte Carlo-based \mathbf{V}^{det} provides only an estimate of \mathbf{V}^{stat} , the cancellation is not perfect. In our case, the resulting $\mathbf{V}^{\text{det}} \leftrightarrow \mathbf{V}^{\text{stat}}$ contour falls slightly to the right of the all-errors contour. Figure 12.16 is related to this diagnostic complexity.

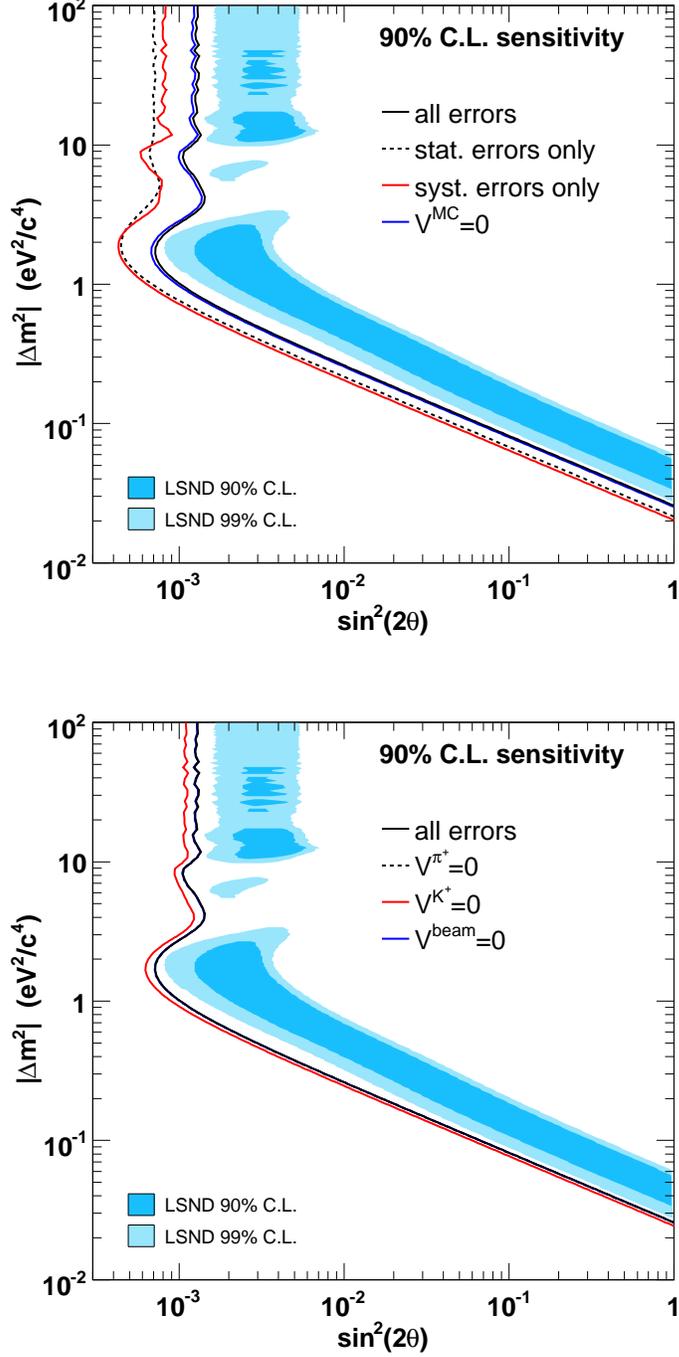


Figure 12.9: (Part 1 of 2.) 90% C.L. sensitivity contours with various error sources removed. In the top panel: (solid black) all errors; (dashed black) statistical errors only [*N.B.* footnote 12]; (red) systematic errors only; (blue) no errors from finite Monte Carlo statistics. In the bottom panel: (solid black) all errors; (dashed black) no π^+ production uncertainty [hidden behind the “all errors” contour]; (red) no K^+ production uncertainty; (blue) no “beamline” errors [nearly hidden].

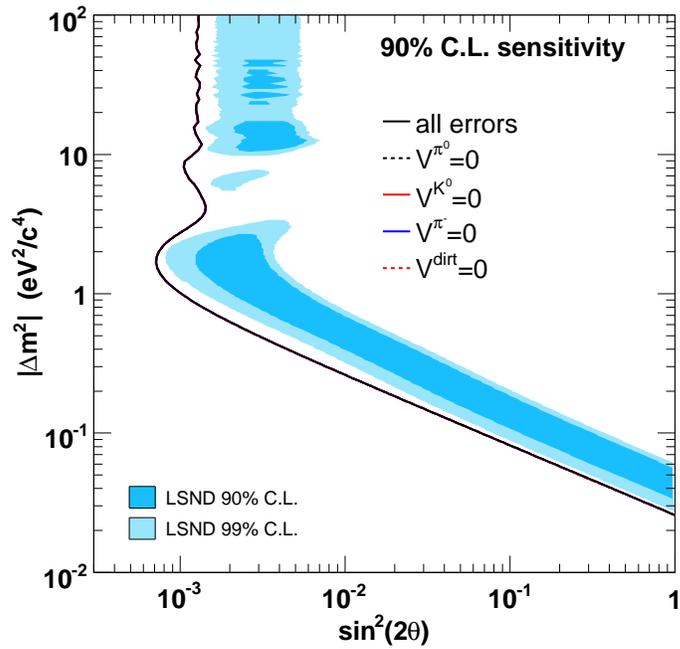
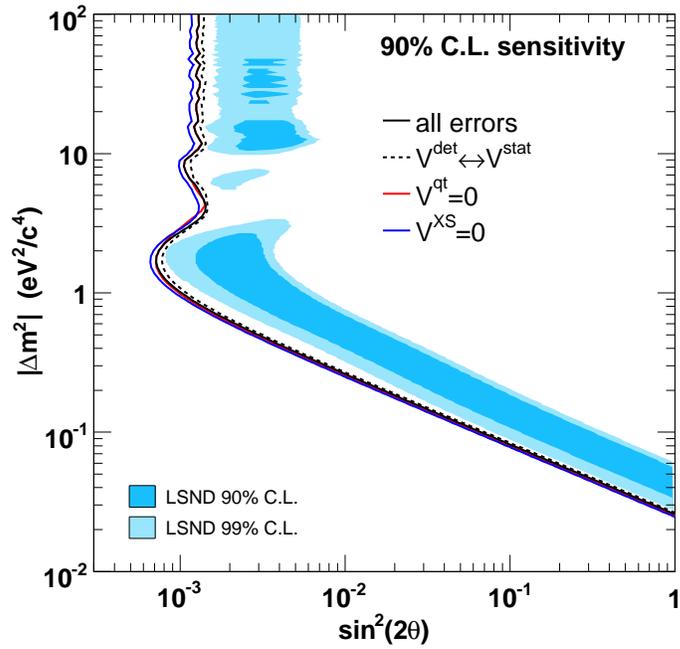


Figure 12.10: (Part 2 of 2.) 90% C.L. sensitivity contours with various error sources removed. In the top panel: (solid black) all errors; (dashed black) no detector model uncertainties [*N.B.* footnote 13]; (red) no charge/time correlation uncertainty; (blue) no cross section uncertainties. In the bottom panel: all contours lie on top of one another. Included are contours for: all errors; no π^0 rate uncertainty; no K^0 production uncertainty; no π^- production uncertainty; no dirt normalization uncertainty.

Counting analysis

A more straightforward signal test, less sensitive than fitting the energy spectrum, is simply to count the number of ν_e candidate events and to compare this with the background expectation. We present counting results in Chapter 14 using events that fall in the energy range $475 \text{ MeV} < E_\nu^{\text{rec}} < 1250 \text{ MeV}$. The upper E_ν^{rec} limit is below than that used in the energy fit (3000 MeV) since negligible oscillation signal can appear above 1250 MeV. So while higher energy events provide normalization information in the energy fit, they contribute only noise to the counting analysis. Figure 12.11 shows the oscillation sensitivity for the counting analysis.

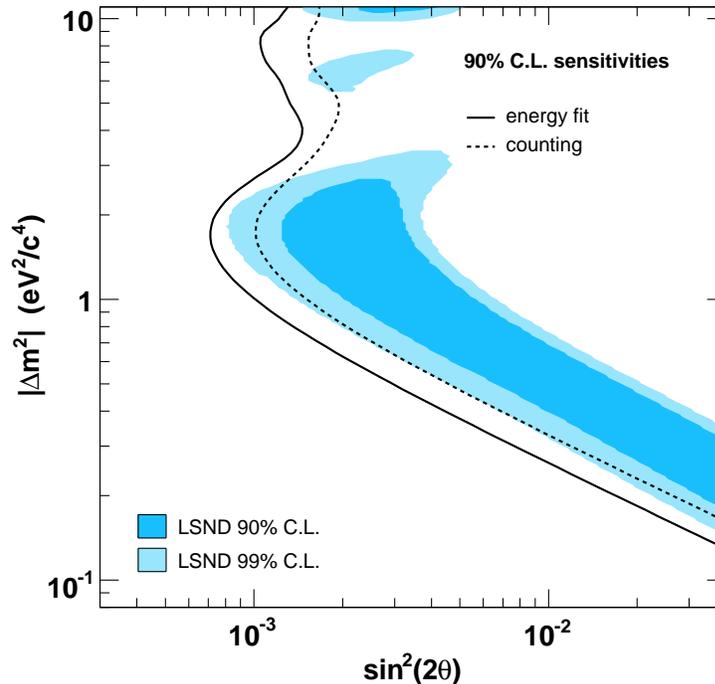


Figure 12.11: 90% C.L. sensitivity contours for the E_ν^{rec} -fit and counting analyses.

12.6 Testing the fit

Simulated experiments, like those that provide the $\Delta\chi_c^2$ surface in Figure 12.7, can also be used to test the signal extraction procedure. We show in Figure 12.12 the distribution of $\Delta\chi^2(\alpha_{\text{true}})$ obtained from 86 600 experiments whose true signals are drawn from all over parameter space, up to $\Delta m^2 \sim 12 \text{ eV}^2$. If fits were getting trapped in non-global minima, $\chi^2(\alpha_{\text{true}})$ would occasionally be lower than χ_{best}^2 . The lack of any negative $\Delta\chi^2(\alpha_{\text{true}})$ values suggests that this does not occur. A prototype minimization algorithm, which involved only a single Minuit call, indeed gave many negative values.

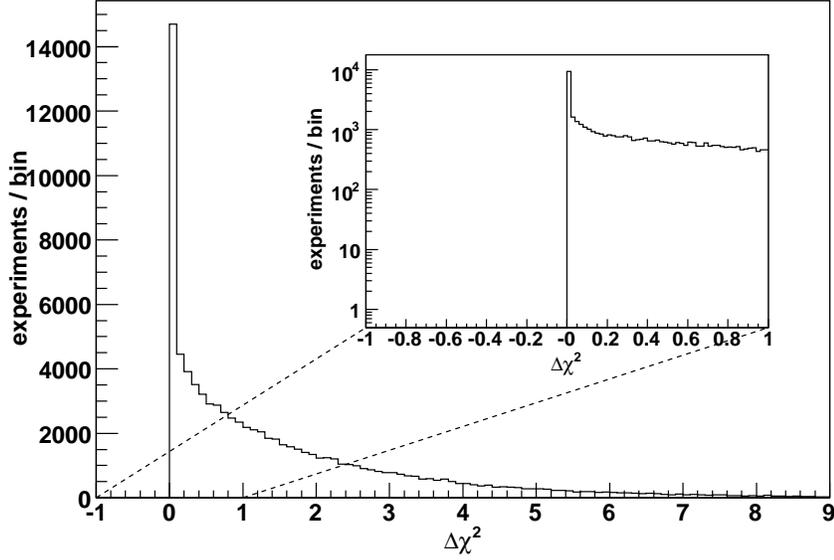


Figure 12.12: $\Delta\chi^2(\boldsymbol{\alpha}_{\text{true}}) = \chi^2(\boldsymbol{\alpha}_{\text{true}}) - \chi_{\text{best}}^2$ for 86 600 fake experiments with signal. No negative values are seen, supporting the claim that the χ^2 minimization algorithm does not get trapped in non-global minima. (The left edge of the first occupied bin has been set to -10^{-9} to absorb floating-point effects.)

Figure 12.13 shows the best-fit parameters extracted from hundreds of experiments at three different true signal points. Systematic errors were turned off both in the fit and in the fake data creation. The scattered points distribute themselves about the true signal points in the expected fashion. The corresponding χ_{best}^2 distributions are also well-behaved, as explained in the caption. Figure 12.14 shows analogous graphs made with both statistical and systematic uncertainties.

In these exercises, the generation of fake data and the calculation of χ^2 both assume a given covariance matrix \mathbf{V} . Thus, while they represent strong tests of the fit machinery, these studies cannot validate the construction of the covariance matrix itself. (Recall that the fake E_ν^{rec} spectra above are created from the p.d.f. in Eq. (12.14).) To fill this gap, we also performed fits on fake data built by directly varying the Monte Carlo simulation according to underlying parameter uncertainties. That is, the parameters controlling the simulation (affecting secondary meson production, neutrino cross sections, the detector model, etc.) were simultaneously adjusted in a fashion dictated by the uncertainties on each. The resulting fake data, which sometimes included signal, were then fit using a covariance matrix formed with the full analysis infrastructure, including new ν_μ CC QE constraints obtained from corresponding fake ν_μ CC QE samples, per Chapter 10. As Figure 12.15 shows, the χ^2 values from these “fully faked” experiments are distributed as expected, indicating that our error propagation procedure indeed yields the intended covariance matrix.

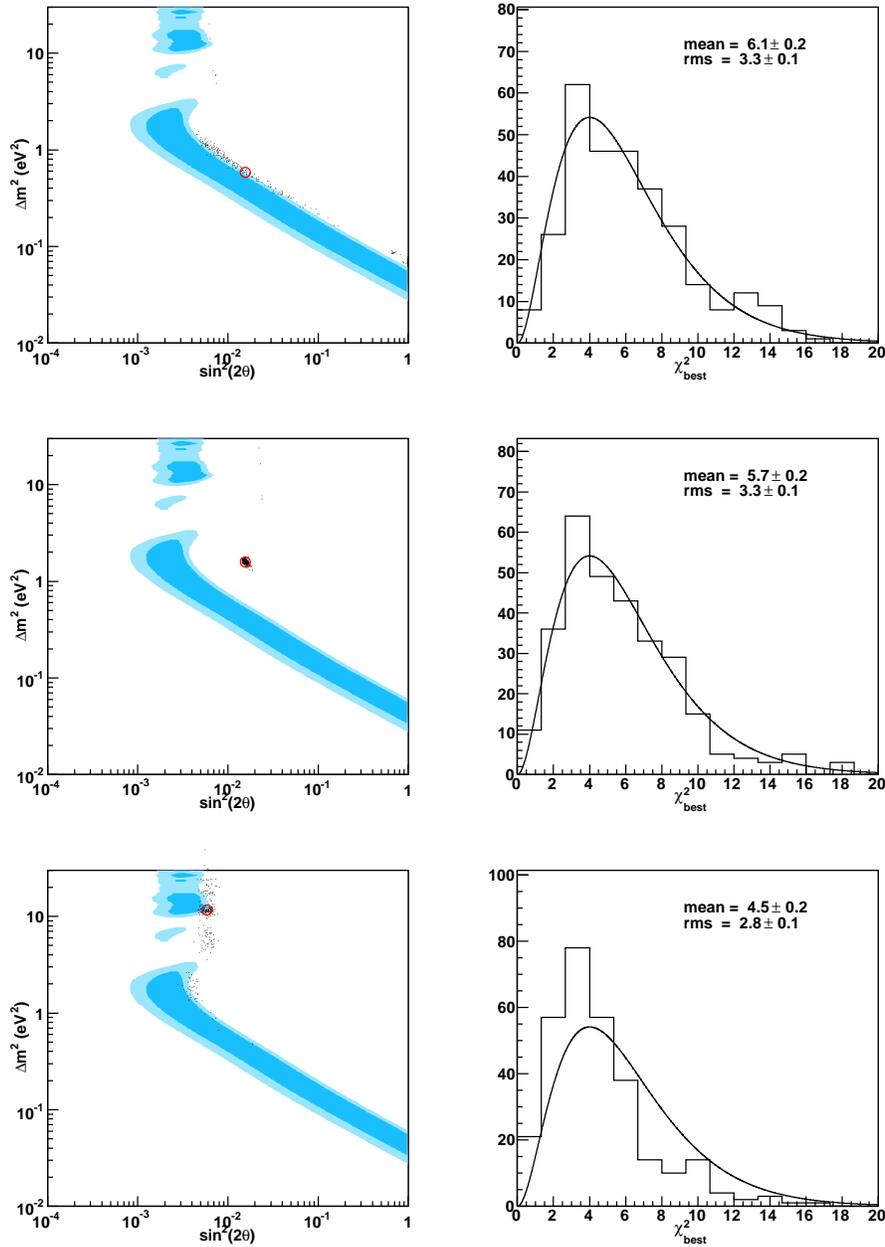


Figure 12.13: 300 fake experiments for three different true signal points, no systematic errors. (Left column) The true parameter values are indicated with red circles. The blue LSND allowed regions are included for context. The scattered black points mark the best-fit parameters extracted in the experiments. The valleys of the χ^2 surface are echoed in the scatter plots. (Right column) Histograms of χ^2_{best} . The smooth curve shows the χ^2 p.d.f. for 6 degrees of freedom, $P(\chi^2; N_{\text{dof}}=6)$ (mean=6, RMS=3.5), scaled to the number of experiments. (Note: 8 bin fit, 2 free parameters.) As the best-fit points in the top two rows typically fall within a single valley, we expect non-parabolic features of the χ^2 surface to play a small role. Thus, the χ^2_{best} distributions follows the naive $P(\chi^2; N_{\text{dof}}=6)$ expectation closely. In the bottom row, χ^2 “valley hopping” is prevalent and gives the fit additional freedom to lower χ^2_{best} . Studies like these were performed throughout parameter space with consistent results.

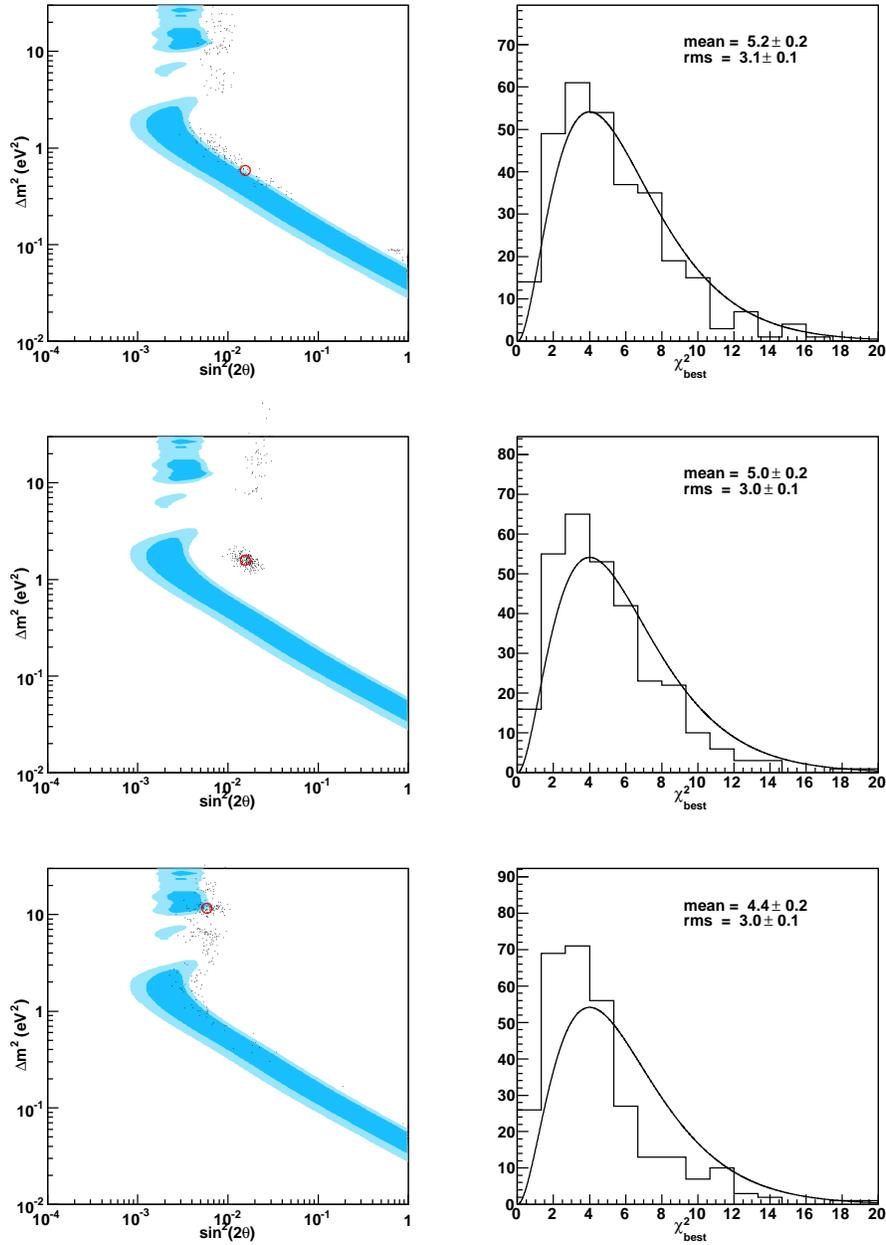


Figure 12.14: 300 fake experiments for three true signal points, all errors included. See Figure 12.13 for descriptions of the graphs. In addition to the nuances mentioned there, large systematic errors (especially at high energies) result in a deflated average χ^2_{best} , as the generated data histograms were not allowed to have negative bin contents.

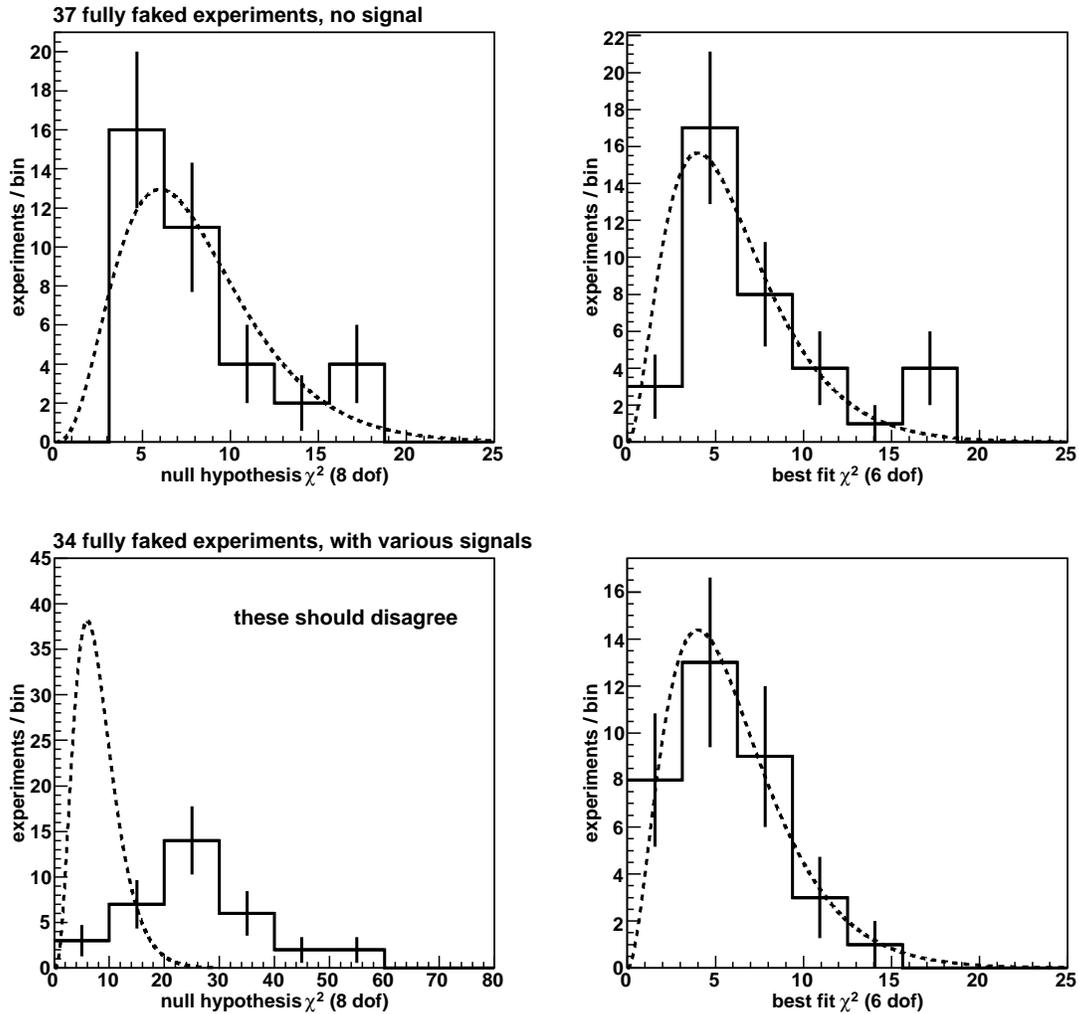


Figure 12.15: (Top row) Fully faked data, no signal. The left panel shows the distribution of χ^2 between fake data and a background-only prediction (no fit). The overlaid curve is proportional to $P(\chi^2; N_{\text{dof}}=8)$. The right panel shows χ^2_{best} along with $P(\chi^2; N_{\text{dof}}=6)$. (Bottom row) Various signals are now present in the fake data. Direct comparisons of data and the background-only prediction (left) should give large χ^2 values, as the null hypothesis is a poor description. The distribution of χ^2_{best} , though, should be (and is) a close match to the overlaid χ^2 function. These studies are CPU limited, and only a few dozen were performed. Relatedly, the seventy-one experiments shown were generated via reweighting from only eleven base Monte Carlo samples. Therefore, only eleven statistical and detector model variants are represented. All other uncertainty components were given independent parameter draws for each experiment.

A final check: Recall that the component matrix \mathbf{V}^{det} is unique in that it includes the statistical error estimate and that it is built from $K=64$ independent, data-sized Monte Carlo samples rather than from a single large sample via weighting. For large covariance matrices (say, for an $N_{\text{bin}}=30$ histogram), impractically large values of K are needed to combat noise in \mathbf{V}^{det} . For fixed K , one must choose a sufficiently small N_{bin} such that χ^2 shows the proper behavior. We use $N_{\text{bin}}=8$ in this analysis. The above “fully faked” experiments demonstrate that our χ^2 is well-behaved, and Figure 12.16 demonstrates that any residual noise in the covariance matrix has no significant impact on the contours derived from the $\Delta\chi^2$ surface. Studies with $N_{\text{bin}}=16$ suggest that little oscillation sensitivity would be gained by increasing K (to allow for higher N_{bin}).

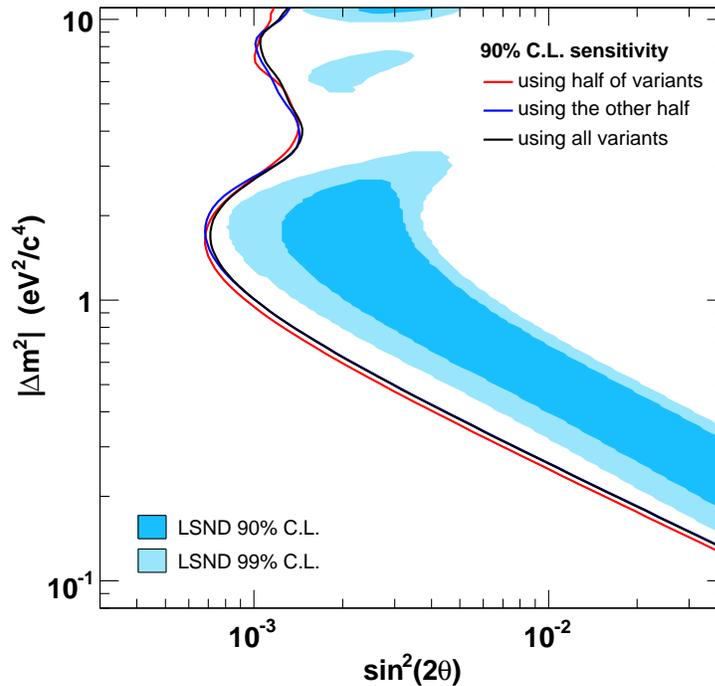


Figure 12.16: 90% C.L. sensitivity contours obtained using thirty-two detector model variants (red), a different thirty-two variants (blue), and all sixty-four variants (black). The differences are small, indicating that the gross shape of the contour is not driven by noise in the covariance matrix.

Chapter 13

Verifying the predictions

Event distributions produced using the tuned Monte Carlo simulation should agree with corresponding distributions from data within the assigned uncertainties. We would like to confirm, while maintaining blindness, that this data/simulation agreement actually occurs.

13.1 Non-signal samples

We begin by defining five event samples from which we create histograms that probe a range of analysis components, including μ and e modeling and reconstruction; π^0 production and misidentification; nuclear effects; error estimation; and the χ^2 procedure itself. Aspects of the analysis that cannot be readily tested while blind (for example, $\Delta \rightarrow N\gamma$ and low-to-mid energy ν_e rates) are addressed indirectly in the next section.

ν_μ CC QE

The first test sample is formed using the ν_μ CC QE selection of Chapter 10. The muons in these events allow for checks of lepton reconstruction in the presence of recoil hadrons; muon particle ID; and quasi-elastic kinematic distributions. The Michel electrons provide checks of low energy electron modeling and reconstruction. For distributions directly involved in Monte Carlo tuning (*e.g.*, Q^2), the comparisons we make here serve to quantify the success of that tuning.

High energy ν_e

High energy events passing the ν_e selection are primarily intrinsic beam ν_e 's. Signal ν_e 's appear at lower energy. We define a high energy ν_e sample consisting of all events that satisfy both $E_\nu^{\text{rec}} > 1.5$ GeV and the ν_e selection (Chapter 12). The expected contributions from (ν_e -from- K) : (ν_e -from- μ) : (NC π^0) are in proportion 1.7 : 1 : 1. The total number of events is small, but one could observe large ν_e or π^0 rate discrepancies, if present.

Sideband A

Three “sidebands” are created using the ν_e selection with modified π^0 cuts. For Sideband A, the standard $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ cut is dropped and these two cuts are added:

- $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0}) < -0.01$
- $M_{\gamma\gamma} < 50 \text{ MeV}/c^2$.

Figure 11.3 shows how sideband A relates to other samples. Sideband A, which consists almost entirely of π^0 events, is small but can potentially reveal π^0 misidentification discrepancies. (Note that these events are themselves partially misidentified, having low reconstructed masses.) When we compare $M_{\gamma\gamma}$ histograms in this sample, we leave out the $M_{\gamma\gamma} < 50 \text{ MeV}/c^2$ sideband cut, letting the E_e -dependent $M_{\gamma\gamma}$ cut from the ν_e selection shape the distribution. (This is indicated with the label A' in tables and figures.)

Sideband B

For Sideband B, the standard $M_{\gamma\gamma}$ cut is dropped and these are added:

- $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0}) > 0$
- $M_{\gamma\gamma} > 50 \text{ MeV}/c^2$.

Figure 11.3 shows this region graphically. When we examine $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$, the label B' is used, indicating that the $\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$ sideband cut is not applied (in analogy with A').

Sideband C

Sideband C replaces both π^0 cuts in the ν_e selection with:

- $-0.09 < \log(\mathcal{L}_e/\mathcal{L}_{\pi^0}) < 0$
- $50 \text{ MeV}/c^2 < M_{\gamma\gamma} < 100 \text{ MeV}/c^2$.

Again, see Figure 11.3. The boundaries are set such that the number of events in sideband C is similar to the number of events in the signal region. This provides a check of our statistical error propagation methods.

13.1.1 The comparisons

Over the next several pages we show the plots that compare prediction and data in the above samples. Each plot includes:

- The data as solid points.
- The Monte Carlo prediction as a red histogram.
- An error band. The half-width of the error band at bin i is $\sqrt{V_{ii}}$, where \mathbf{V} is the full covariance matrix.

- χ^2 , calculated with the full covariance matrix. This quantifies the data/simulation agreement. (The error band alone is insufficient since correlations can be important.)
- The χ^2 cumulative probability: $p = \int_{\chi^2}^{\infty} P(\eta) d\eta$, where $P(\eta)$ is the χ^2 p.d.f.

Table 13.1 provides a summary of the comparisons. The pages that follow show the plots. All predictions and covariance matrices are created following the procedure outlined in Chapter 12 and used for the signal fit itself, with the exception that the detector model variants (§12.3.2) in the ν_μ CC QE sample were smoothed via the procedure described in §12.3.3.¹

The agreement is excellent, and the cumulative χ^2 probabilities are well distributed from zero to one (Figure 13.1). The lowest p among the forty-six calculated is 0.004 for the U_x distribution in sideband B. This distribution's agreement was deemed tolerable, since the U_x distributions in other samples give no cause for alarm, since 250-to-1 odds are not exceedingly long given forty-six trials, and since close inspection of the distribution and its errors offered no reasons to change anything.

In closing, we note that the particular list of quantities tested in each sample is largely historical, as the samples were created for a variety of original purposes.

¹This is necessary if we are to compare more than eight bins or so, something we would like to do with the high-statistics ν_μ CC QE sample. We have too few detector model variants to estimate covariance matrices for histograms much larger than ~ 8 bins. Smoothing mitigates this issue, as it approximately eliminates the statistical contribution to \mathbf{V}^{det} , which we can replace by adding in a traditionally calculated \mathbf{V}^{stat} . Note that all other χ^2 comparisons, either in other test samples or in the signal extraction, use nine or fewer bins (usually eight), for which our ensemble of detector model variants is sufficiently large. (See §12.6.)

sample	quantity	χ^2	N_{dof}	$P(\chi^2+)$
ν_μ CC QE	Michel E	19.1	32	0.96
.	E	40.0	33	0.19
.	Michel R^3	43.0	37	0.23
.	R^3	34.4	37	0.59
.	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$	52.0	50	0.40
.	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$ (e -like portion)	11.1	21	0.96
.	Q^2	49.8	50	0.48
.	E_ν^{rec}	28.3	29	0.50
.	line cut residual	37.4	50	0.91
.	U_x	62.1	50	0.12
.	U_y	53.3	50	0.35
.	U_z	47.5	50	0.57
High energy ν_e	E_ν^{rec}	0.5	5	0.99
.	U_z	2.8	5	0.74
.	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$	8.3	6	0.22
.	$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$	7.8	5	0.17
.	$M_{\gamma\gamma}$	5.1	6	0.54
Sideband A	E_ν^{rec}	4.9	8	0.77
.	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$	2.6	8	0.96
.	$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$	6.7	9	0.67
(A')	$M_{\gamma\gamma}$	5.7	8	0.69
.	E	8.6	7	0.28
.	U_x	15.6	8	0.05
.	U_y	7.6	8	0.47
.	U_z	3.4	8	0.91
Sideband B	E_ν^{rec}	5.8	8	0.67
.	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$	5.0	8	0.76
(B')	$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$	10.8	8	0.21
.	$M_{\gamma\gamma}$	6.3	6	0.39
.	E	4.8	7	0.69
.	U_x	22.6	8	0.00
.	U_y	7.9	8	0.45
.	U_z	3.3	8	0.92
Sideband C	R_e^3	5.1	8	0.75
.	(endpoint R_μ) ³	8.4	8	0.39
.	E_ν^{rec}	7.3	8	0.51
.	E	8.1	7	0.33
.	X	1.8	8	0.99
.	Y	5.2	8	0.73
.	Z	5.8	8	0.67
.	U_x	4.0	8	0.86
.	U_y	15.8	8	0.04
.	U_z	4.4	8	0.82
.	$\log(\mathcal{L}_e/\mathcal{L}_\mu)$	5.9	8	0.66
.	$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$	1.6	4	0.80
.	$M_{\gamma\gamma}$	1.8	4	0.77

Table 13.1: Summary of comparisons between data and simulation. The columns in the table indicate (1) the test sample, (2) the variable being examined, (3) the χ^2 quantifying the agreement, (4) the number of degrees of freedom in the comparison, and (5) the cumulative χ^2 probability.

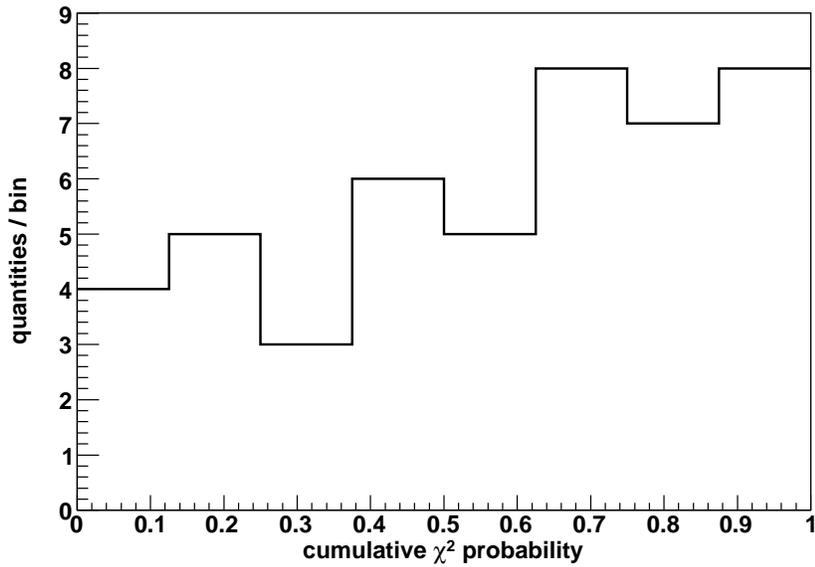
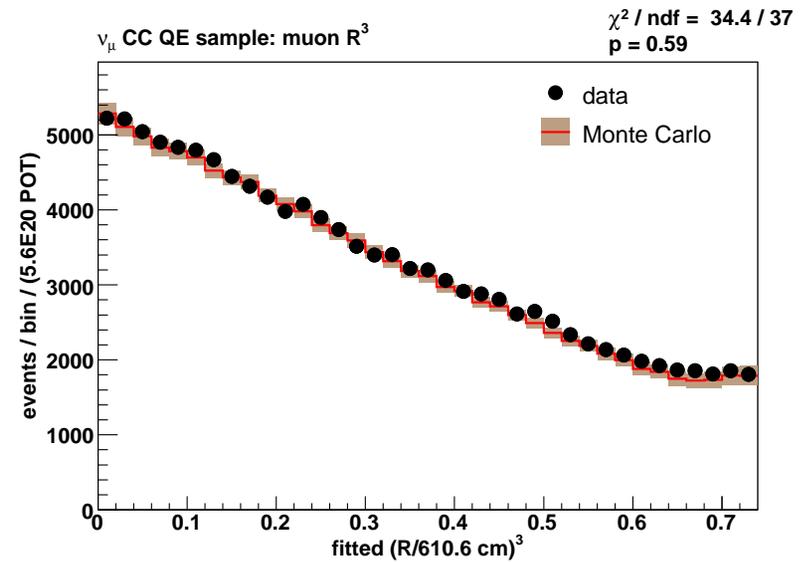
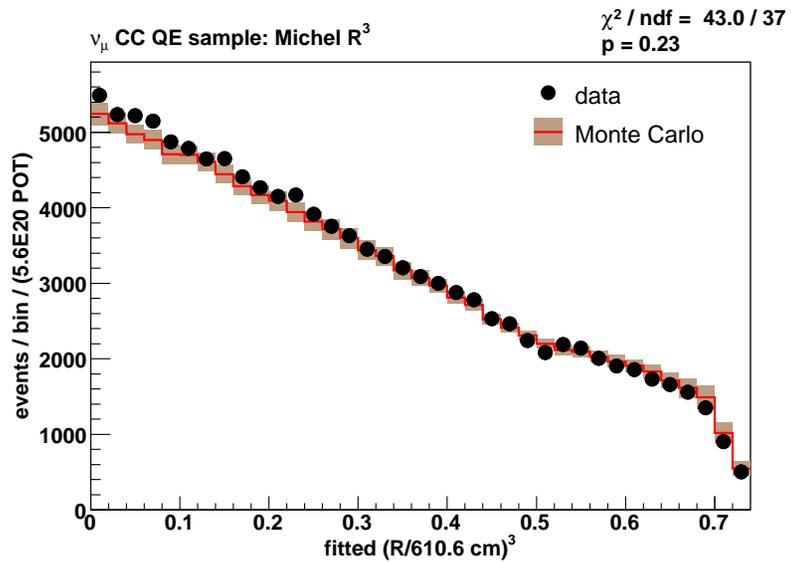
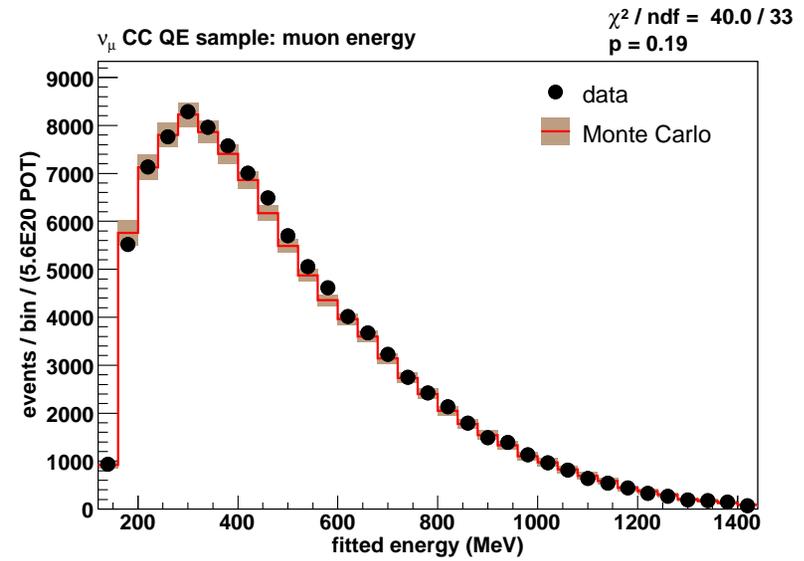
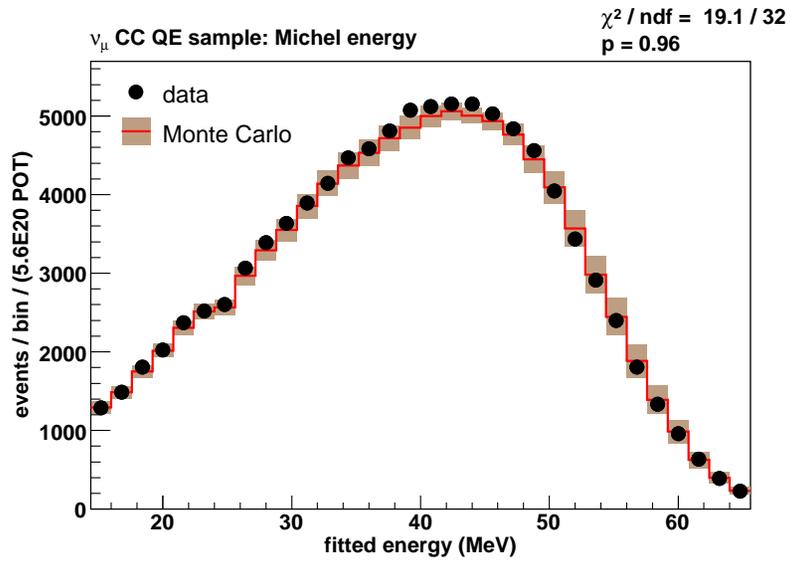


Figure 13.1: Distribution of cumulative χ^2 probabilities. Independent draws from a χ^2 distribution would yield a flat histogram here. Our situation involves correlated errors and non-independent quantities, so deviations from flat are possible. Additionally, the smoothing that facilitates the many-bin comparisons in the ν_μ CC QE sample results in slightly overestimated \mathbf{V}^{det} matrices, unavoidably biasing p upward for that sample. (An alternative is to compare fewer bins. This was done, and good data/simulation agreement persisted.)

ν_μ CC QE plots

Next 3 pages

Figure 13.2: ν_μ CC QE plots (Part 1).

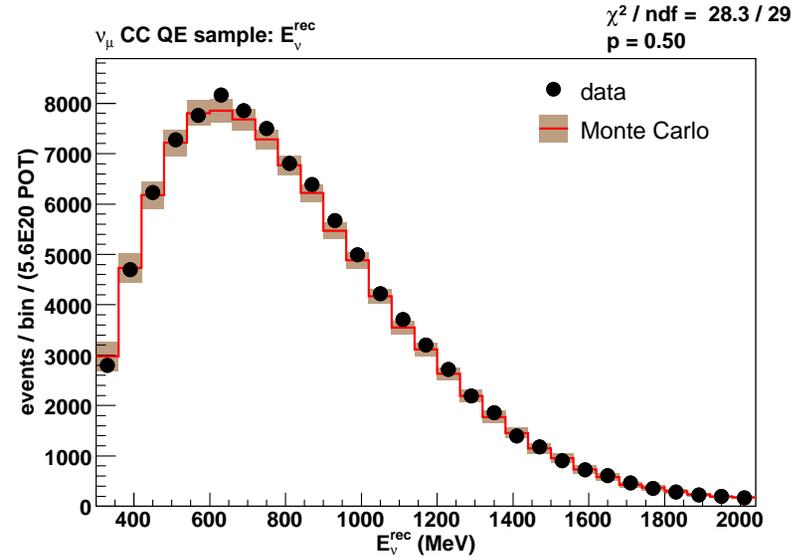
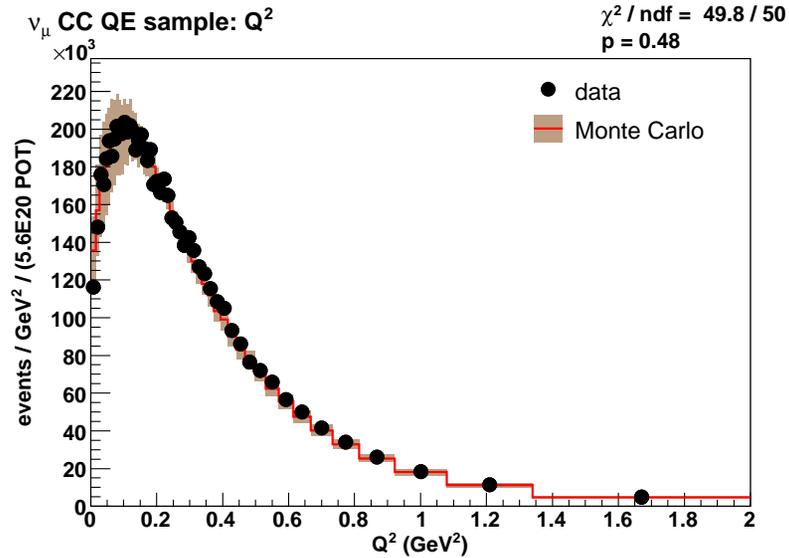
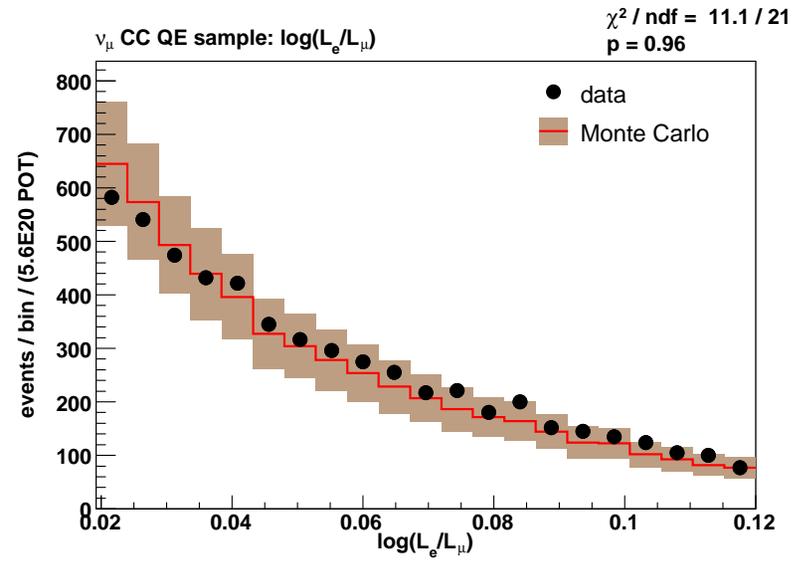
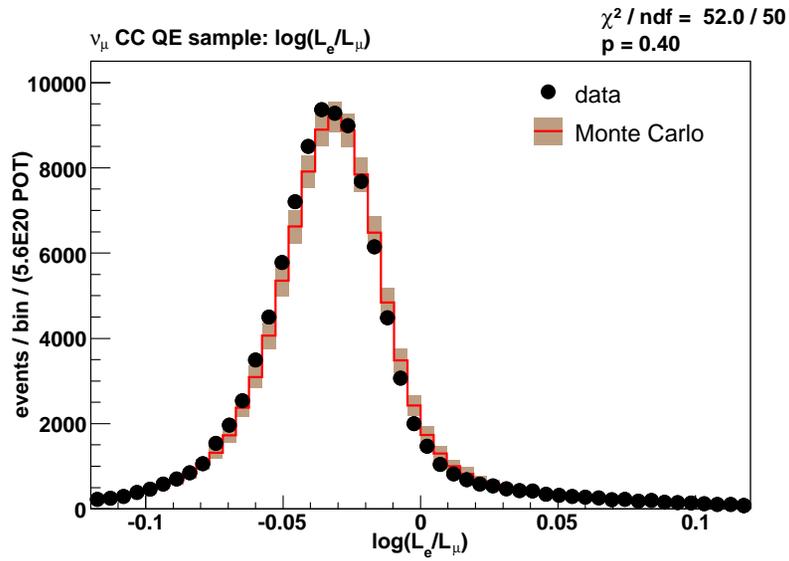
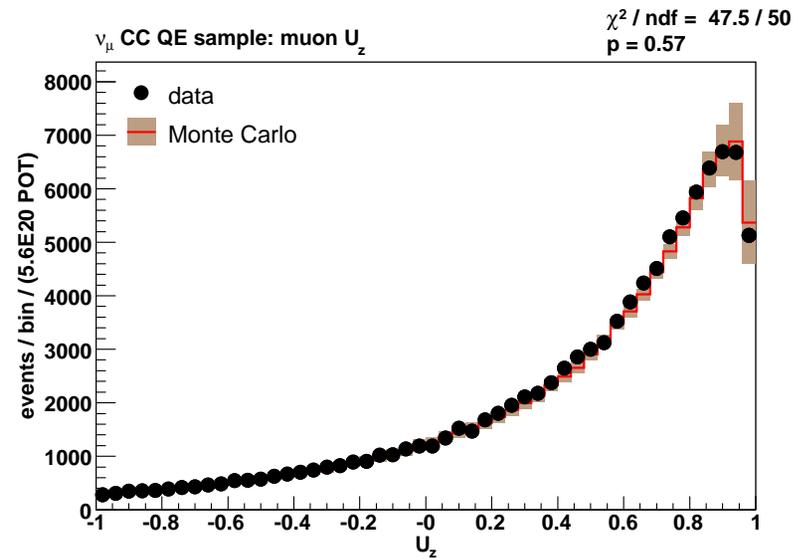
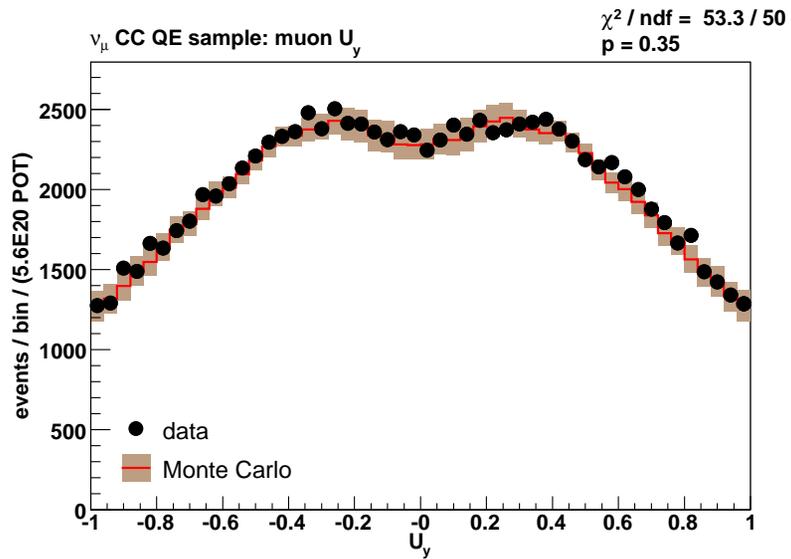
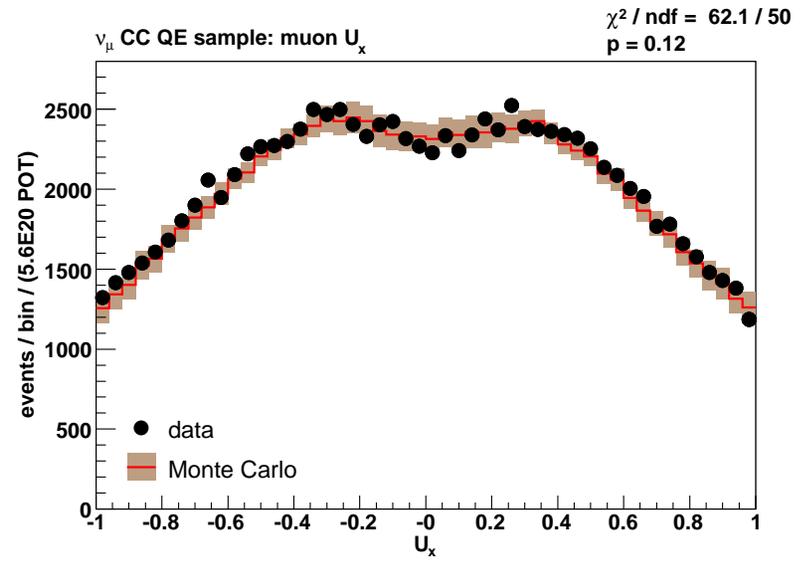
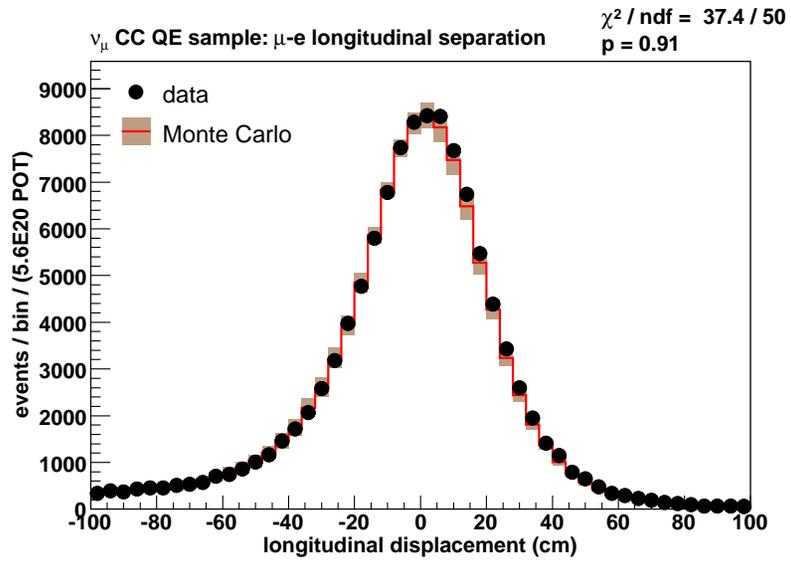
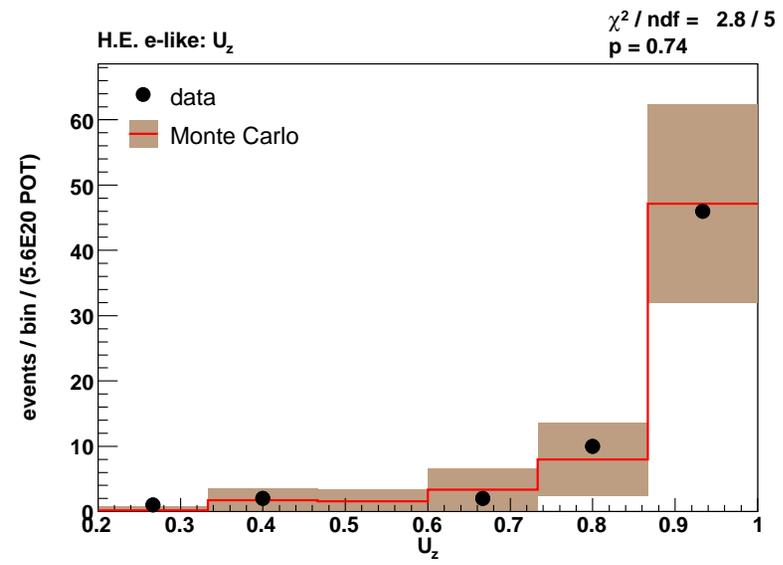
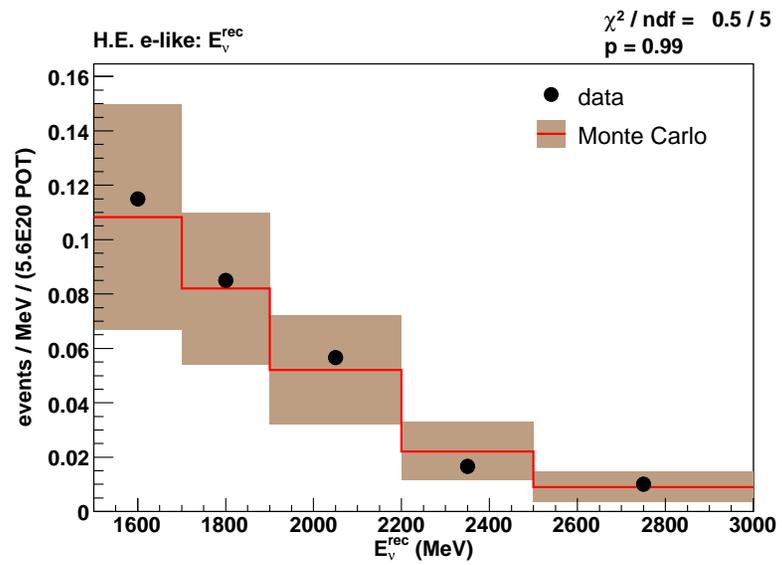


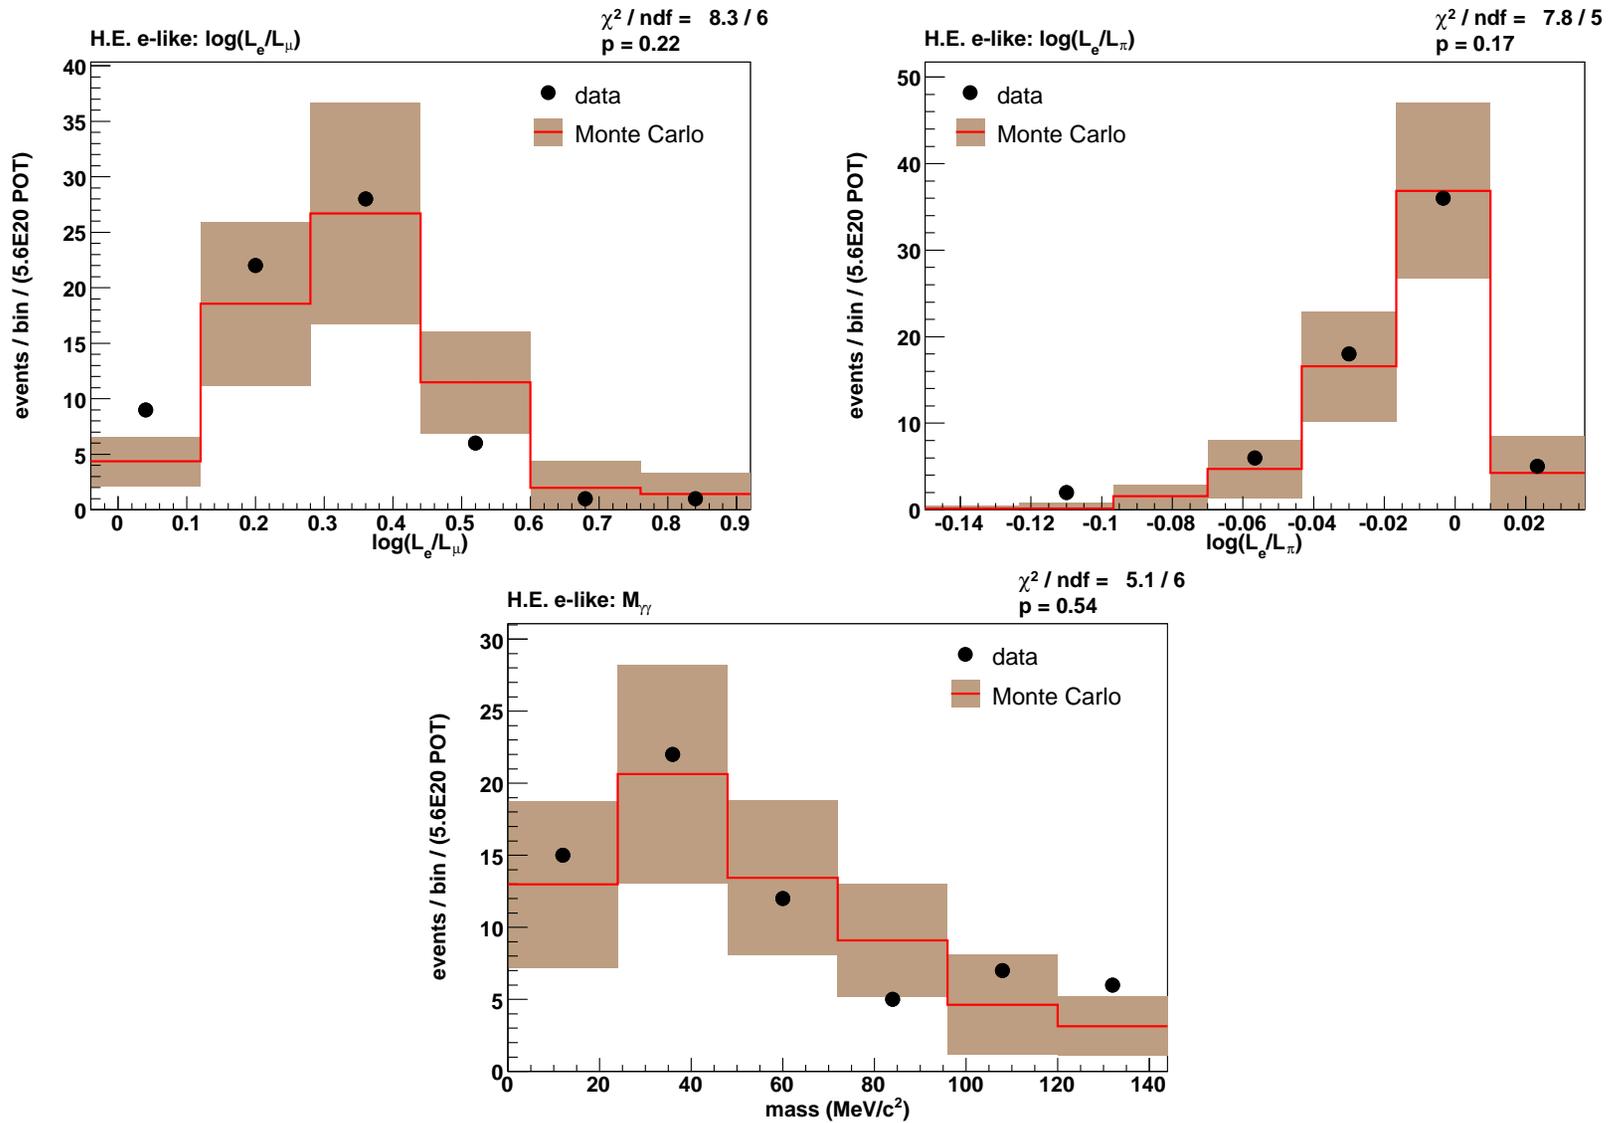
Figure 13.3: ν_μ CC QE plots (Part 2). The upper-right plot is a zoomed-in version of the upper-left plot, showing the e -like portion of the distribution (where muon misidentification occurs).

Figure 13.4: ν_μ CC QE plots (Part 3).

High energy ν_e plots

Next 2 pages

Figure 13.5: High energy ν_e plots (Part 1).

Figure 13.6: High energy ν_e plots (Part 2).

Sideband A plots

Next 2 pages

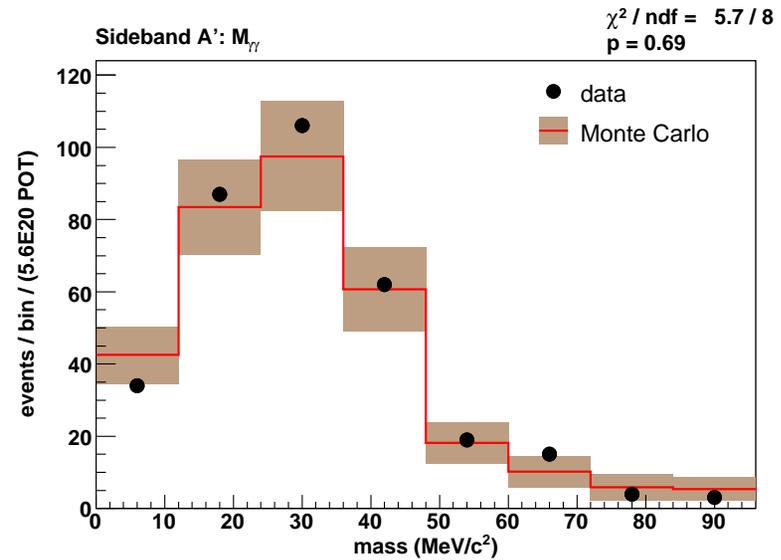
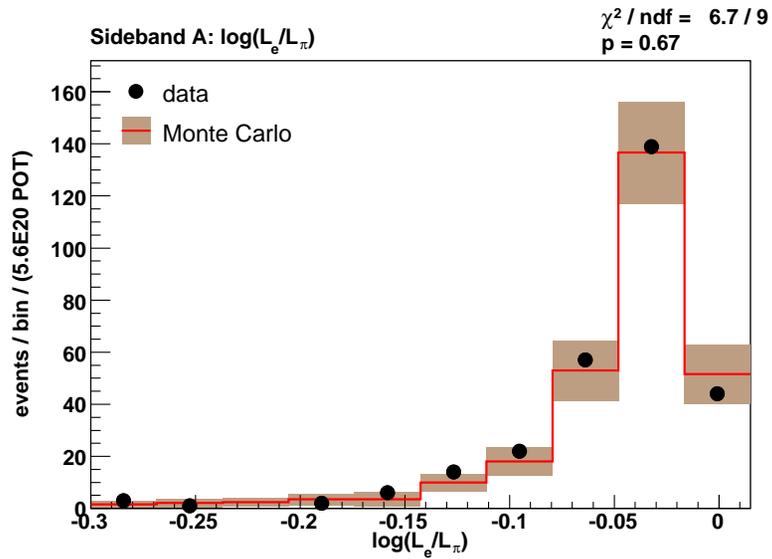
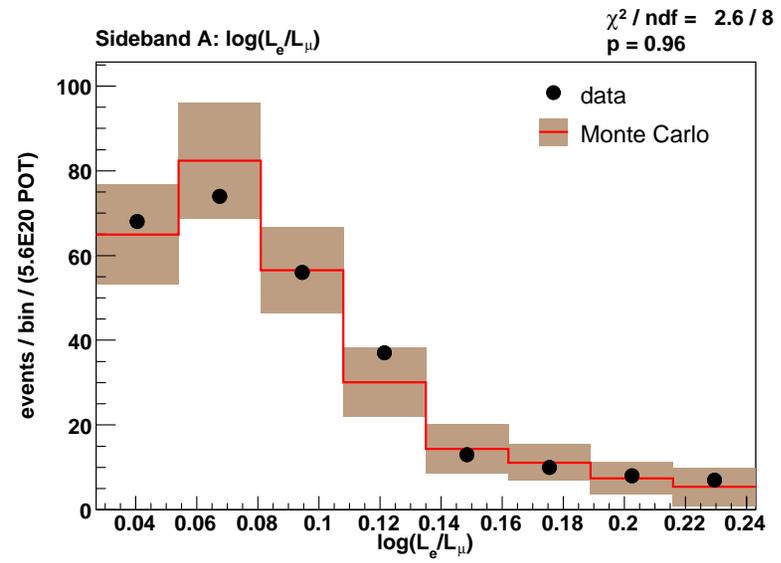
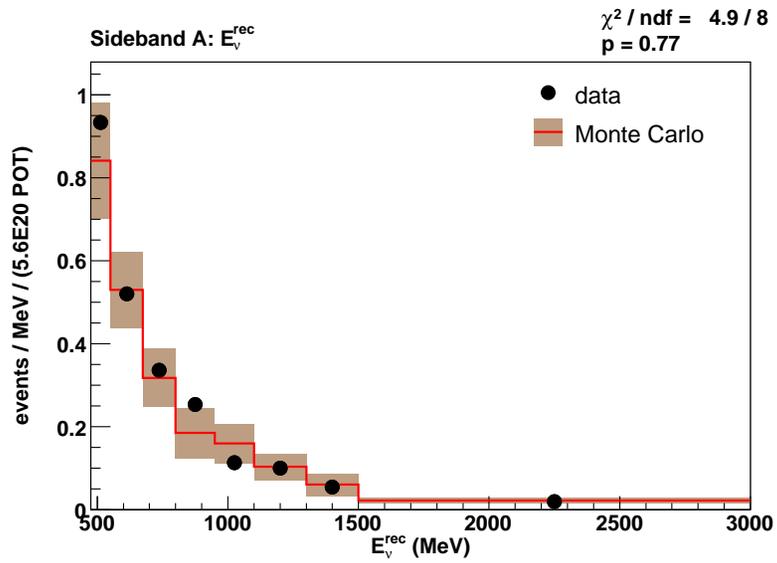


Figure 13.7: Sideband A plots (Part 1).

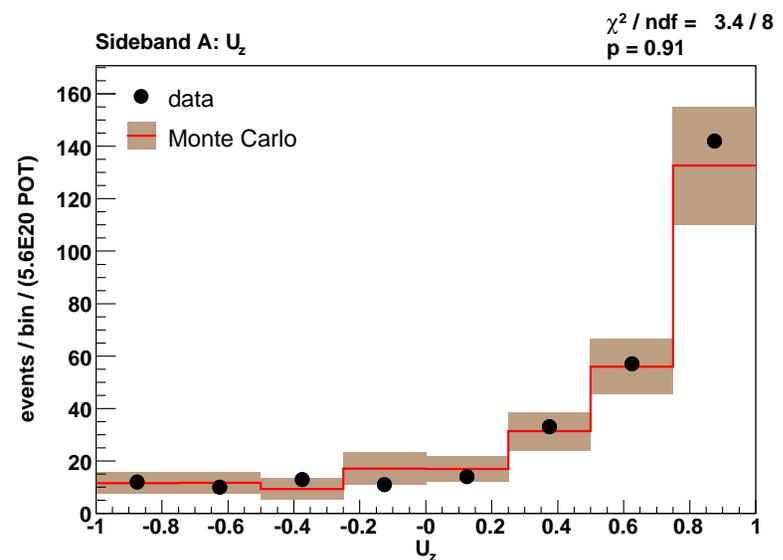
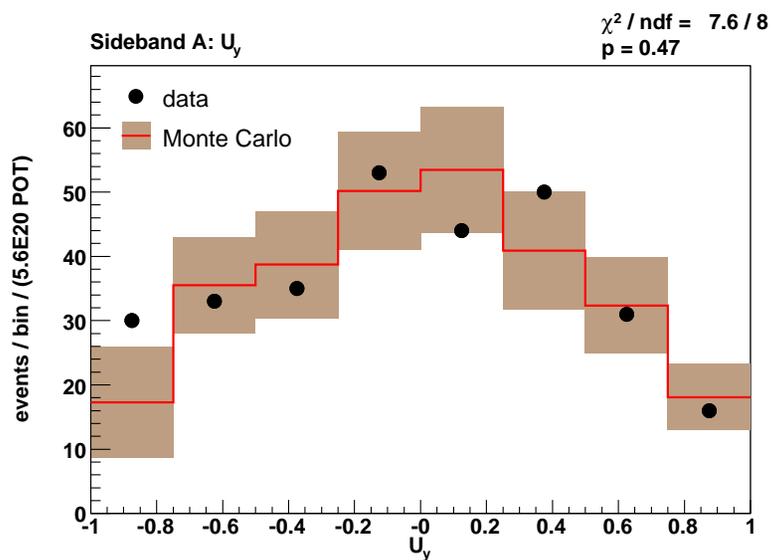
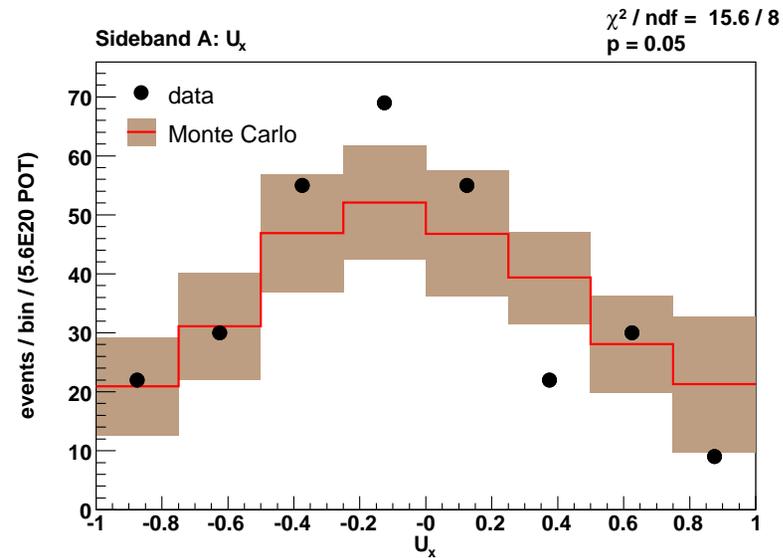
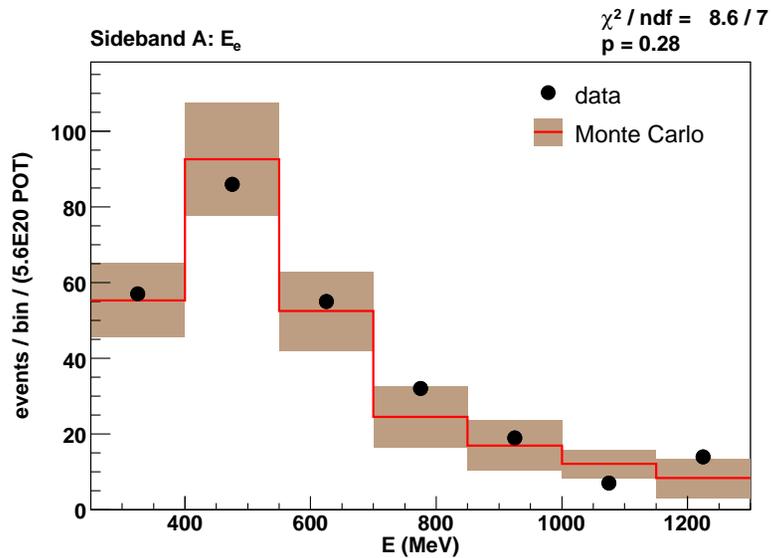


Figure 13.8: Sideband A plots (Part 2).

Sideband B plots

Next 2 pages

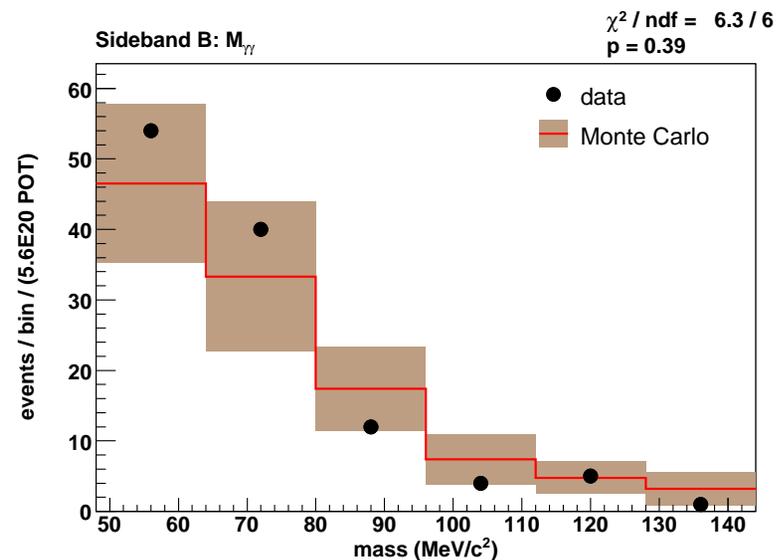
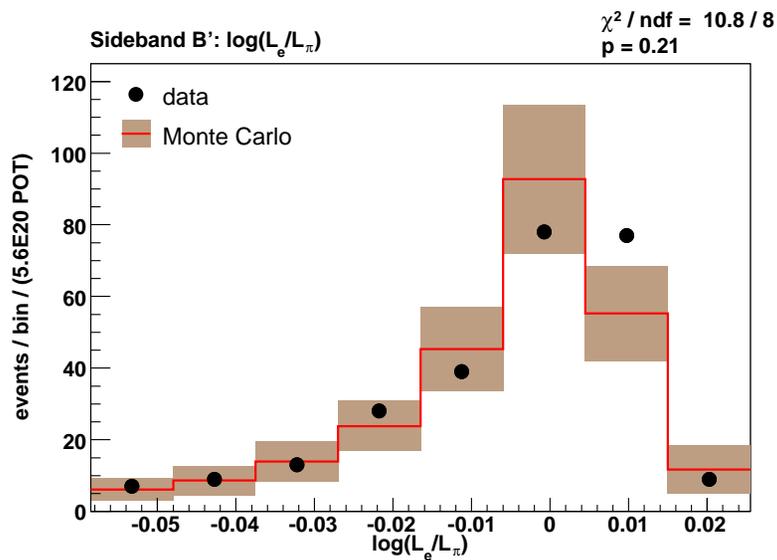
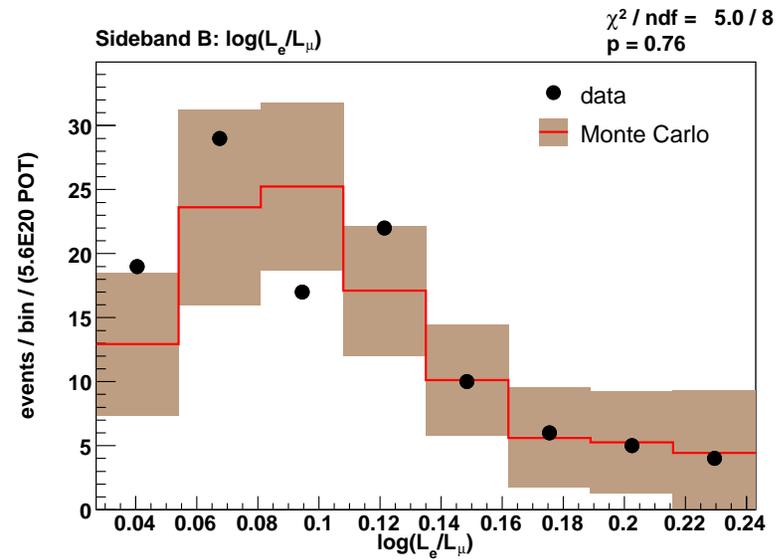
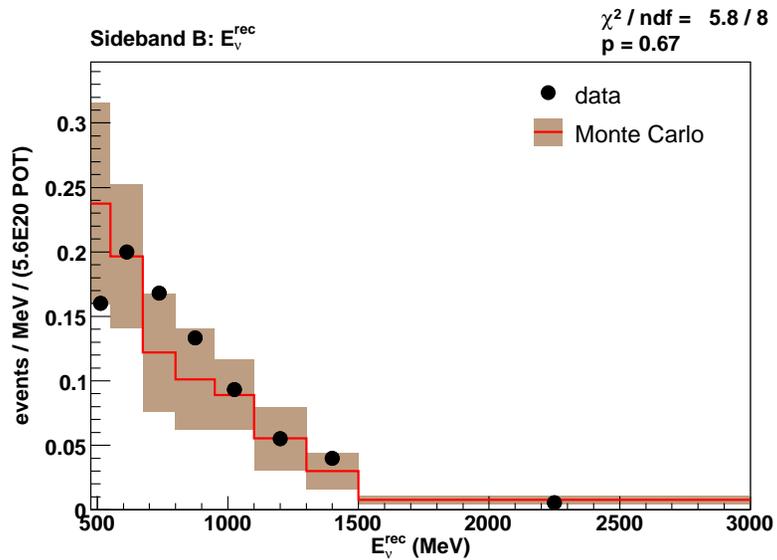


Figure 13.9: Sideband B plots (Part 1).

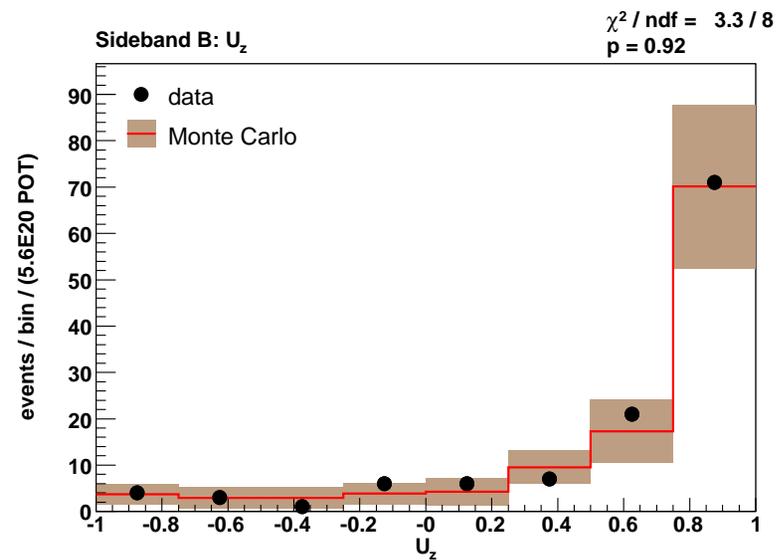
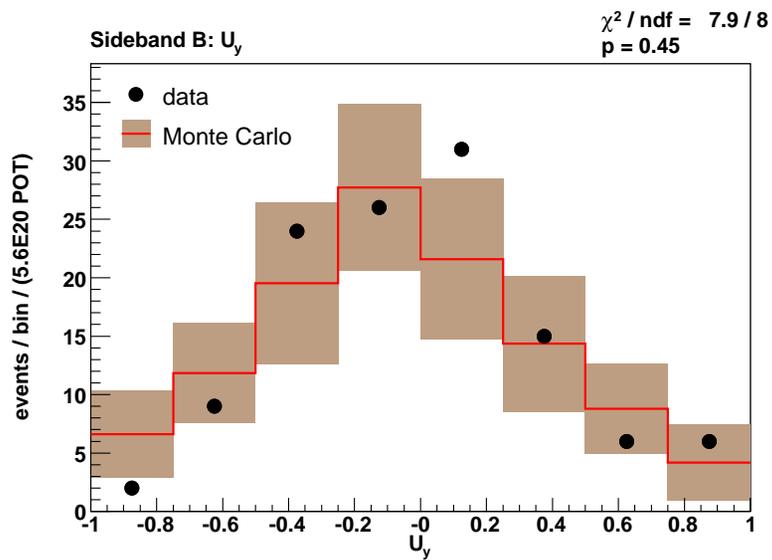
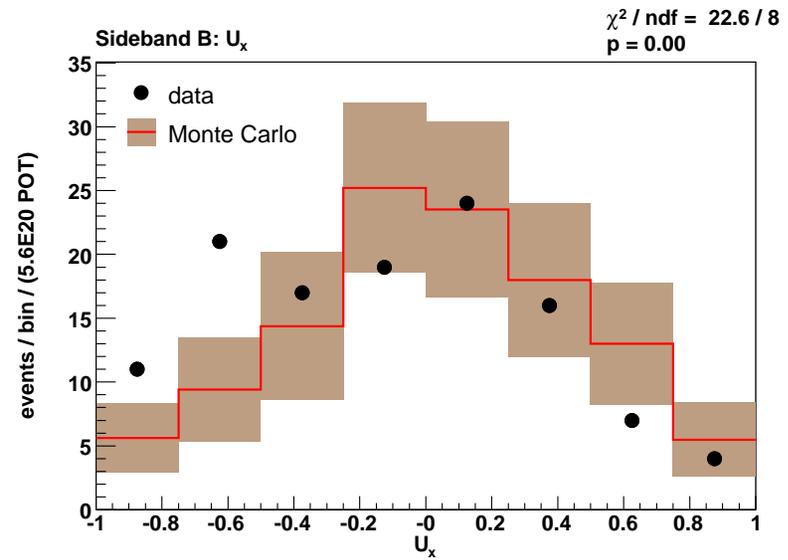
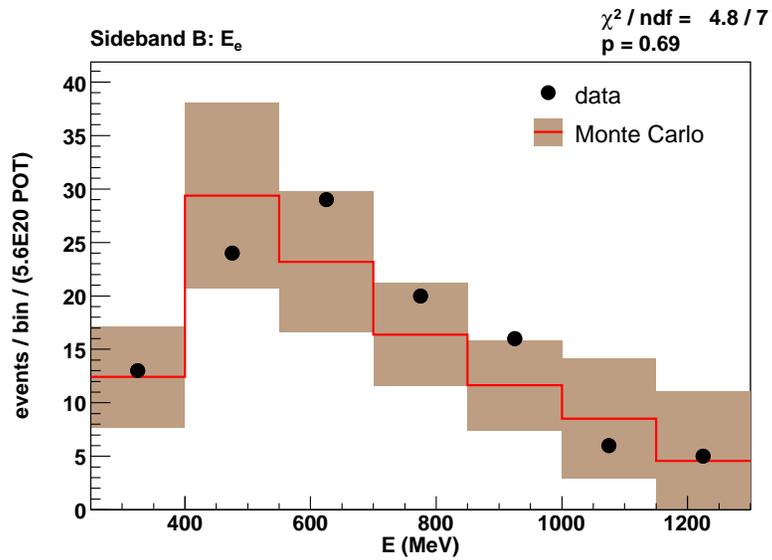


Figure 13.10: Sideband B plots (Part 2).

Sideband C plots

Next 4 pages

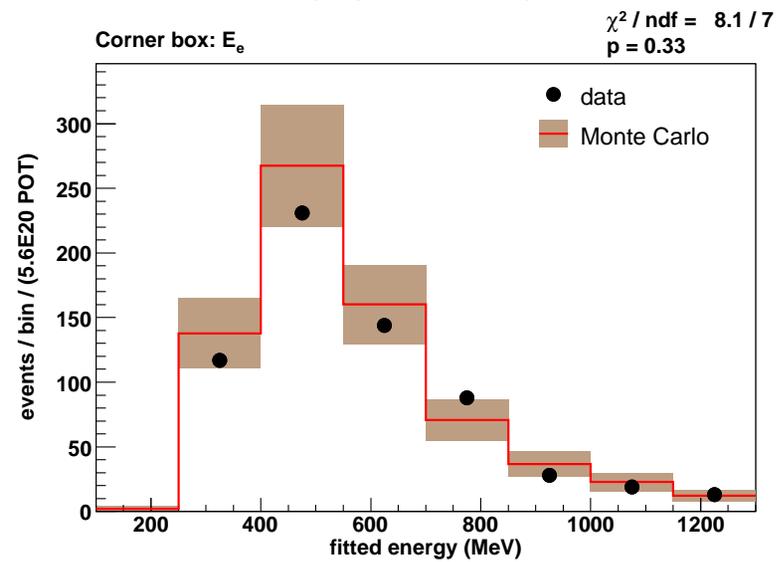
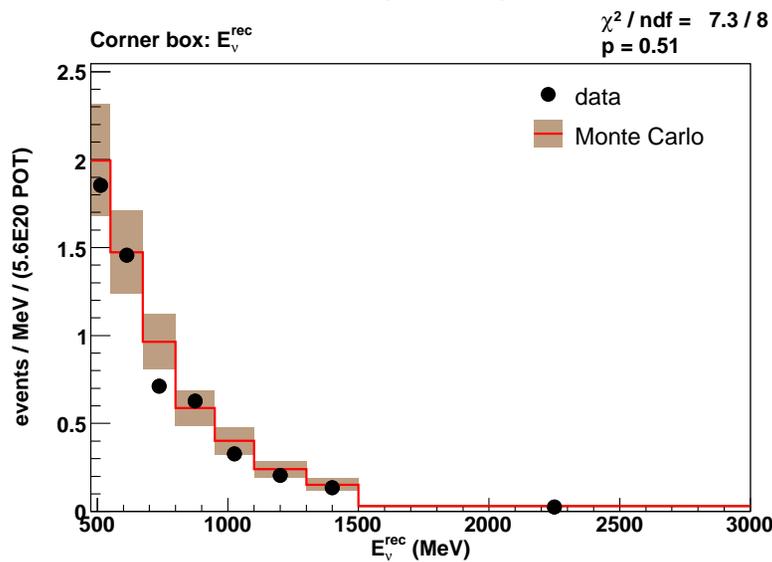
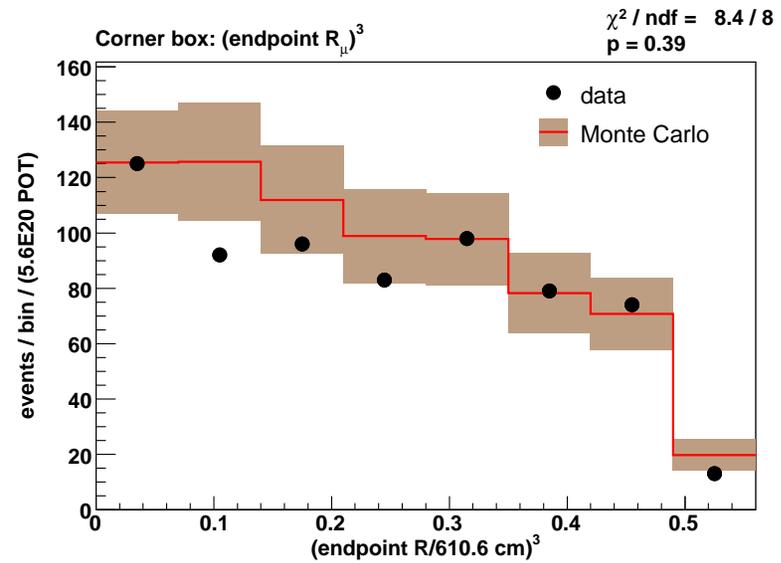
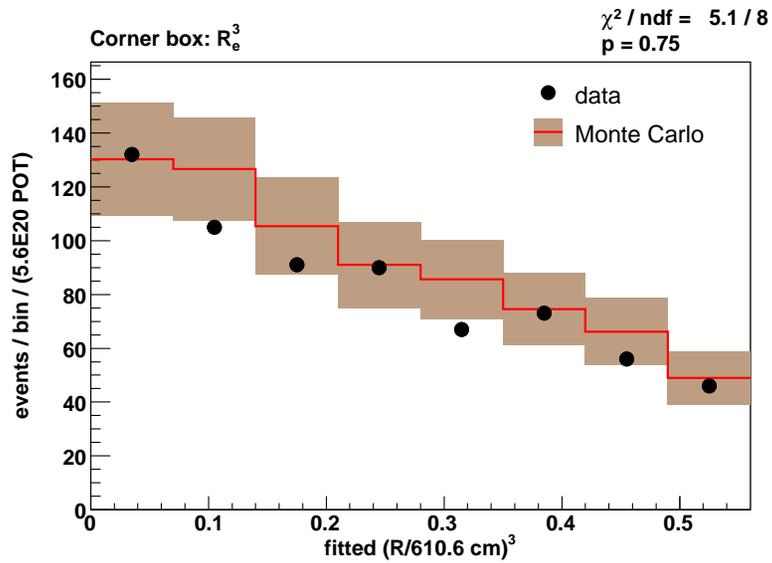


Figure 13.11: Sideband C plots (Part 1).

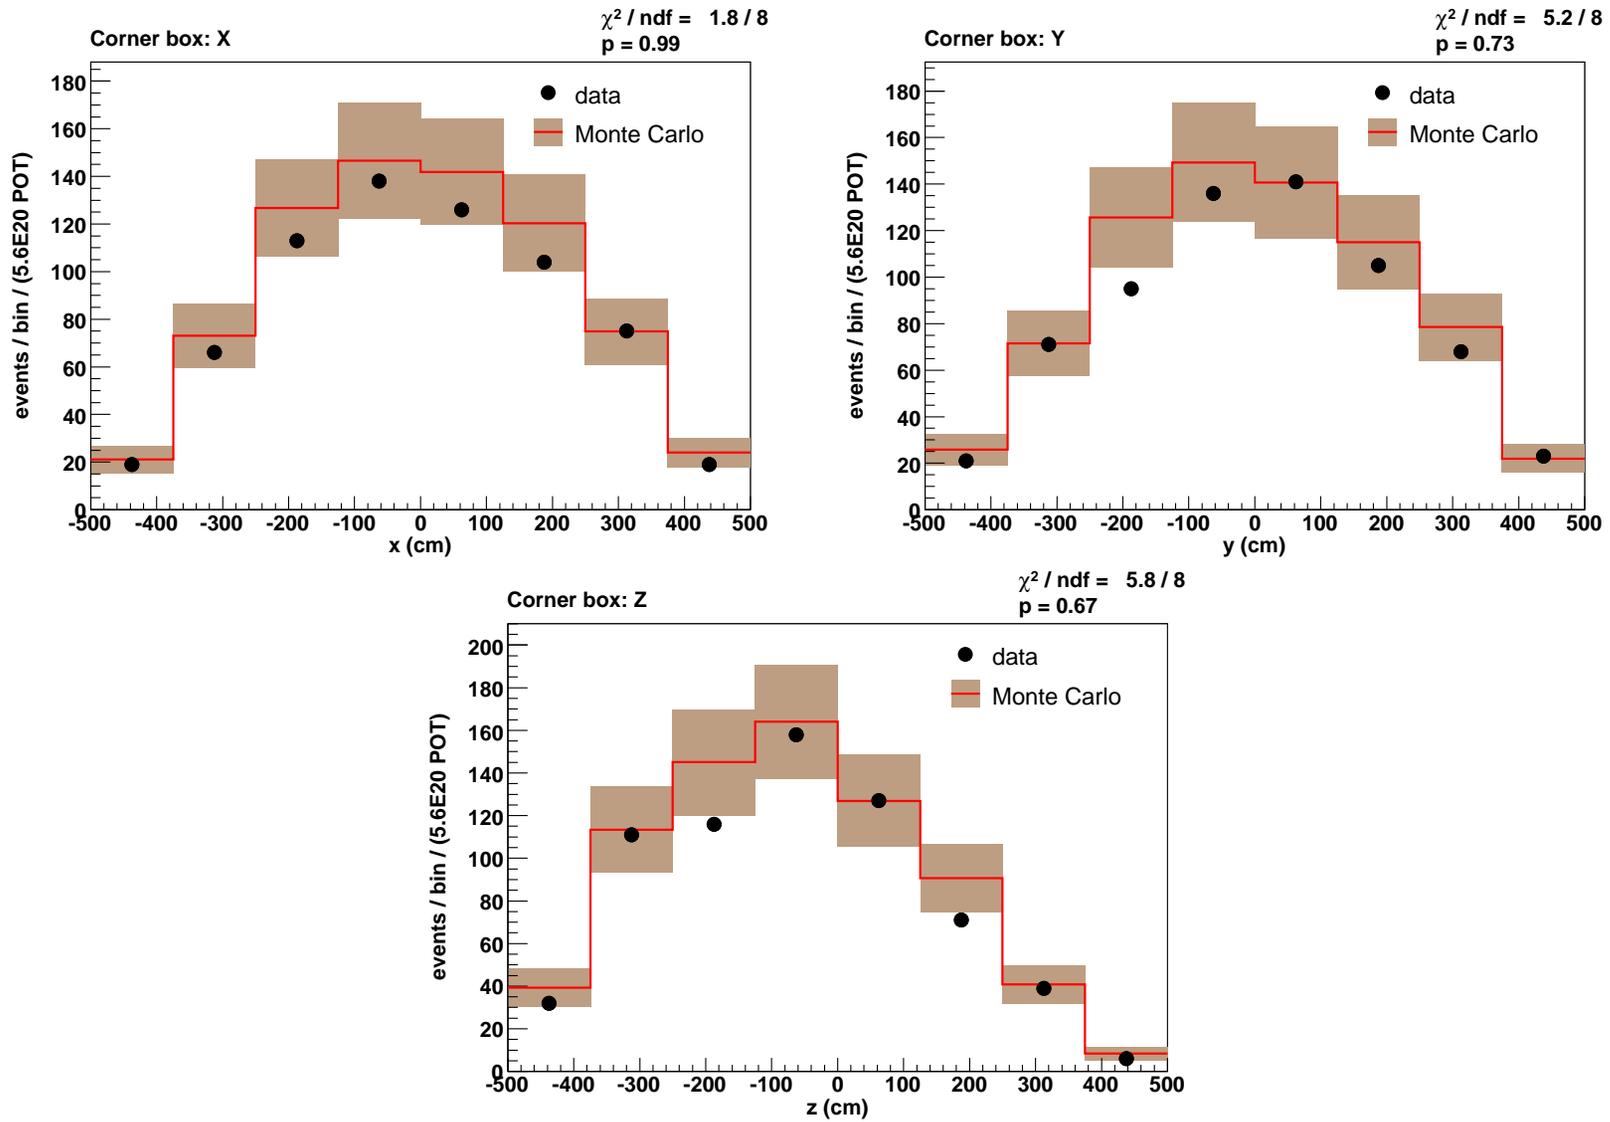


Figure 13.12: Sideband C plots (Part 2).

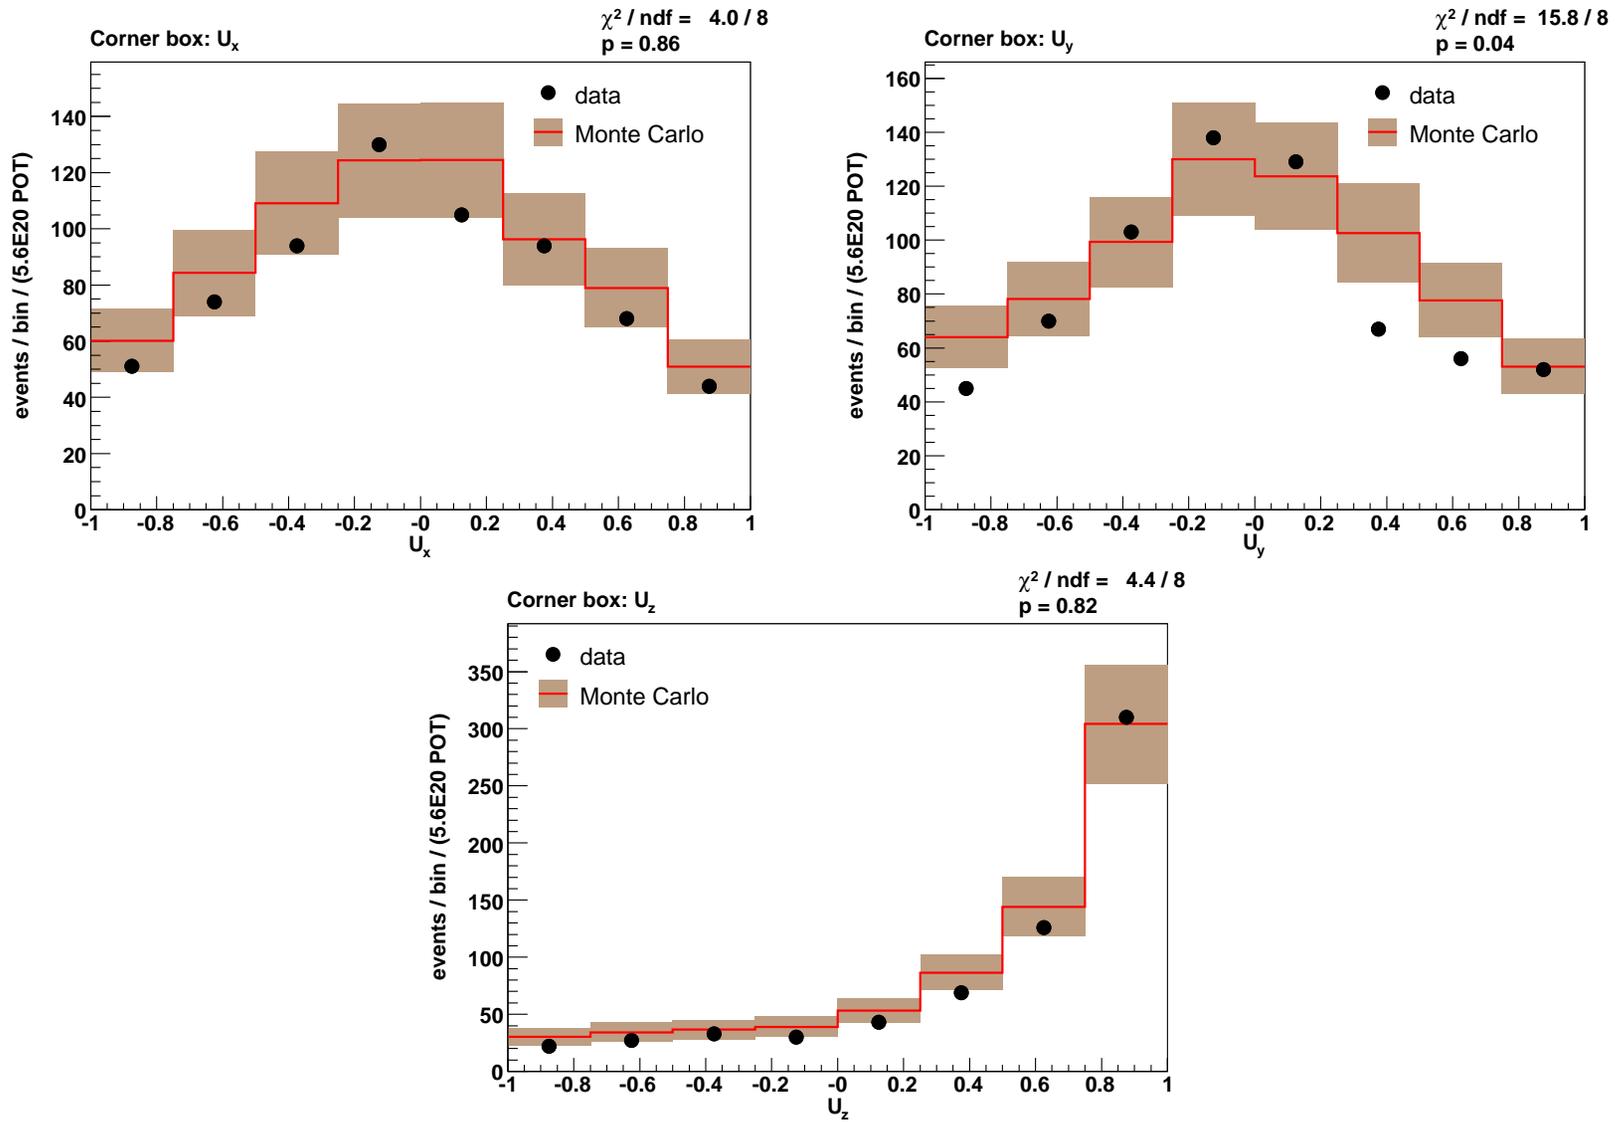


Figure 13.13: Sideband C plots (Part 3).

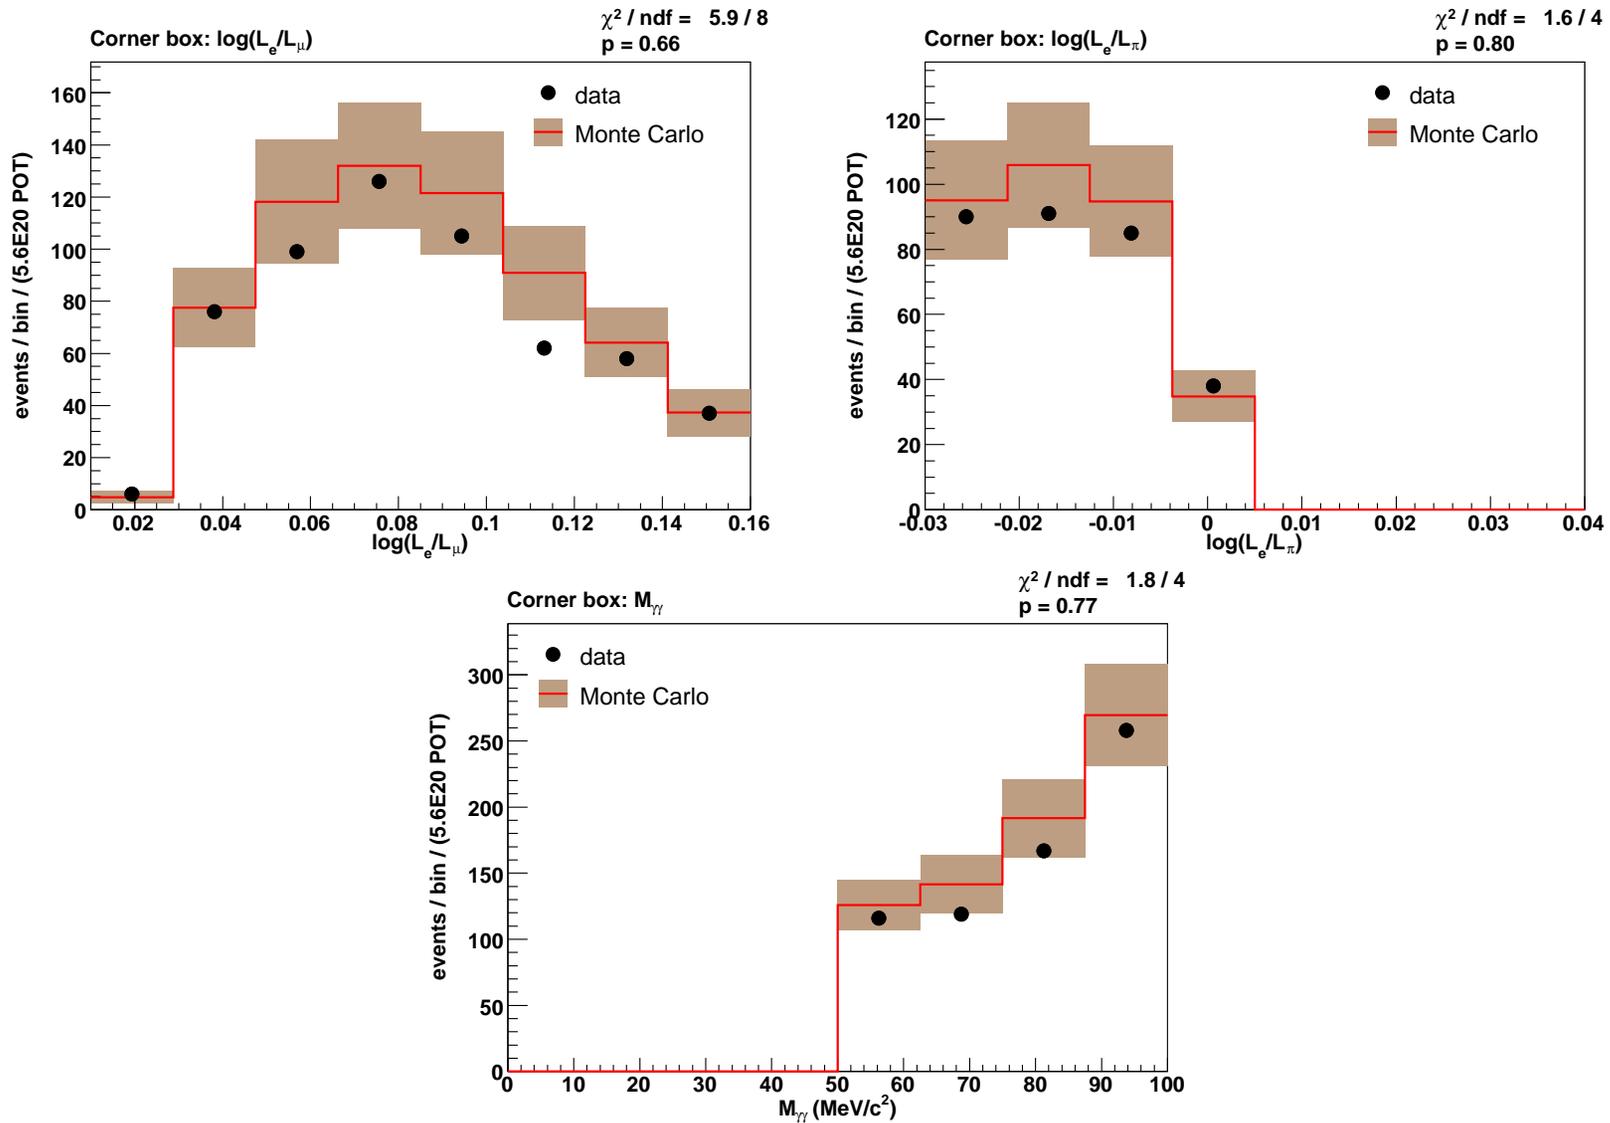


Figure 13.14: Sideband C plots (Part 4).

13.2 Lifting blindness

We can perform similar checks with the ν_e sample, but we must consider the possibility of an oscillation signal. We find the best-fit oscillation parameter set α_{best} by minimizing the $E_\nu^{\text{rec}} \chi^2$, per §12.4.1. However, *neither the minimum χ^2 nor the parameter values are revealed*. Rather, we use the unseen best-fit parameters (which may correspond to a small, large, or non-existent signal) to make data/simulation comparisons, like those above, with the candidate ν_e events. If our simulation correctly predicts key distributions – accounting for any oscillation signal via α_{best} – then we can have confidence continuing the unblinding process.²

Table 13.2 shows the χ^2 values obtained with an $E_\nu^{\text{rec}} = 300$ MeV energy threshold. While most of the distributions show good agreement, E has a cumulative χ^2 probability of only 1%. Since E and E_ν^{rec} are closely related, this low probability suggests that the yet-to-be-seen χ^2 from the signal fit might also be poor.³

quantity	χ^2	$P(\chi^2+)$
R^3	2.5	0.96
(endpoint R_μ) ³	2.7	0.95
E	20.2	0.01
X	3.8	0.87
Y	5.4	0.71
Z	5.2	0.74
U_x	5.1	0.74
U_y	5.5	0.71
U_z	5.6	0.70
$\log(\mathcal{L}_e/\mathcal{L}_\mu)$	13.7	0.09
$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$	6.7	0.57
$M_{\gamma\gamma}$	3.3	0.91

Table 13.2: First data/simulation comparisons from the ν_e sample, accounting for the best-fit signal. $N_{\text{dof}} = 8$ for all rows.

After an initial review, nothing amiss was found in background estimates or systematic error assessments. To aid the investigation, the following measure of data/simulation agreement was released for each bin of the E histogram:

$$\delta_i = \frac{|d_i - p_i(\alpha_{\text{best}})|}{\sqrt{V_{ii}}}, \quad (13.1)$$

where d_i is the height of the i -th E -histogram bin in data, p_i is the same for the simulation

²Before performing this blind fit and making the subsequent comparisons, we tested the code on 10% of the available neutrino data. While this 10% sample is statistics limited and, thus, provides minimal analysis validation, it helps verify our ability to extract best-fit parameters and to perform data/simulation comparisons while staying signal-blind.

³Indeed, it was revealed some weeks later: $\chi_{\text{best}}^2/N_{\text{dof}} = 15.5/6$, $p = 1.7\%$.

assuming signal α_{best} , and V_{ii} is the no-signal covariance matrix. δ_i gives the magnitude of the discrepancy in bin i in units of the (no-signal) standard error for that bin. Figure 13.15 shows $\{\delta_i\}$. Because bin-to-bin correlations are important, the individual δ_i 's are of limited use. Applying the full covariance matrix to the discrepancies, however, reveals that bin 2 ($E \sim 300$ MeV) causes the most trouble, increasing χ^2 by 5 – 11 units (recall $N_{\text{dof}}=8$) in the presence of any of the other seven discrepancies. Noting that bin 1 has few events and cannot tell us much⁴, our best lead after this study was only that low energy background estimates might be involved in the data/simulation disagreement. (See Appendix D for the present state of knowledge.)

It was found that raising the E_{ν}^{rec} threshold as high as 500 MeV had little impact on the $\nu_{\mu} \rightarrow \nu_e$ oscillation sensitivity. Figure 13.16 compares the oscillation sensitivities for $E_{\nu}^{\text{rec}} > 300$ MeV and $E_{\nu}^{\text{rec}} > 475$ MeV analysis cuts. The similarity in the contours implies that events below 475 MeV add negligible signal information in the fit. Given this, we increased the analysis threshold from 300 MeV to 475 MeV, and we repeated the blind signal fit and subsequent data/simulation comparisons. Table 13.3 shows the results. Given the good agreement (including E , with $p=28\%$), we proceeded by viewing the signal fit's χ_{best}^2 , where this agreement continued ($\chi_{\text{best}}^2/N_{\text{dof}}=0.6/6$, $p=99\%$). We then lifted blindness completely and viewed the oscillation result.

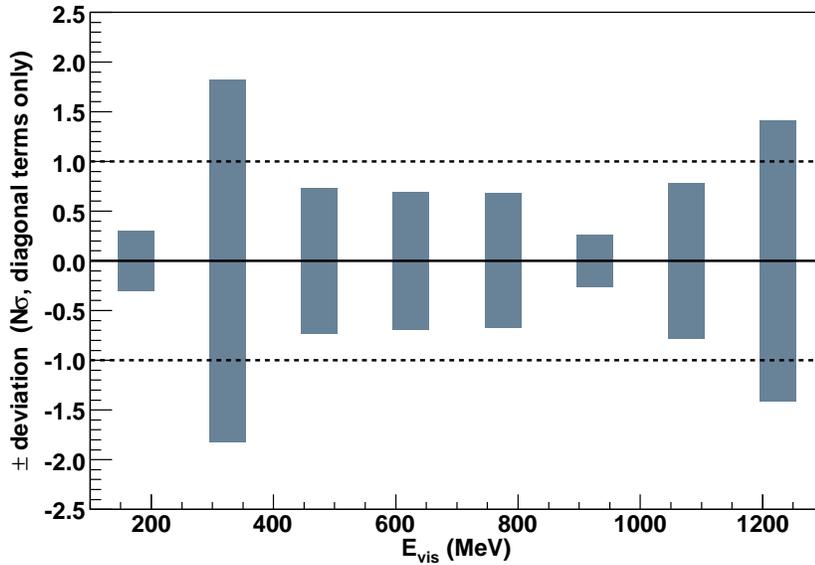


Figure 13.15: A view of the $\{\delta_i\}$. Since the signs of the discrepancies are not known, δ_i is plotted symmetrically about zero.

⁴The E_{ν}^{rec} threshold decimates this first E bin, and its statistical error is therefore large.

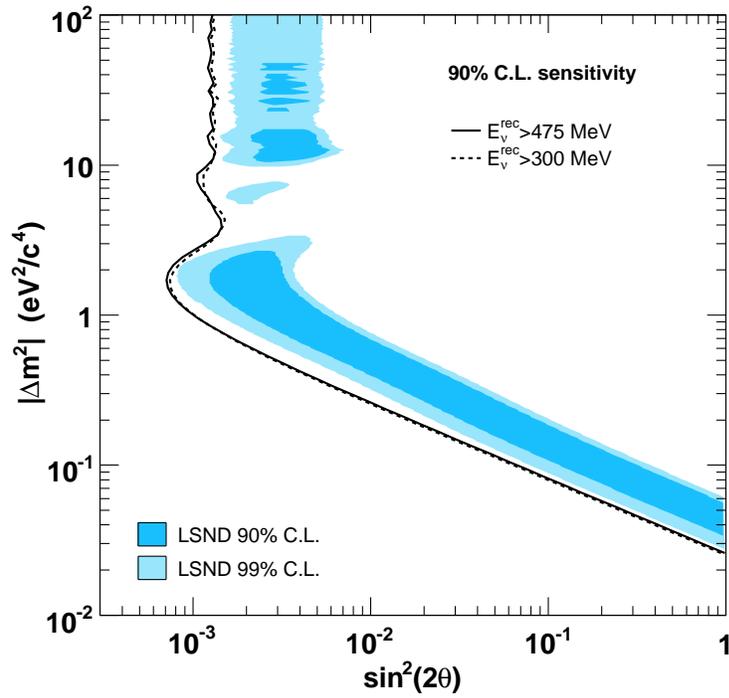


Figure 13.16: Oscillation sensitivity for two different E_{ν}^{rec} thresholds.

quantity	χ^2	$P(\chi^2+)$
R^3	9.0	0.34
(endpoint R_{μ}) ³	6.4	0.60
E	8.7	0.28
X	8.0	0.43
Y	9.0	0.34
Z	2.9	0.94
U_x	8.5	0.38
U_y	6.9	0.54
U_z	7.6	0.47
$\log(\mathcal{L}_e/\mathcal{L}_{\mu})$	7.2	0.51
$\log(\mathcal{L}_e/\mathcal{L}_{\pi^0})$	9.7	0.21
$M_{\gamma\gamma}$	0.8	1.00
E_{ν}^{rec}	0.6	0.99

Table 13.3: Second data/simulation comparisons from the ν_e sample, using a cut of $E_{\nu}^{\text{rec}} > 475$ MeV. The E_{ν}^{rec} fit's χ^2 is also shown.

Chapter 14

Results

Below are the findings from the $\nu_\mu \rightarrow \nu_e$ oscillation search.

14.1 Summary

The number of ν_e candidate events in the $475 \text{ MeV} < E_\nu^{\text{rec}} < 1250 \text{ MeV}$ range is

$$\begin{aligned} \text{observed: } & 380 \\ \text{expected background: } & 358 \pm 19(\text{stat.}) \pm 35(\text{syst.}) \end{aligned}$$

The excess of $(380 - 358) = 22$ events is insignificant (0.55σ). Figure 14.1 shows the E_ν^{rec} distribution of the ν_e candidate events along with the background prediction and an example LSND-sized signal. The data are well-described by the background-only prediction, with $\chi^2/N_{\text{dof}} = 1.93/8$ ($p = 98\%$). Figure 14.2 shows the the best-fit E_ν^{rec} spectrum, along with a breakdown of the expected backgrounds. The fit gives:

$$\begin{aligned} \text{best-fit parameters: } & (0.0011, 4.1 \text{ eV}^2) \\ \chi_{\text{best}}^2/N_{\text{dof}} &: 0.99/6 \\ \Delta\chi^2(\boldsymbol{\alpha}_{\text{null}}) &: 0.83 \mapsto 60\% \text{ C.L.} \end{aligned}$$

The last line says that signal-free experiments will give $\Delta\chi^2(\boldsymbol{\alpha}_{\text{null}}) \leq 0.83$ 60% of the time (§12.4.2). Equivalently, the null hypothesis is rejected at only the 60% confidence level. The best-fit and null parameters, then, provide comparably good descriptions of the data. Table 14.1 discusses the unusually low values of χ^2 .

Figure 14.3 shows the 90% C.L. $\sin^2 2\theta$ upper limit obtained as described in §12.4.3. The LSND 90% and 99% C.L. allowed regions are almost entirely excluded at 90% C.L.

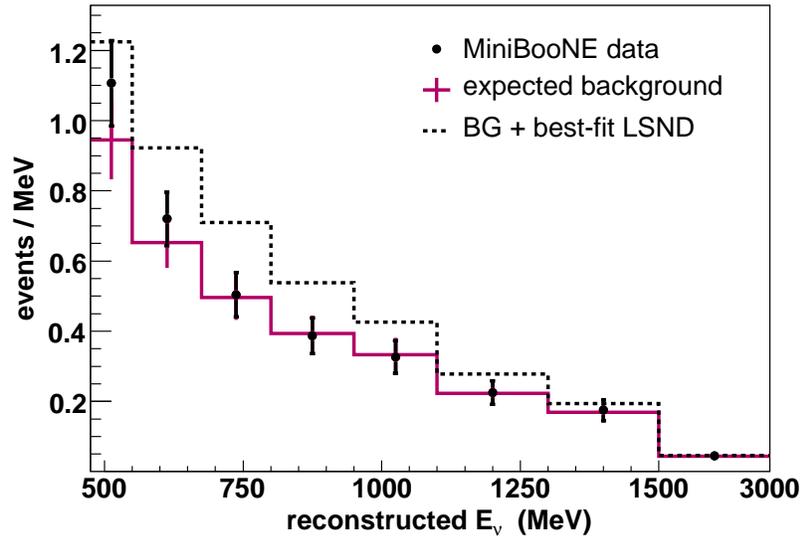


Figure 14.1: E_ν^{rec} histograms for data (black points with statistical errors), expected background (magenta histogram with systematic error band), and an example signal (dashed black histogram).

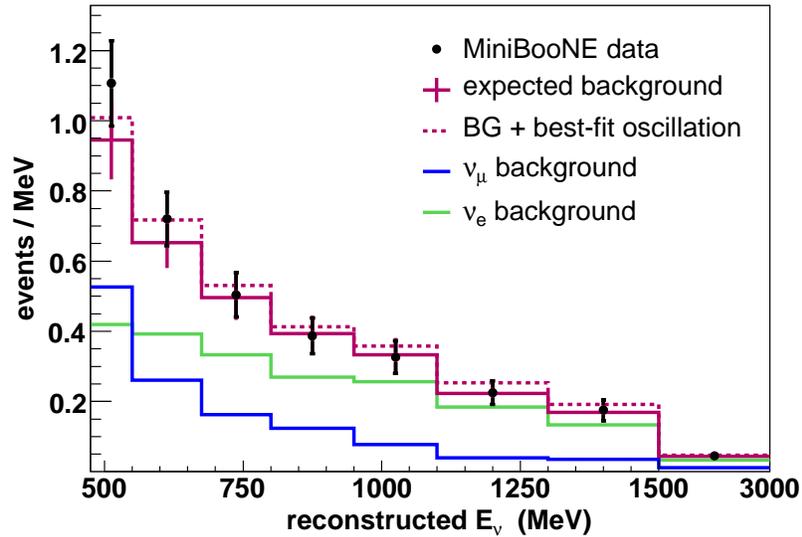


Figure 14.2: E_ν^{rec} spectra, including data and the best-fit prediction. The misidentification (ν_μ) and intrinsic (ν_e) background contributions are also shown.

low edge (MeV)	data	predicted background	(data-pred.)	statistical error	χ^2 contribution
475	83	70.91	-12.09	8.42	2.06
550	90	81.51	-8.49	9.03	0.88
675	63	61.94	-1.06	7.87	0.02
800	58	58.95	0.95	7.68	0.02
950	49	50.01	1.01	7.07	0.02
1100	45	44.49	-0.51	6.67	0.01
1300	35	33.73	-1.27	5.81	0.05
1500	67	65.27	-1.73	8.08	0.05

Table 14.1: A look at the low χ^2 values. In each bin of E_ν^{rec} , we list (4th column) the difference between the observed and predicted event counts. We also compute the statistical error for each bin from the prediction directly: $\text{error} = \sqrt{(\text{predicted background})}$. Comparing these columns reveals that the data fall unexpectedly close to the predictions in six of the eight bins. The last column gives each bin's contribution to χ^2 , assuming these statistical errors only. (The contribution is just $[(\text{data} - \text{pred.})/\text{error}]^2$.) These values should typically be ~ 1 (or larger since we have ignored systematic errors), yet most are 20–100 times smaller. This implies that the low χ^2 values result from a statistical fluke rather than, say, from overestimated systematic uncertainties.

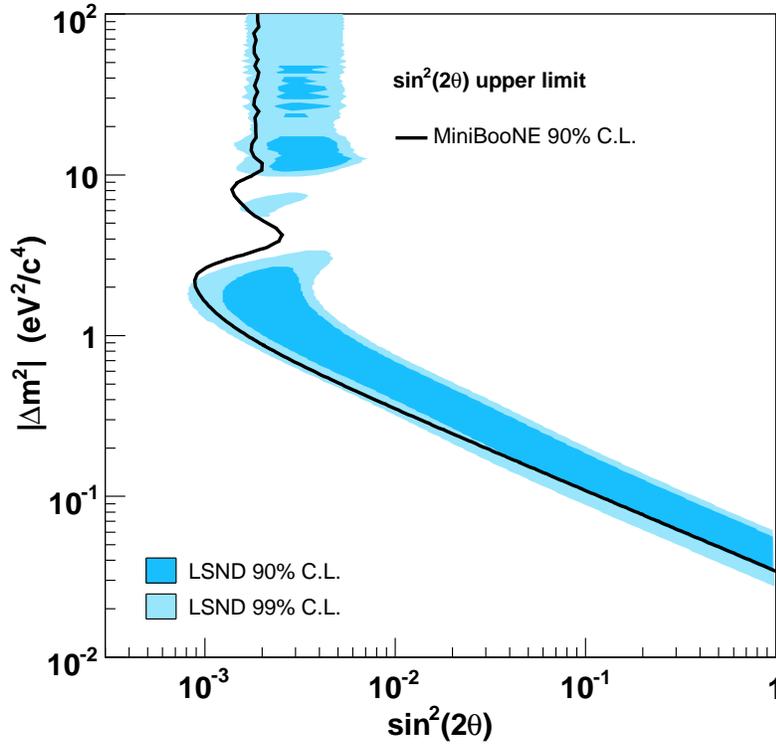


Figure 14.3: 90% upper limit on $\sin^2 2\theta$ as a function of Δm^2 .

14.2 Additional discussion

In Figure 14.4, we show a collection of relevant exclusion curves overlaid on the LSND allowed regions:

- **MiniBooNE 90% C.L.** The official MiniBooNE limit and the result presented here.
- **MiniBooNE 90% C.L. sensitivity.** The sensitivity curve lies to the left of the established limit curve due to the slight upward fluctuation in the data sample. A downward fluctuation would result in the converse arrangement.
- **BDT analysis 90% C.L.** This contour shows the result of a largely independent MiniBooNE analysis. The BDT analysis had different algorithms for reconstruction, particle identification, and signal extraction and used a different scheme for incorporating ν_μ CC QE information. Particle identification was based on a 172-variable “boosted” decision tree, a machine learning algorithm [140]. Systematic errors were dominated by detector model uncertainties. The long list of differences makes the BDT analysis a valuable cross check of the analysis presented in this thesis.¹
- **KARMEN2 90% C.L.** The KARMEN result discussed in §1.3.5.
- **Bugey 90% C.L.** The Bugey result discussed in §1.3.5.

Note that KARMEN reports a global scan contour while Bugey reports a $\sin^2 2\theta$ upper limit contour. The MiniBooNE curves in Figure 14.4 are of the latter variety. Figure 14.5 presents the global scan contour for our analysis, using $\Delta\chi_c^2 = 4.61$. Figure 14.6 compares the global scan contour to one obtained using the approximate cutoff surface of Figure 12.7. The similarity in these contours is our primary justification for using a constant $\Delta\chi^2$ cutoff. Figure 14.7 shows our exclusion curves alongside the 90% C.L. allowed region resulting from a 2002 LSND-KARMEN joint analysis [141]. The global scan contour is more meaningfully compared to the LSND-KARMEN allowed region.

Figure 14.8 shows the behavior of the E_ν^{rec} spectrum below the 475 MeV threshold. The observed excess at low energies, which is responsible for the poor energy χ^2 of §13.2, does not follow a two-neutrino $\nu_\mu \rightarrow \nu_e$ oscillation shape. Appendix D has more on this anomaly.

14.3 Conclusions

The positive oscillation signal reported by the LSND collaboration has sat for some time in need of independent and definitive confirmation or refutation. Verification of the signal would require an entirely new class of particles, the breaking of fundamental symmetries, or more exotic standard model extensions not yet theorized. Rebuttal of the LSND evidence would leave intact a three-generation neutrino oscillation picture consistent with all other observations to date.

¹The choice of which analysis to use for the official result was driven by the relative oscillation sensitivities.

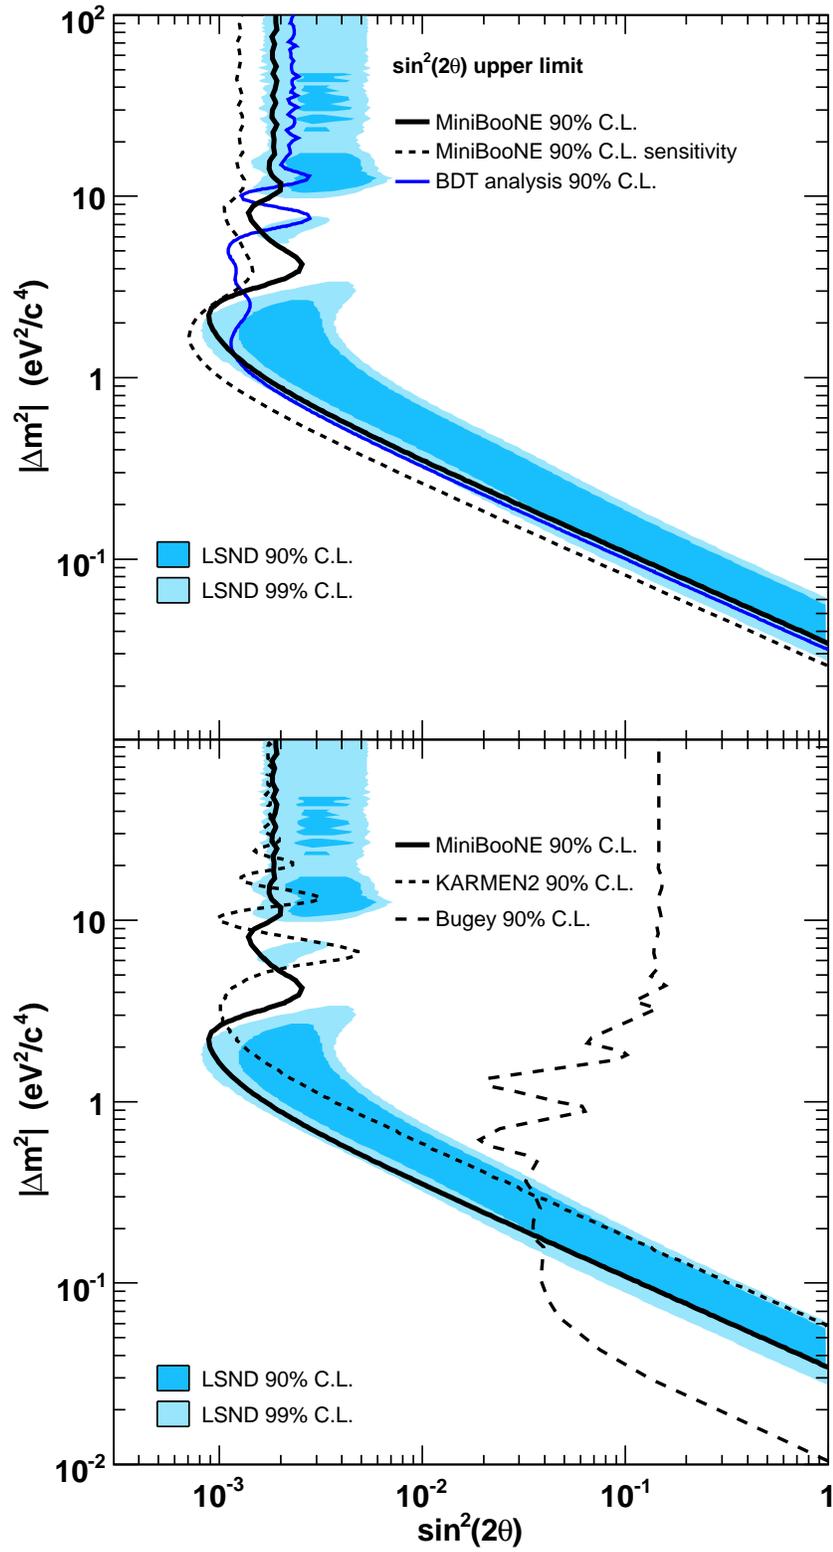


Figure 14.4: Several exclusion curves and one sensitivity curve, as described in the text.

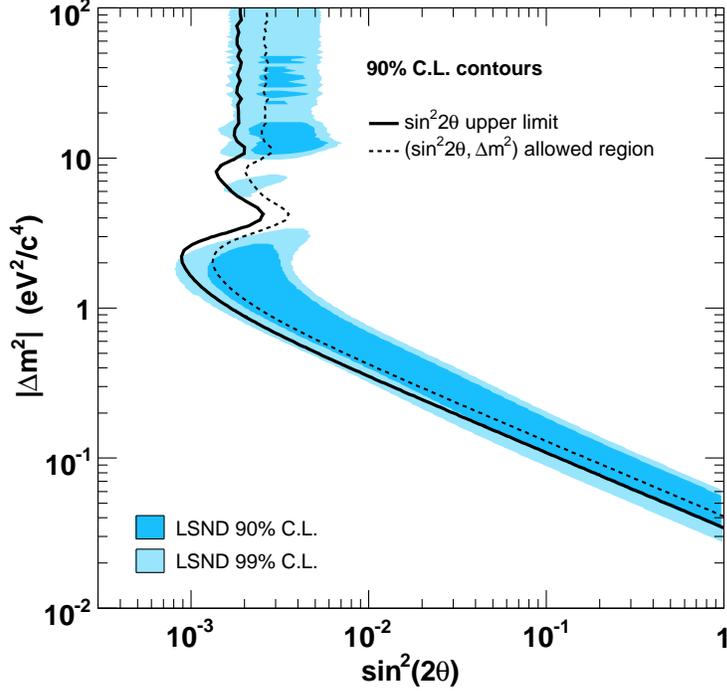


Figure 14.5: The upper limit (solid) and global scan (dashed) 90% contours obtained in this analysis. The former represents the 90% C.L. upper limit on $\sin^2 2\theta$ as a function of Δm^2 . The latter defines the 90% C.L. 2D allowed region for $(\sin^2 2\theta, \Delta m^2)$. See §12.4 for more.

We have presented in this thesis the apparatus, analysis, and results of a $\nu_\mu \rightarrow \nu_e$ oscillation search conducted at FNAL, with sensitivity to an LSND-like signal. Our analysis used 1.7×10^6 neutrino interactions at ~ 1 GeV collected with an 800 ton mineral oil Cherenkov detector. Statistical and systematic errors were comparable, with the latter mostly due to neutrino cross sections and secondary K^+ production.

We find no significant excess of ν_e candidate events, and the reconstructed neutrino energy spectrum above 475 MeV is consistent with expected backgrounds. Our results strongly disfavor a CPT -conserving, two-neutrino oscillation interpretation of the LSND excess. Models that violate CPT or that introduce a rich spectrum of sterile neutrino states are not addressed in this work but are being explored within the MiniBooNE collaboration and elsewhere in the neutrino community. Appendix D has more.

An article related to the work in this thesis can be found in Physical Review Letters [142].

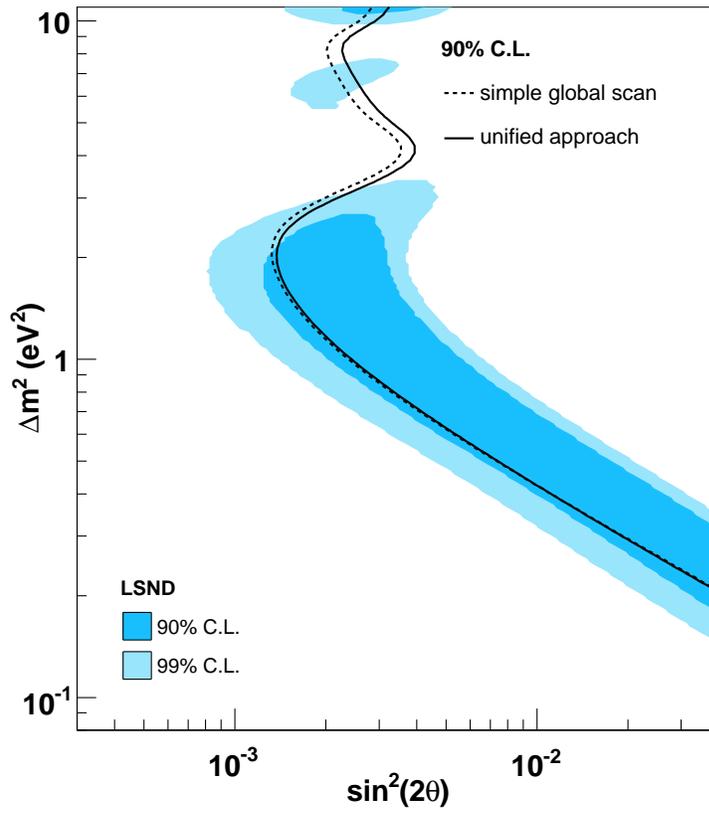


Figure 14.6: Comparing the global scan contour to an approximate “unified approach” contour (per Feldman and Cousins [139]).

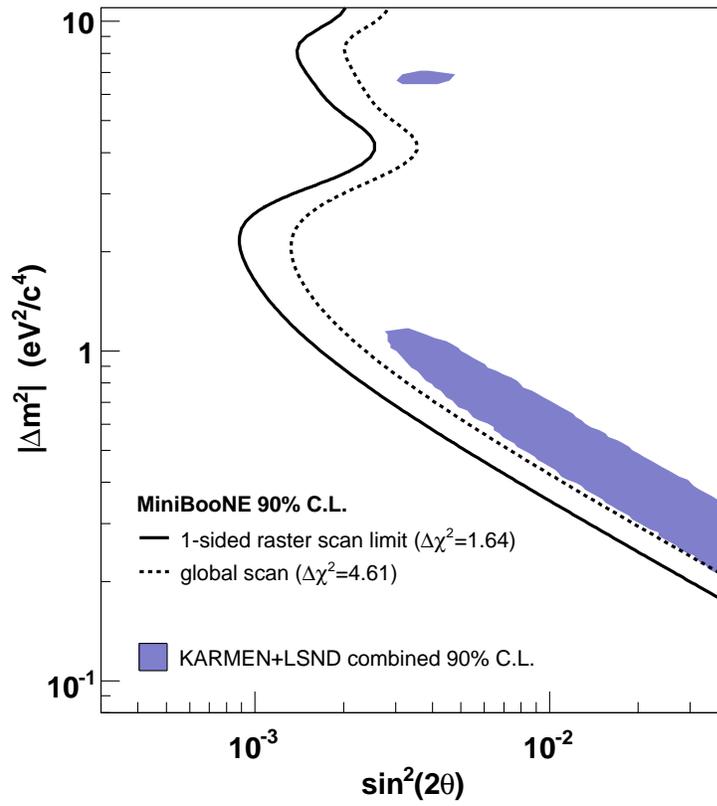


Figure 14.7: 90% C.L. combined LSND-KARMEN allowed region [141] and the results of our analysis.

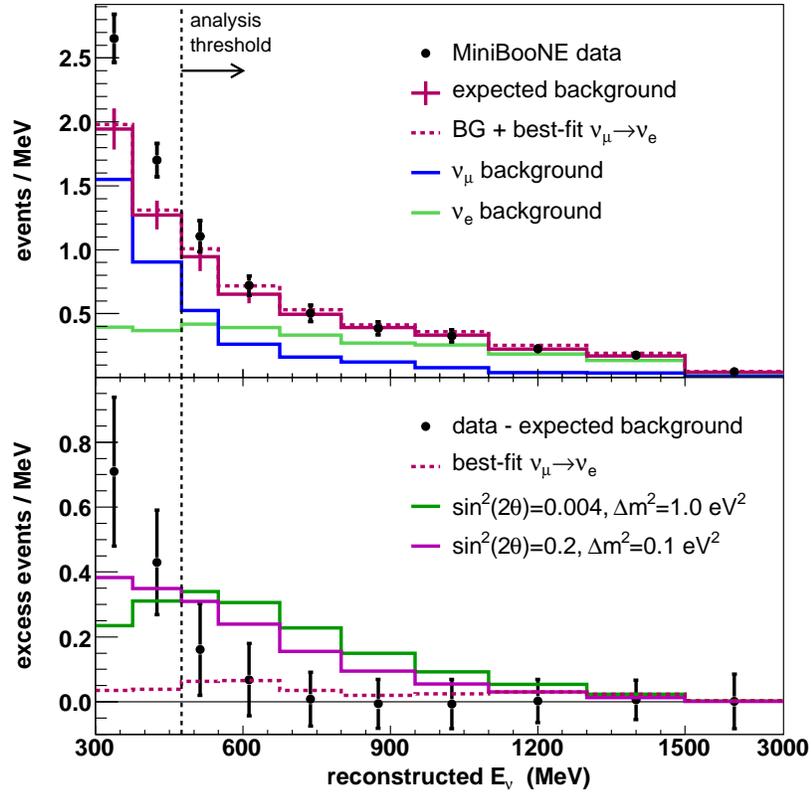


Figure 14.8: (Top) E_ν^{rec} histogram extended below the 475 MeV analysis threshold. The distribution observed in data rises more quickly at low energy than does the predicted one. (Bottom) The excess over background expectation (black points) and example oscillation scenarios. The low- Δm^2 curve (magenta) gives the steepest low- E_ν^{rec} slope obtainable with two-neutrino oscillations. (Still lower Δm^2 values do not increase the slope appreciably.)

Part III

Appendices

Appendix A

Data integrity

In this short appendix, we outline the data integrity checks applied to neutrino events.

- **Data stream merging.** All expected components of the detector and beamline data streams must be present and must have matching GPS time stamps.
- **Beam toroids.** Readings from both toroids in the neutrino beamline – TOR860 and TOR875 – are required, and they must agree to within 10%:

$$2 \times \frac{|\text{TOR860} - \text{TOR875}|}{(\text{TOR860} + \text{TOR875})} < 0.1 . \quad (\text{A.1})$$

- **Horn current.** The horn current must fall between 170 kA and 180 kA.
- **Targeting.** At least 95% of protons must be on course to pass through the length of the target. This is calculated using BPM readings (devices HP875, VP875, HPTGTL, and VPTGTL), beam transport matrices, and a model of emittance growth with beam intensity [143].
- **Latency.** If an event sits in the QT boards' circular buffers for more than 204 μs before getting read out, the event is rejected. (For times longer than this, the event gets written over by new digitizations.) This affects $\sim 0.1\%$ of triggers.

Appendix B

Charge and time miscellany

This appendix details a few topics related to the electronics and the calculation of PMT charge.

B.1 The bleed-off time τ

In various examples in Chapter 4, we use $\tau = 1200$ ns. This section discusses that choice. (Note that we do not assume any particular τ in the analysis, since we construct V_q^{ref} empirically. This section simply supports the arguments made in §4.2.1.)

Figure B.1 shows scope traces taken from a bench-top QT board. The input signal V_{pmt} is not from a PMT; it is from a pulse generator. The V_q trace shows the post-integrator voltage.¹ The clean input signal allows us to see the exponential bleed-off clearly. This particular channel has $\tau = 1120$ ns. The value of τ varies from channel to channel, spanning roughly 1100 ns to 1400 ns.

Interestingly, our laser system allows us to verify the bench-top measurements with our data stream. The laser control box sends a “sync.” pulse to a QT board every time the laser fires. The sync. pulse is a clean, narrow, logic pulse, so the $\{\text{Q_ADC}n\}$ recorded for that channel can be used to extract τ cleanly. We find that the laser sync. channel has a bleed-off time $\tau = 1345$ ns, as shown in Figure B.2.

B.2 Criteria for using a $\text{Q_ADC}n$ in the Q_{raw} fit

Q_{raw} is calculated using a fit to the $\{\text{Q_ADC}n'\}$, with the fit result given by Eq. (4.11). In this section, we list the reasons a $\text{Q_ADC}n'$ may be excluded from a fit.

- $\text{Q_ADC}1$ is never included in a fit. The primary reason for this, as described in §B.3, is that $\text{Q_ADC}1$ often contains no usable charge information at all. Additionally, $\text{Q_ADC}1$ samples a particularly volatile region of the V_q curve – 10 ns of baseline, 40 ns of a sharply rising edge, 50 ns of fast fluctuations (see Figure 4.4) – and discarding it wholesale avoids complications these features would bring.

¹More specifically, this trace shows the voltage at pin 3 of the AD829 op-amp that feeds the charge ADC.

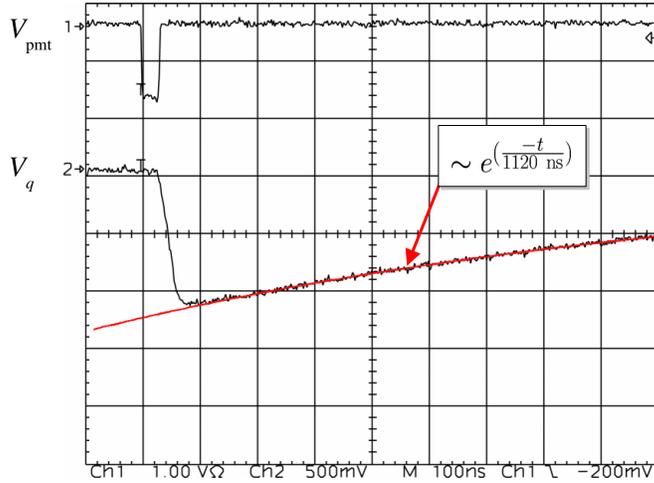


Figure B.1: V_q for a simple pulse. The input pulse V_{pmt} is from a pulse generator. The V_q trace shows the voltage just before the charge ADC and the op-amp that feeds it. (The voltage *after* the op-amp is discussed in §B.3.) The bleed-off time for this channel is measured as $\tau = 1120$ ns. The superimposed curve is proportional to $e^{\left(\frac{-t}{1120 \text{ ns}}\right)}$.

- Naturally, saturated hits have some saturated values: $\text{Q_ADC}n = 0$. These are excluded.
- For saturated hits, we exclude $n = 3$. (Note that $n = 1, 2$ are already excluded for saturated hits by the above cuts.)
- The DAQ occasionally and inexplicably sets $\text{Q_ADC}3 = 255$ for the first (and sometimes the second) quad in a saturated hit. This phenomenon is not understood, mostly because it is so rare. However, evidence suggests that these are simply spurious digitizations. Indeed, no healthy hit can have $\text{Q_ADC}3 = 255$. We do not actually exclude these values in the code; we mention it for documentation's sake.

In summary, non-saturated hits use the values at $n = 2, 3$. Saturated hits use any good values at $n > 3$.

B.3 Op-amp slew rate limitation

The charge ADC does not directly digitize the voltage across the integrating capacitor. Instead, it digitizes an amplified and baseline-shifted version of this voltage. Part of the amplification is provided by an AD829 op-amp configured as shown in Figure B.3. Ideally, the feedback voltage at point *B* in the figure is always equal to the input voltage at point *A*. However, no op-amp can maintain this equality for arbitrarily fast input signals. We show in this section that hits with $q \gtrsim 3$ PE produce V_q pulses that rise too quickly for this configuration of the AD829.

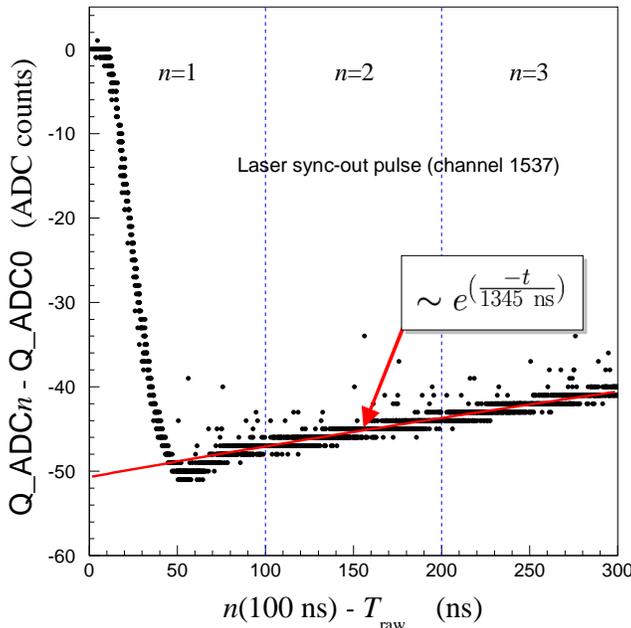


Figure B.2: Bleed-off time for the laser sync. channel. The laser sync. pulse (and, consequently, its V_q) is expected to have a consistent shape and size. Further, the laser fires asynchronously with the 10 MHz clock, so over many events the channel will sample the entire V_q curve. A scatter plot of $Q_ADCn' \equiv Q_ADCn - Q_ADC0$ versus $n(100 \text{ ns}) - T_{\text{raw}}$ for $n = \{1, 2, 3\}$ reveals the sampled V_q curve, as shown. The vertical dashed lines separate the contributions from the three n values. The bleed-off time is determined using $\tau = \left\langle \frac{100 \text{ ns}}{\ln(Q_ADC2'/Q_ADC3')} \right\rangle$, where $\langle \dots \rangle$ indicates the average over many sync. pulses. We find $\tau = 1345 \text{ ns}$. The superimposed curve is proportional to $e^{\frac{-t}{1345 \text{ ns}}}$ (interpreting the abscissa as t).

We first look at the on-board signals from a clean pulse $V_{\text{pmt}}^{\text{fake}}$. Figure B.4 shows the voltages at the input and feedback pins of the op-amp. The left panel of the figure shows that the upstream electronics respond quickly enough, as the rise-time of trace A is the same as the width of the input pulse (as indicated by Δt in the figure). The rise is too fast for the op-amp, though, as evidenced by trace B (which would look identical to trace A in the absence of a slew rate limitation.) We determine in Figure B.5 that the op-amp, as configured, has a maximum slew rate of $\sim 19 \text{ V}/\mu\text{s}$. This number is consistent with the slew rates given in the AD829 data sheet, although that document does not provide a value for our exact configuration [144].

To see when this limit becomes important, we first write the slew rate in terms of ADC counts. Knowing that the 8-bit ADC has a 2-volt range and noting the $\times \frac{15}{16}$ voltage divider between the op-amp and the ADC, we obtain a slew rate of $\sim 230 \text{ ADC counts}/(100 \text{ ns})$. Figure B.6 shows that a hit with charge $q = 1 \text{ PE}$ on a new tube results in a V_q that rises

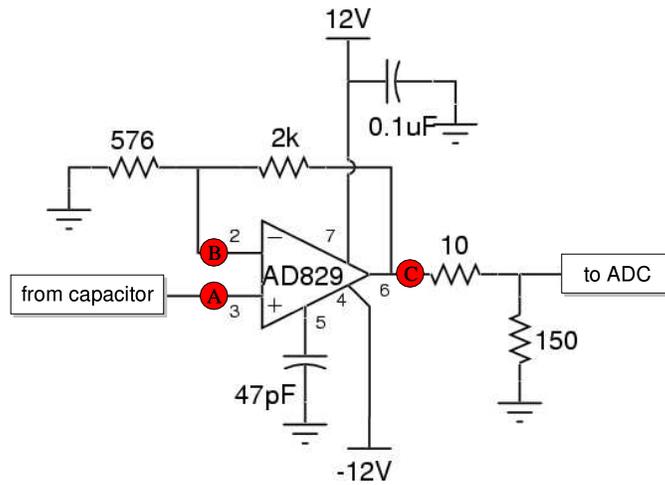


Figure B.3: The V_q op-amp. The “from capacitor” signal is a baseline-shifted version of the voltage across the integrating capacitor. The “to ADC” signal feeds directly into the charge ADC. Other figures in this section show scope traces taken at the points labeled *A*, *B*, and *C*. (Adapted from Ref. [63].)

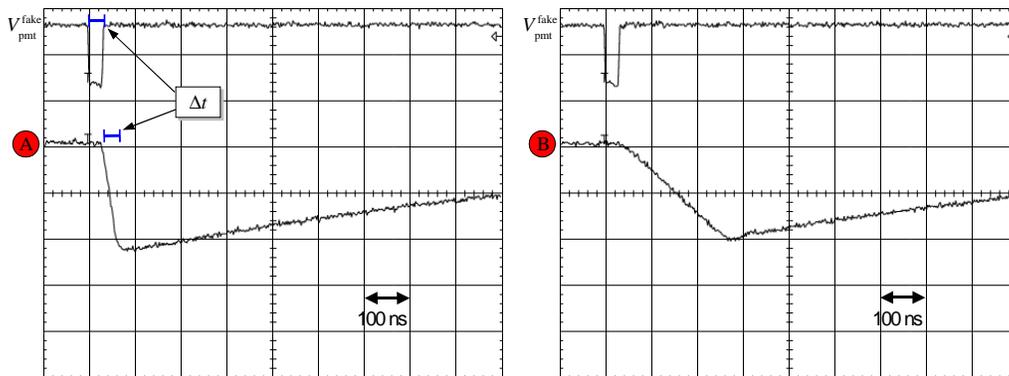


Figure B.4: Slew rate limitation for a sample pulse. $V_{\text{pmt}}^{\text{fake}}$ is a signal from a pulse generator that we have routed into a benchtop QT board. The post-integrator signal, as measured at point *A* in the above circuit diagram, looks as expected, with a risetime comparable to the width of the input signal (indicated by Δt). The signal at point *B* reveals the slew rate of the op-amp. We measure this slew rate in Figure B.5.

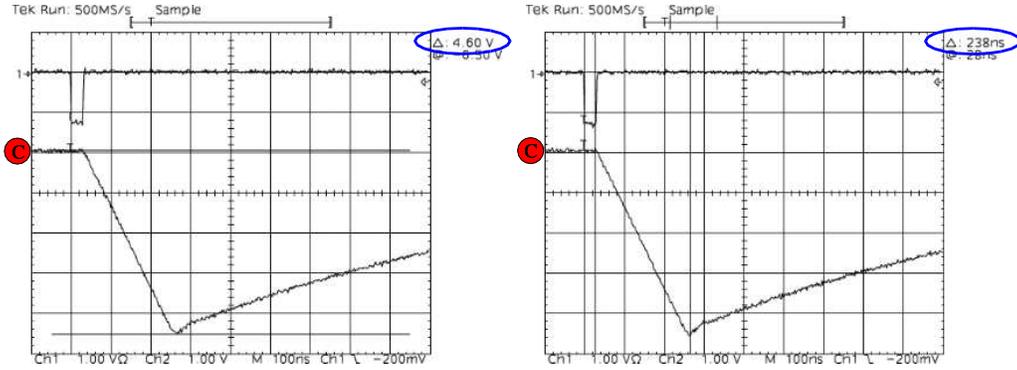


Figure B.5: Measuring the slew rate. By taking traces at point C in the circuit we can directly measure the slew rate. The two plots show the scope cursor measurements (circled in blue) that yield a result of $\sim 19 \text{ V}/\mu\text{s}$.

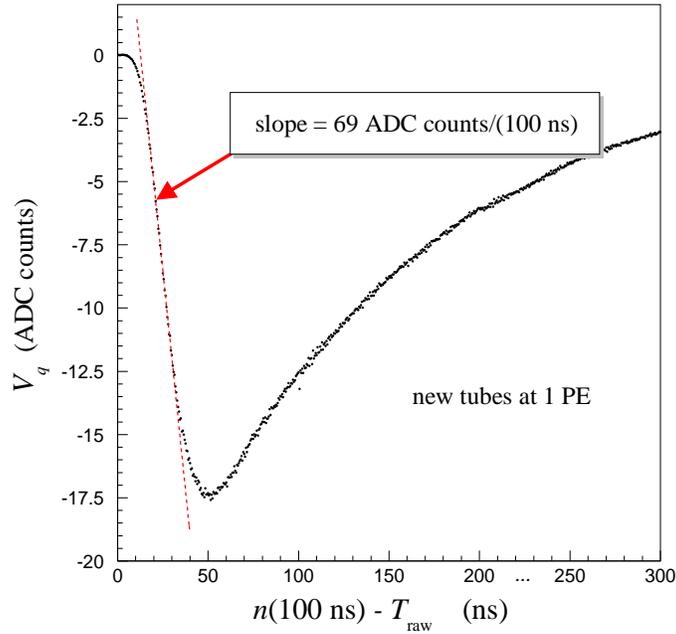


Figure B.6: The rising edge of V_q at 1 PE. The trace shown is the V_q curve for new tubes normalized to 1 PE. The slope of the rising edge reaches $\sim 70 \text{ ADC counts}/(100 \text{ ns})$.

at $\sim 70 \text{ ADC counts}/(100 \text{ ns})$.² Thus, the rising edge of V_q is limited by the slew rate for hits with charge $q \gtrsim \frac{230}{70} \text{ PE} = 3 \text{ PE}$.

The effect is indeed present in the data. One dramatic way to show this is to obtain the average V_q for very high intensity laser events following the procedure described in §4.2.3 for obtaining V_q^{ref} . Figure B.7 shows exactly this for a laser run with hits of mean charge

²For old tubes, this number is $\sim 50 \text{ ADC counts}/(100 \text{ ns})$.

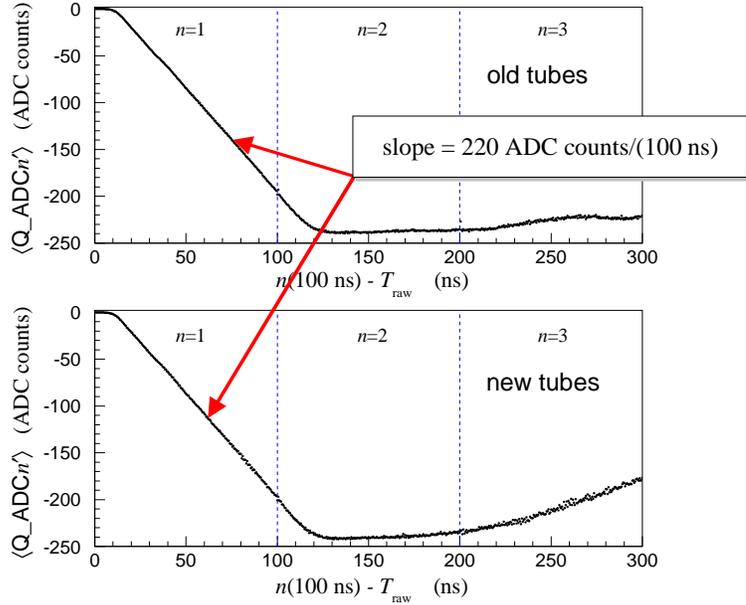


Figure B.7: Seeing the slew rate in data. These traces show the results of the V_q^{ref} -determining procedure when it is applied to high-charge ($\langle q \rangle \approx 20$ PE) hits. The slew rate limitation is apparent.

$\langle q \rangle \approx 20$ PE. We see that the slope of V_q is capped at the expected level and that this level is independent of tube type (*i.e.*, old versus new), as expected.

We conclude this section by noting a few of the ways this slew rate limitation impacts our charge calculations:

- If Q_ADC1 lies on the rising edge of a slope-limited V_q , it provides no useful charge information. For this (and other) reasons, Q_ADC1 is not used in the Q_{raw} fit.
- For high-charge hits, Q_ADC2 can fall on a slope-limited rising edge, as seen in Figure B.7. If Q_{raw} and T_{raw} are in an appropriate (though small) range, this can occur even if the hit is non-saturated. In such cases, only Q_ADC3 has any charge information.
- In principle, the algorithm could be extended to recognize multiple hits on a tube by looking at the fit residuals $\{\Delta_n\}$ defined in §4.2.5. For slew-rate-limited hits, the residual Δ_1 is no longer useful for identifying multiple hits.

B.4 Additional QT board effects

It is worth documenting the following two phenomena observed on the QT boards.

1. Figure B.8 shows a fake PMT signal (from a pulse generator) and the post-integrator voltage measured at point B in the circuit diagram of Figure B.3. As long as the input

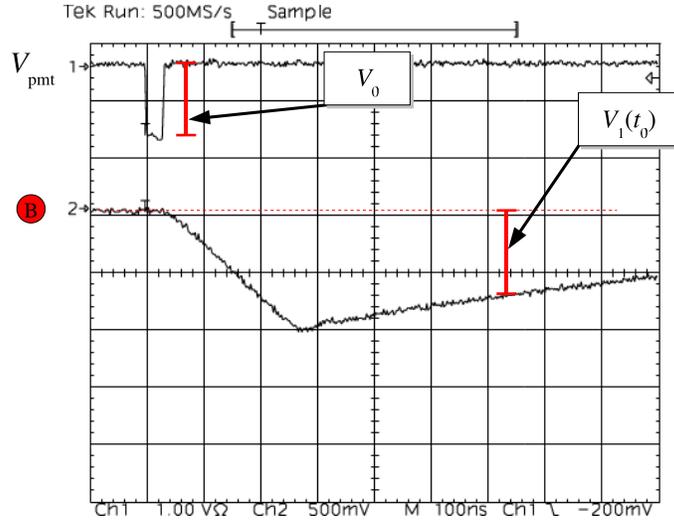


Figure B.8: Proportionality breakdown. V_0 labels the height of the input pulse. $V_1(t_0)$ labels the height of (something proportional to) V_q at a given time t_0 after the pulse arrives. The expected proportionality $V_1(t_0) \propto V_0$ breaks down at large V_0 .

signal's shape does not change, we expect the following proportionality to hold:

$$V_1(t_0) \propto V_0, \quad (\text{B.1})$$

where V_0 is the height of the input pulse and where $V_1(t_0)$ is the voltage at point B some fixed time t_0 after the discriminator fires. We observe a breakdown of this proportionality at large V_0 , with a 30-PE-equivalent pulse yielding output that is $\sim 8\%$ high compared to the small-pulse behavior.

2. At even higher charges, the response voltage $V_1(t_0)$ plateaus. The left panel of Figure B.9 shows our sample input pulse and the voltage $V_1(t_0)$ now measured at point A in the circuit diagram of Figure B.3. In the right panel, we double the input pulse height with no resulting change at A . Note that the traces shown were taken upstream of the op-amp, so this is not a saturation of that device. This hard limit seems to vary channel to channel, falling somewhere between 40-PE-equivalent and 80-PE-equivalent. This represents an upper limit on reconstructible charge imposed by the QT board electronics.

These and other effects contribute to the charge nonlinearities discussed in Chapter 8.

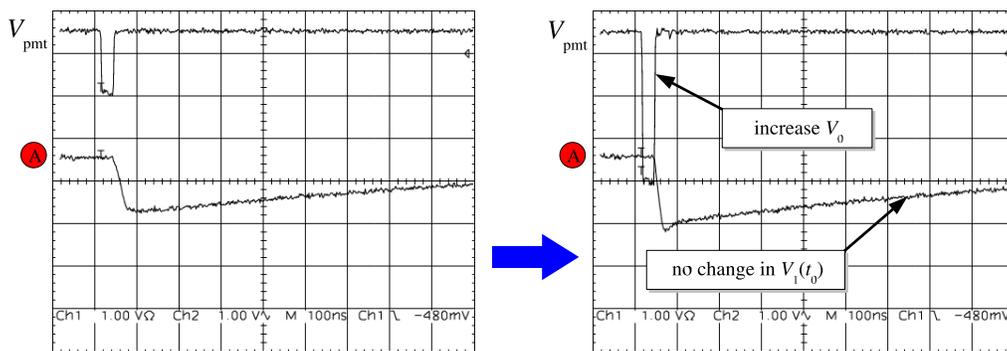


Figure B.9: Hard limit on charge. The left panel shows a sample input pulse V_{pmt} and the voltage measured at the input pin of the charge op-amp. The right panel shows that doubling this input pulse's height does not double $V_1(t_0)$.

Appendix C

Additional reconstruction details

C.1 Solid angle $\Omega(r)$ and acceptance $\epsilon(\eta)$

In measurements performed during the summer of 2004, PMTs were illuminated with an approximately parallel beam of low intensity light, and the rate of photoelectrons was recorded for various incident angles of that light [129]. The angular efficiency function $\epsilon(\eta)$ comes from these measurements.

The distance-dependent PMT solid angle is given by

$$\Omega(r) \equiv 1 - \frac{r+h}{\sqrt{(r+h)^2 + a^2}}, \quad (\text{C.1})$$

where r is the distance to the apex of the PMT,¹ $h = 7.5$ cm is the globe half-height, and $a = 10$ cm is the globe radius. Overall constant factors are taken to be part of the flux Φ .

C.2 Transmission functions $T_{\text{sci}}(r)$ and $T_{\text{Ch}}(r)$

Our detector Monte Carlo simulation (**BooDetMC**) incorporates all of our understanding of the optical properties of the detector. This makes it a convenient tool for establishing tables that involve averages or integrals of underlying optical properties. We describe here how we use **BooDetMC** (v01-11-01) to establish the transmission curves $T_{\text{sci}}(r)$ and $T_{\text{Ch}}(r)$.

We create two **BooDetMC** cardfiles modified from the default as follows:

- All indirect light is turned off. (**RSL1**, **RAM1**, **FLU3**, and **UVF1** are all set to zero and the fifth argument of **OPAR** is set to unity. We set the components of **FLU3** to zero rather than setting **FLU2** to zero because the latter breaks older versions of the **BooDetMC** code.) With these processes off, all extinction becomes absorption.
- Hadronic interactions and particle decays are turned off. (**HADR** and **DCAV** are set to zero. This change should make no significant difference.)
- Only one tank tube is installed. (**PMTR** is set to “1 11” and **PHT1** is reduced to “0.”.)

¹This is the same r used in Chapter 9.

Cherenkov				scintillation			
f_1	0.057	λ_1	18 cm	f_1	0.683	λ_1	44 cm
f_2	0.328	λ_2	349 cm	f_2	0.157	λ_2	3048 cm
f_3	0.603	λ_3	2878 cm	f_3	0.157	λ_3	2907 cm

Table C.1: Parameters for the Cherenkov and scintillation transmission curves shown in Figure C.1. The scintillation transmission is well described with (effectively) only two decay constants.

- The flasks and cubes are moved out of the way. (`GPFx` and `GPCx` are suitably modified.)

One of these cardfiles is further modified to have (essentially) no extinction by setting

```
AL01 1
AL02 0. 10000. 1000000. 0. 0. 0. 0. 0. 0. 0.
```

These two cardfiles are then used to create isotropic Cherenkov and scintillation point sources situated a fixed distance r from the front face of the lone PMT. The sources are “electron bombs”: one hundred simultaneous 3 MeV electrons with isotropically chosen directions. The ratio of the numbers of Cherenkov- or scintillation-induced photoelectrons observed in the two runs (one with extinction, one without) provides the transmission for that distance. Figure C.1 shows the ratio obtained for various distances. Also shown is a fitted parametrization of the form

$$T(r) = f_1 e^{-\frac{r}{\lambda_1}} + f_2 e^{-\frac{r}{\lambda_2}} + f_3 e^{-\frac{r}{\lambda_3}} . \quad (\text{C.2})$$

Table C.1 gives the parameter values. Note that this procedure incorporates all spectral dependences (from production, propagation, and detection) into $T(r)$.

C.3 Scattering tables

§9.2.1 incorporates indirect light into the predicted charge using the scattering tables $A_{\text{sci}}(R, \cos \Theta)$ and $A_{\text{Ch}}(R, \cos \Theta, \cos \theta, \phi)$. As with the transmission curves, we use `BooDetMC` to establish these tables.

We begin with the scintillation scattering table. Our source events are single 3 MeV electrons distributed uniformly in R with isotropically chosen directions. Each PMT in each event corresponds to a particular $(R, \cos \Theta)$ pair. Using the `BooDetMC` output `Ntuples`, we construct a two-dimensional table $Q_{\text{all}}(R, \cos \Theta)$ that holds the average number of scintillation photoelectrons seen by PMTs in each bin of $(R, \cos \Theta)$. We then repeat the exercise with direct light only.² We call this second table $Q_{\text{direct}}(R, \cos \Theta)$. The excess $Q_{\text{all}} - Q_{\text{direct}}$, when normalized by the average direct charge Q_{direct} , is our desired $(R, \cos \Theta)$ -dependent

²§C.2 lists the `BooDetMC` cards that control the indirect light.

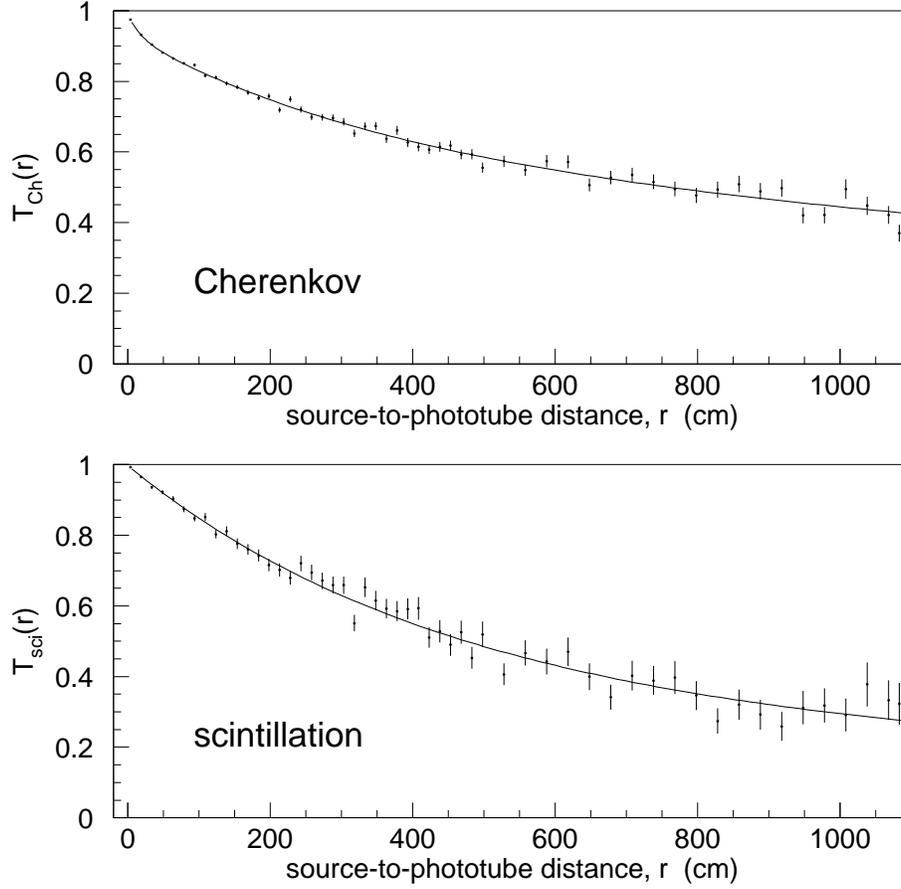


Figure C.1: Transmission curves for Cherenkov and scintillation light. These two panels show the photoelectron ratios (points) described in the text and the fitted three-exponential curve versus the propagation distance r for Cherenkov light (top) and scintillation light (bottom).

function

$$A_{\text{sci}}(R, \cos \Theta) = \frac{Q_{\text{all}}(R, \cos \Theta) - Q_{\text{direct}}(R, \cos \Theta)}{Q_{\text{direct}}(R, \cos \Theta)}. \quad (\text{C.3})$$

Figure C.2 shows the result.

For the Cherenkov table, we create $Q_{\text{direct}}(R, \cos \Theta)$ as above but we also make two tables binned further in $\cos \theta$ and ϕ : $Q_{\text{all}}(R, \cos \Theta, \cos \theta, \phi)$ and $Q_{\text{direct}}(R, \cos \Theta, \cos \theta, \phi)$. The difference $Q_{\text{all}} - Q_{\text{direct}}$ provides the indirect light excess which we normalize to the directionless average charge $Q_{\text{direct}}(R, \cos \Theta)$. Our Cherenkov scattering table, then, is

$$A_{\text{Ch}}(R, \cos \Theta, \cos \theta, \phi) = \frac{Q_{\text{all}}(R, \cos \Theta, \cos \theta, \phi) - Q_{\text{direct}}(R, \cos \Theta, \cos \theta, \phi)}{Q_{\text{direct}}(R, \cos \Theta)}. \quad (\text{C.4})$$

Figure C.3 shows a portion of the resulting table.

Before ending this section, we make rigorous our definition of ϕ . Let \mathbf{u} be a vector

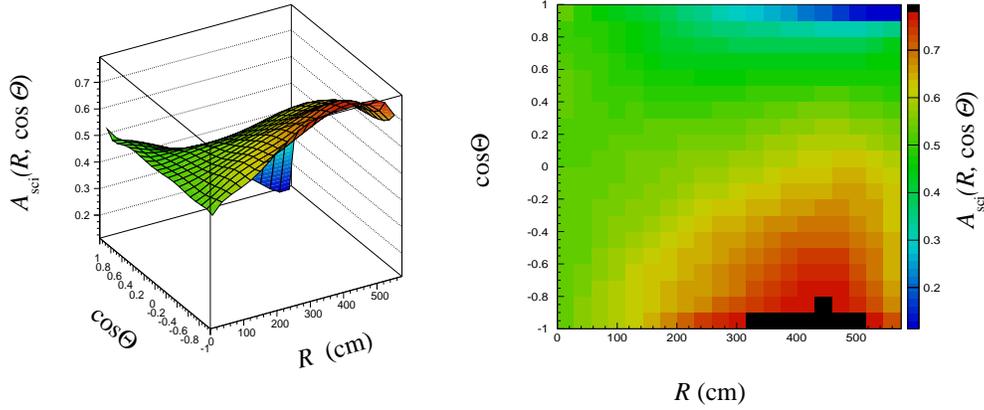


Figure C.2: The scintillation scattering table.

parallel to the track direction and let \mathbf{h} be a vector pointing from the tank center to the PMT. Further, let $\hat{\mathbf{R}}$ be a unit vector pointing from the tank center to the source. The components of \mathbf{u} and \mathbf{h} perpendicular to $\hat{\mathbf{R}}$ are

$$\mathbf{u}_\perp \equiv \mathbf{u} - (\mathbf{u} \cdot \hat{\mathbf{R}}) \hat{\mathbf{R}} \quad (\text{C.5})$$

and

$$\mathbf{h}_\perp \equiv \mathbf{h} - (\mathbf{h} \cdot \hat{\mathbf{R}}) \hat{\mathbf{R}} . \quad (\text{C.6})$$

We set

$$\phi \equiv \cos^{-1}(\mathbf{u}_\perp \cdot \mathbf{h}_\perp) . \quad (\text{C.7})$$

C.4 Production profiles $\rho_{\text{sci}}(s)$, $\rho_{\text{Ch}}(s)$ and $g(\cos \theta; s)$

To create the spatial and angular production profiles, we use a modified version of `BooDetMC` that books and fills `HBOOK` [145] histograms with geometric information about produced photons. We run, at one of various energies, many muon or electron events with

- HADR and DCAY set to zero,
- indirect light turned off (although this should make no difference since we are working with photons at production),
- the flasks and cubes moved out of the way using `GPFx` and `GPCx`, and
- the tank inflated to a radius of 1609.6 via the first argument of `GPAR`.

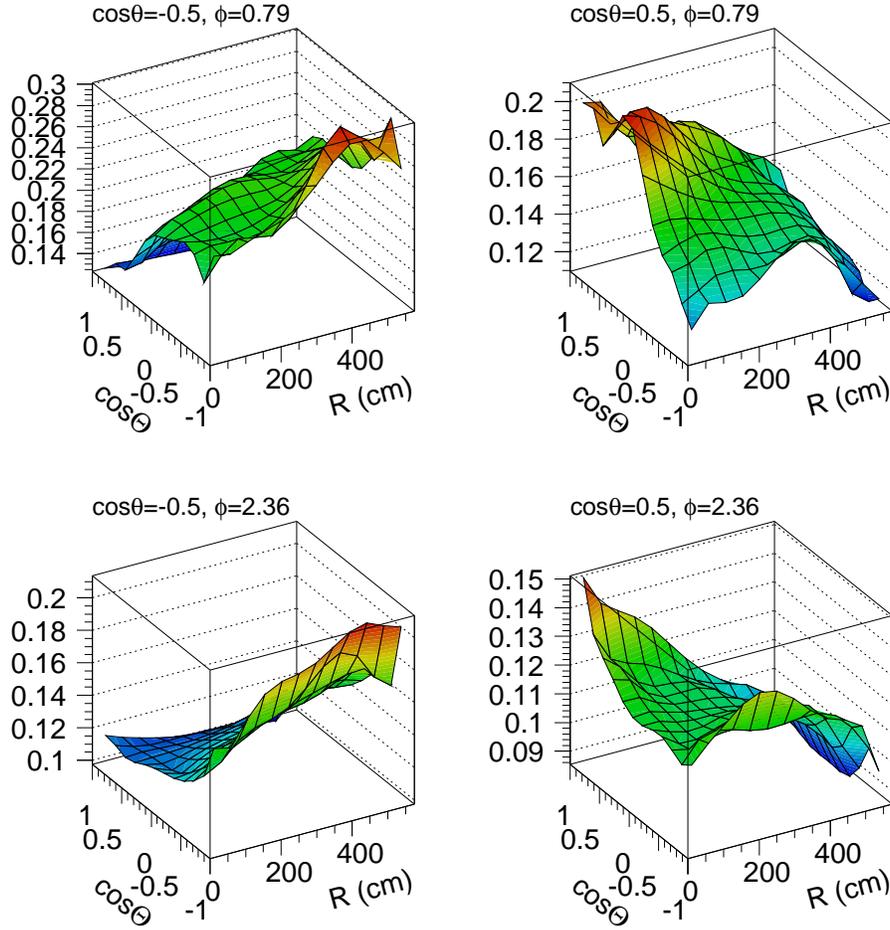


Figure C.3: Part of the Cherenkov scattering table. The R and $\cos \Theta$ dependence of A_{Ch} is shown in four of the one hundred $(\cos \theta, \phi)$ bins.

We start each particle at $(x_0, y_0, z_0) = (0 \text{ cm}, 0 \text{ cm}, -600 \text{ cm})$ with momentum parallel to the $+z$ axis.

The spatial distributions $\rho_{\text{sci}}(s; E_0)$ and $\rho_{\text{Ch}}(s; E_0)$ for a given energy E_0 come directly from finely-binned histograms of the z components of photon production points. For the angular distribution $g(\cos \theta; s, E_0)$, we fill a two-dimensional histogram with produced Cherenkov photons' positions and directions, z and $\cos \theta$. Since tracks can scatter and shower away from their original $+z$ trajectories, θ is taken relative to a vector that points from the event starting vertex to the “center-of-mass” of photon production points. (The waist seen in $g(\cos \theta; s, E_0)$ in Figure 9.5 is an artifact of this.)

This is done separately at energies ranging from 5 MeV to 3100 MeV, in steps that start at 5 MeV and gradually increase to 200 MeV. Note that none of the production profiles is input directly into the fitter. They are used only to create the track integrals of §9.2.2, the energy dependences of which are parametrized to provide a smooth likelihood surface.

C.5 Glossary for Chapter 9

This reference section summarizes the functions and tables used by the reconstruction algorithm. In Chapter 9, we often use modified function arguments for clarity. This section shows all true functional dependences.

Properties of the detector

$\Omega(r)$

Solid angle subtended by a PMT at a distance r .

$\epsilon(\eta)$

PMT acceptance as a function of the photon incident angle η .

$T_{\text{sci}}(r)$

Transmission of scintillation light over a distance r . This function accounts for the wavelength dependence of scintillation light production, propagation, and detection.

$T_{\text{Ch}}(r)$

Transmission of Cherenkov light over a distance r . See previous item.

$A_{\text{sci}}(R, \cos \Theta)$

Scintillation scattering table.

$A_{\text{Ch}}(R, \cos \Theta, \cos \theta, \phi)$

Cherenkov scattering table.

Light production (particle dependent)

These functions exist separately for μ and e/γ track types.

$\Phi_{\text{sci}}(E_0)$

Number of scintillation photons produced by a particle with energy E_0 . This and the next item are smooth functions of E_0 derived from Monte Carlo runs at many energies.

$\Phi_{\text{Ch}}(E_0)$

Number of Cherenkov photons produced by a particle with energy E_0 . See previous item.

$\rho_{\text{sci}}(s; E_0)$

Spatial distribution of scintillation light production. (Not a direct fitter input – see

integrals below.)

$\rho_{\text{Ch}}(s; E_0)$

Spatial distribution of Cherenkov light production. (Not a direct fitter input – see integrals below.)

$g(\cos \theta; s, E_0)$

Angular distribution of Cherenkov light production. This depends both on the energy E_0 and on the position along the track s . (Not a direct fitter input – see integrals below.)

Production profile integrals (particle dependent)

These functions exist separately for μ and e/γ track types. They are smooth functions of E_0 , although we have not described the energy parametrizations used. They are built from integrals involving the production profiles, per §9.2.2.

$I_i^{\text{sci}}(E_0)$

$\int_{-\infty}^{\infty} ds \rho_{\text{sci}}(s; E_0) s^i$, evaluated numerically in 2.6 cm steps. $i \in \{1, 2\}$.

$I_i^{\text{Ch}}(E_0)$

$\int_{-\infty}^{\infty} ds \rho_{\text{Ch}}(s; E_0) s^i$, evaluated numerically in 2.6 cm steps. This item is used in conjunction with the Cherenkov scattering table. $i \in \{1, 2\}$.

$\mathcal{I}_i^{\text{Ch}}(E_0, r(0), \cos \theta(0))$

$\int_{-\infty}^{\infty} ds \rho_{\text{Cer}}(s; E_0) g(\cos \theta(s); s, E_0) s^i$, evaluated numerically in 2.6 cm steps. The s dependence of $\cos \theta$ is specified by the $r(0)$ and $\cos \theta(0)$ arguments of \mathcal{I} . The integrals are repeated in each of 100×200 bins of $[r(0), \cos \theta(0)]$. The $r(0)$ tabulation stops at 13 m. $i \in \{0, 1, 2\}$.

Corrected time p.d.f.'s (particle dependent)

$G_{\text{Ch}}(t^c; E_0, \mu_{\text{prompt}})$

The prompt (“Cherenkov”) primitive distribution. This would be the corrected time p.d.f. if late light were absent. This is a smooth function of all its arguments, per §9.3.1.

$G_{\text{sci}}(t^c; E_0, \mu_{\text{late}})$

The late (“scintillation”) primitive distribution. This would be the corrected time p.d.f. if prompt light were absent. This is a smooth function of all its arguments, per §9.3.1.

Appendix D

Low energy region

Understanding the excess of ν_e candidates over expected backgrounds below $E_\nu^{\text{rec}} = 475$ MeV (Figure D.1, Table D.1) is a high priority within the MiniBooNE collaboration. As Figure 14.8 suggests, the energy dependence of the excess is inconsistent with a two-neutrino oscillation explanation. A recent paper by M. Maltoni and T. Schwetz demonstrates that 3+2 and 3+3 models, which posit multiple sterile neutrinos, can accommodate the MiniBooNE and LSND observations together, but only if the null results of short-baseline disappearance experiments are ignored [146]. If the low energy discrepancy is caused by overlooked or misunderstood backgrounds, upcoming long-baseline experiments could face similar issues.

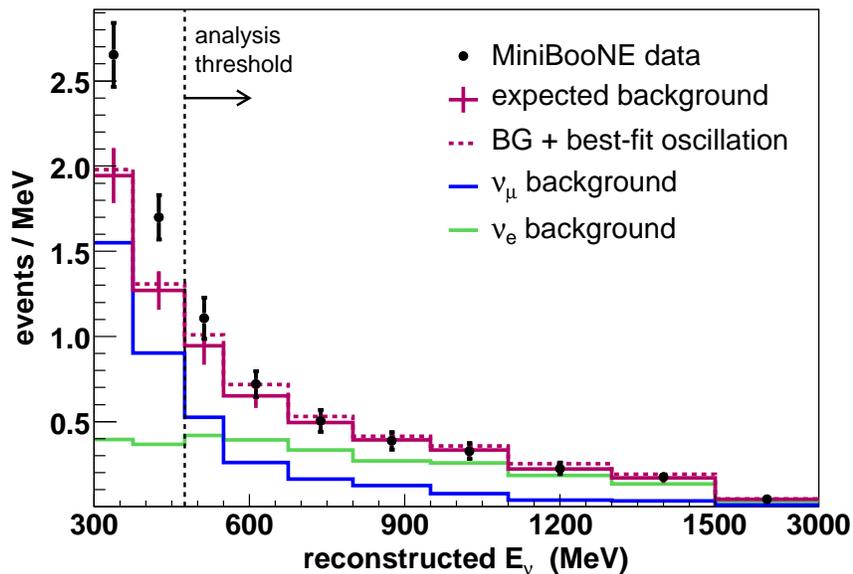


Figure D.1: Distribution of reconstructed neutrino energies for ν_e candidate events. An excess is seen at lower energies where ν_μ misidentification dominates.

ν_μ CC QE	12 ± 4
$\nu_\mu e \rightarrow \nu_\mu e$	5 ± 2
other ν_μ events	13 ± 5
NC π^0	76 ± 10
NC $\Delta \rightarrow N\gamma$	51 ± 8
dirt events	50 ± 5
$\mu \rightarrow \nu_e$	48 ± 7
$K^+ \rightarrow \nu_e$	14 ± 6
$K_L^0 \rightarrow \nu_e$	4 ± 2
$\pi^+ \rightarrow \nu_e$	1 ± 1
total background	273 ± 21
observed	369

Table D.1: Events in the reconstructed neutrino energy range $300 \text{ MeV} < E_\nu^{\text{rec}} < 475 \text{ MeV}$. (The systematic errors listed are correlated with one another.)

Photonuclear absorption

Figure D.2 shows the cross section for photonuclear absorption

$$\gamma + {}^{12}\text{C} \rightarrow X + (\text{no similar } \gamma) \quad (\text{D.1})$$

along with cross sections for electromagnetic processes. Photonuclear absorption is not modeled in GEANT3 and was absent from all background predictions. This reaction can essentially remove a γ from an event, leaving in its stead nearly invisible nucleons. Since our π^0 misidentification rate of $\mathcal{O}(0.01)$ is comparable to $\sigma_{\text{nuc}}/\sigma_{\text{elec}}$, a significant π^0 -based background could result.

An upgraded Monte Carlo which includes this process is now available, and samples are being generated. A preliminary version that leaves out the photonuclear final state particles suggests no increase in background estimates above 475 MeV but gives ~ 40 additional background events in the $300 \text{ MeV} < E_\nu^{\text{rec}} < 475 \text{ MeV}$ range – 40% of the observed excess. Including final state particles will increase reconstructed energies, bringing more of these events above the 300 MeV threshold and into the sample. The magnitude of the increase will be known soon.

Some disfavored explanations

- *Additional ν_e flux.* Figure D.3 shows the reconstructed energy and direction (E and U_z) of ν_e candidates. The low- E_ν^{rec} U_z spectrum is particularly interesting, as the expected ν_e and ν_μ backgrounds have quite different shapes. The U_z distribution of the observed excess disfavors a ν_e flux explanation (oscillation-based or otherwise).
- *Beam-off activity.* Running “random” triggers (§3.7) through the analysis chain suggests that ~ 2 beam-unrelated events should appear in the ν_e candidate sample.
- *Dirt.* Dirt events are examined directly in Figure 11.10. No evidence for a dirt excess

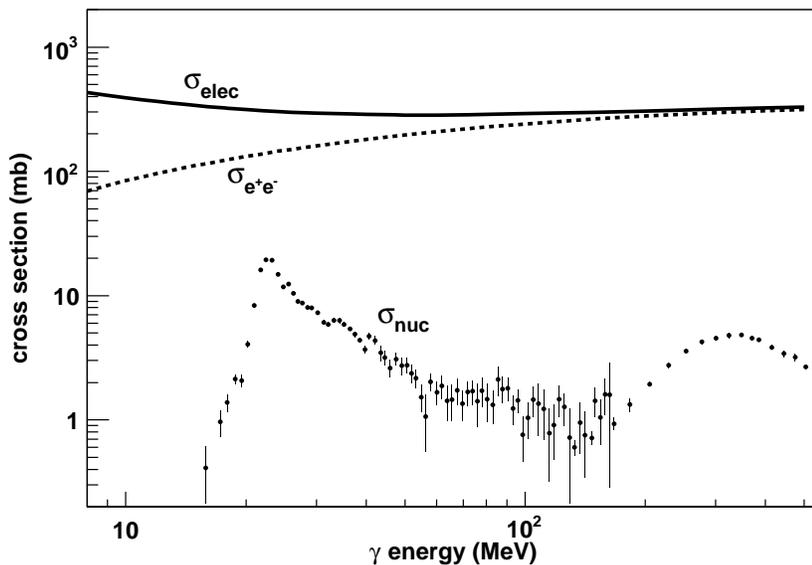


Figure D.2: Cross sections for γ interactions on carbon as a function of γ energy. Electromagnetic processes (“elec”), including the dominant pair production (“ e^+e^- ”), result in visible showers. Photonuclear absorption (“nuc”) is rarer but can cause a γ to go undetected. Data are from Refs. [147] and [148].

is seen (although we would have corrected one away if it were).

- ν_μ CC QE mis-ID. We verify the simulation of muon mis-ID by applying the ν_e particle ID cuts to the large sample of decay-tagged ν_μ CC QE events. No significant data/simulation discrepancies are seen.
- Hyperon decay ν_e flux. The beam simulation does not propagate or decay any strange baryons produced in the target. This absent ν_e flux component, however, is found to be negligible.

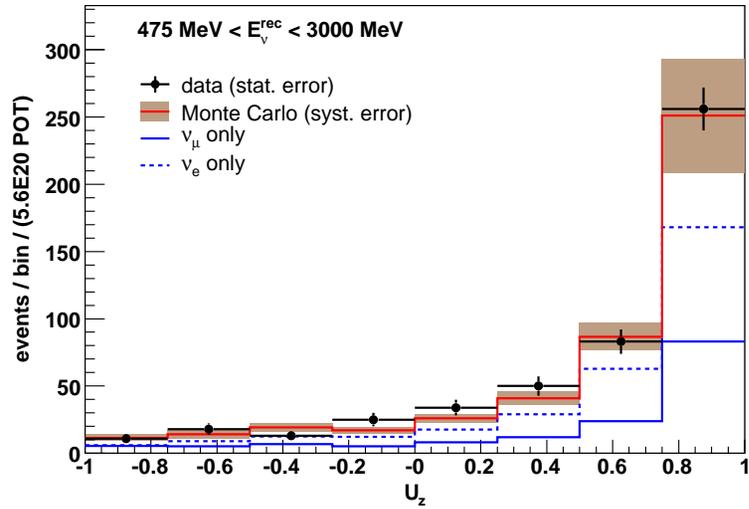
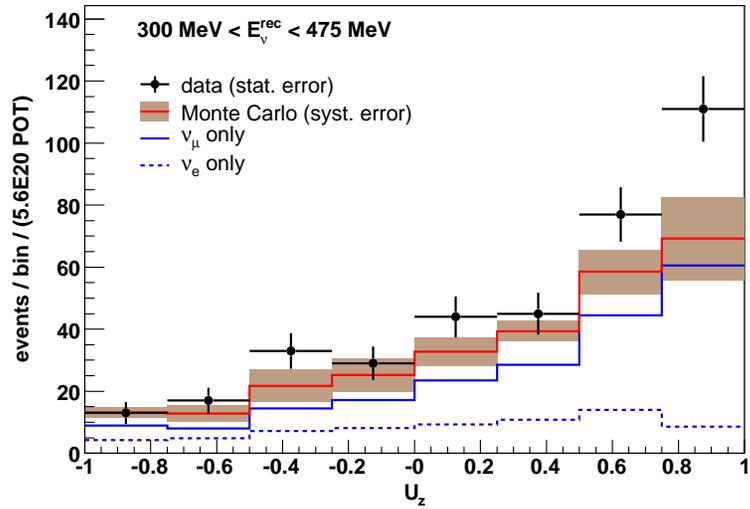
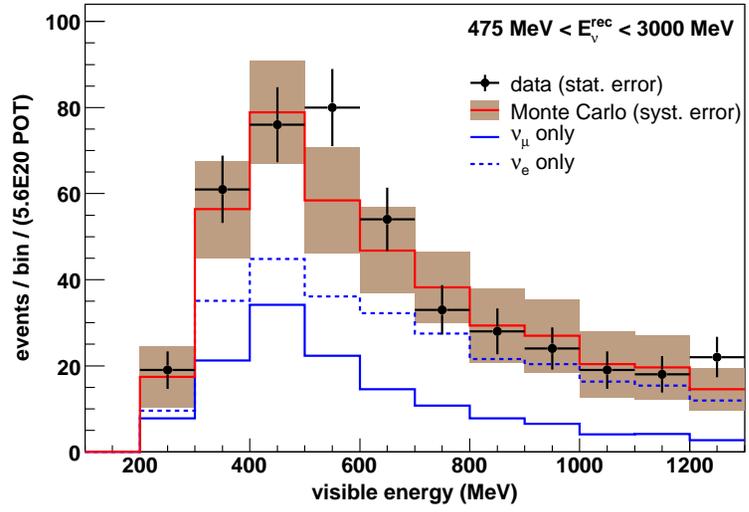
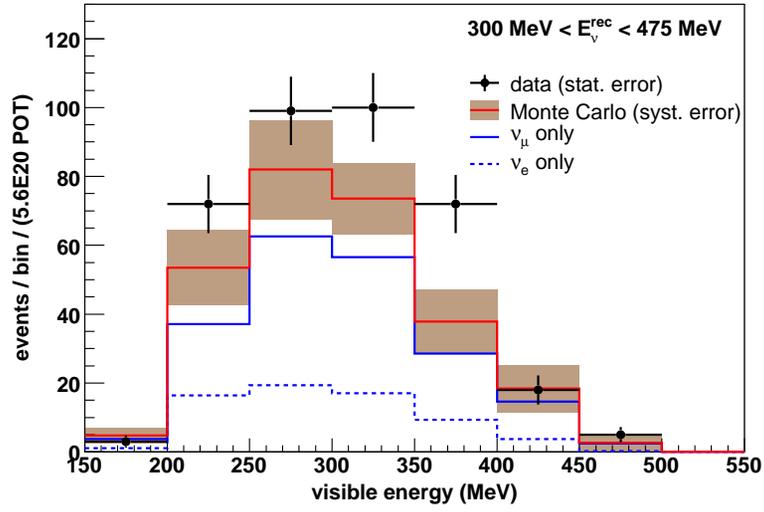


Figure D.3: Reconstructed energy E and direction U_z (top and bottom) for ν_e candidates in the low-energy and main analysis regions (left and right). The agreement in the right plots is good, and the excess in the left can be seen.

Bibliography

- [1] Pauli's letter, dated December 4, 1930, can be found in *Physics Today* (Sept. 1978).
- [2] F. Reines and C. L. Cowan, *Phys. Rev.* **92**, 830 (1953).
- [3] C. L. Cowan *et al.*, *Science* **124**, 103 (1956).
- [4] S. L. Glashow, J. Iliopoulos, and L. Maiani, *Phys. Rev.* **D2**, 1285 (1970); S. Weinberg, *Phys. Rev. Lett.* **19**, 1264 (1967); A. Salam, p. 367 of *Elementary Particle Theory*, ed. N. Svartholm (Almquist and Wilksells, Stockholm, 1969).
- [5] A. Zee, *Quantum Field Theory in a Nutshell*, Princeton University Press, Princeton (2003).
- [6] K. Zuber, *Neutrino Physics*, IOP, Bristol (2004).
- [7] B. Kayser, F. Gibrat-Debu, and F. Perrier, *World Sci. Lect. Notes Phys.* **25**, 1 (1989).
- [8] F. Boehm and P. Vogel, *Physics of Massive Neutrinos*, Cambridge University Press, Cambridge (1992).
- [9] W.-M. Yao *et al.*, *Journal of Physics G* **33**, 1 (2006).
- [10] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*, Addison-Wesley, Reading (1995).
- [11] R. N. Mohapatra and G. Senjanović, *Phys. Rev. Lett.* **44**, 912 (1980).
- [12] C. W. Kim and A. Pevsner, *Contemp. Concepts Phys.* **8**, 1 (1993).
- [13] Z. Maki, M. Nakagawa, and S. Sakata, *Prog. Theor. Phys.* **28**, 870 (1962).
- [14] B. H. J. McKellar, *Phys. Lett.* **B97**, 93 (1980).
- [15] J. Angrik, *KATRIN Design Report 2004*, KZKA-7090 (2004).
- [16] K. S. Krane, *Introductory Nuclear Physics*, Wiley, New York (1987).
- [17] S. R. Elliott and P. Vogel, *Ann. Rev. Nucl. Part. Sci.* **52**, 115 (2002).
- [18] D. Akimov *et al.*, *Nucl. Phys. Proc. Suppl.* **138**, 224 (2005).

- [19] H. V. Klapdor-Kleingrothaus *et al.*, Phys. Lett. **B586**, 198 (2004).
- [20] D. N. Spergel *et al.*, astro-ph/0603449 (2006).
- [21] S. Hannestad, JCAP **0305**, 004 (2003).
- [22] H. T. Wong *et al.*, Phys. Rev. **D75**, 012001 (2007).
- [23] D. A. Krakauer *et al.*, Phys. Rev. **D44**, 6 (1991).
- [24] M. Koshiya, Phys. Rept. **220**, 229 (1992).
- [25] B. Kayser, Phys. Rev. D **24**, 110 (1981).
- [26] L. Wolfenstein, Phys. Rev. **D17**, 2369 (1978); S. P. Mikheev and A. Y. Smirnov, Sov. J. Nucl. Phys. **42**, 913 (1985).
- [27] J. N. Bahcall, Phys. Rev. Lett. **12**, 300 (1964).
- [28] R. Davis, Phys. Rev. Lett. **12**, 303 (1964).
- [29] R. Davis, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).
- [30] B. T. Cleveland *et al.*, Astrophys. J. **496**, 505 (1998).
- [31] J. N. Bahcall, M. H. Pinsonneault, and S. Basu, Astrophys. J. **555**, 990 (2001).
- [32] W. Hampel *et al.*, Phys. Lett. **B447**, 127 (1999).
- [33] M. Altmann *et al.*, Phys. Lett. **B616**, 174 (2005).
- [34] V. N. Gavrin, Nucl. Phys. Proc. Suppl. **138**, 87 (2005).
- [35] K. S. Hirata *et al.*, Phys. Rev. Lett. **63**, 16 (1989).
- [36] S. Fukuda *et al.*, Phys. Lett. **B539**, 179 (2002).
- [37] S. N. Ahmed *et al.*, Phys. Rev. Lett. **92**, 181301 (2004).
- [38] B. Aharmim *et al.*, Phys. Rev. C **72**, 055502 (2005).
- [39] K. Eguchi *et al.*, Phys. Rev. Lett. **90**, 021802 (2003).
- [40] T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005).
- [41] K. S. Hirata *et al.*, Phys. Lett. **B205**, 416 (1988).
- [42] K. S. Hirata *et al.*, Phys. Lett. **B280**, 146 (1992).
- [43] R. Becker-Szendy *et al.*, Phys. Rev. D **46**, 3720 (1992).
- [44] Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998).

- [45] S. Fukuda *et al.*, Phys. Rev. Lett. **85**, 3999 (2000).
- [46] Y. Ashie *et al.*, Phys. Rev. Lett. **93**, 101801 (2004).
- [47] M. Apollonio *et al.*, Phys. Lett. **B466**, 415 (1999).
- [48] D. G. Michael *et al.*, Phys. Rev. Lett. **97**, 191801 (2006).
- [49] H. Murayama, <http://hitoshi.berkeley.edu/neutrino/>.
- [50] A. Aguilar *et al.*, Phys. Rev. D **64**, 112007 (2001).
- [51] C. Athanassopoulos *et al.*, Phys. Rev. Lett. **75**, 2650 (1995).
- [52] C. Athanassopoulos *et al.*, Phys. Rev. C **54**, 2685 (1996).
- [53] B. Armbruster *et al.*, Phys. Rev. D **65**, 112001 (2002).
- [54] Y. Declais *et al.*, Nucl. Phys. B **434**, 503 (1995).
- [55] G. L. Fogli *et al.*, Nucl. Phys. Proc. Suppl. **66**, 315 (1998).
- [56] J. Drees, Int. J. Mod. Phys. **A17**, 3259 (2002).
- [57] M. Sorel, J. M. Conrad, and M. Shaevitz, Phys. Rev. **D70**, 073004 (2004).
- [58] J. Pulido and B. C. Chauhan, Nucl. Phys. Proc. Suppl. **145**, 37 (2005).
- [59] H. Murayama and T. Yanagida, Phys. Lett. **B520**, 263 (2001).
- [60] G. Barenboim, L. Borissov, and J. D. Lykken, Phys. Lett. **B534**, 106 (2002).
- [61] I. Stancu *et al.*, *Technical Design Report for the MiniBooNE Neutrino Beam* (2001) (unpublished). Available at <http://www-boone.fnal.gov/publicpages/target.tdr.ps.gz>.
- [62] A. A. Aguilar-Arevalo *et al.*, *The MiniBooNE Run Plan* (2003) (unpublished). Available at <http://www-boone.fnal.gov/publicpages/runplan.ps.gz>.
- [63] I. Stancu *et al.*, *The MiniBooNE Detector Technical Design Report*, FERMILAB-TM-2207 (2001).
- [64] J. L. Raaf *et al.*, IEEE Trans. Nucl. Sci. **49**, 957 (2002).
- [65] B. T. Fleming *et al.*, IEEE Trans. Nucl. Sci. **49**, 984 (2002).
- [66] Ludox[®] is registered trademark of Grace Davison.
- [67] Available at http://sales.hamamatsu.com/assets/pdf/parts_R/R5912.pdf.
- [68] J. R. Musser, Phys. Rev. Lett. **94**, 101805 (2005), and Ref. 1 within.

- [69] S. Agostinelli *et al.*, Nucl. Instrum. Meth. **A506**, 250 (2003).
- [70] D. Casper, Nucl. Phys. Proc. Suppl. **112**, 161 (2002).
- [71] R. Brun *et al.*, *GEANT3*, CERN-DD/EE/84-1.
- [72] H. A. Tanaka, MiniBooNE Technical Note 177 (unpublished).
- [73] J. R. Sanford and C. L. Wang, Brookhaven National Laboratory, AGS internal reports 11299 and 11479 (1967) (unpublished).
- [74] M. Tzanov, MiniBooNE Technical Note 204 (unpublished).
- [75] New analysis by J. Link of data from E910 collaboration, Brookhaven National Laboratory (2005) (unpublished). For a description of the E910 experiment: I. Chemakin *et al.*, Nucl. Phys. **A639**, 407 (1998).
- [76] M. G. Catanesi *et al.*, hep-ex/0702024 (2007).
- [77] M. H. Shaevitz, MiniBooNE Technical Note 188 (unpublished).
- [78] F. Abe *et al.*, Phys. Rev. **D36**, 1302 (1987).
- [79] R. P. Feynman, Phys. Rev. Lett. **23**, 1415 (1969).
- [80] M. H. Shaevitz, MiniBooNE Technical Note 209 (unpublished).
- [81] T. Abbott *et al.*, Phys. Rev. **D45**, 3906 (1992).
- [82] A. Aleshin *et al.*, ITEP-77-80 (1977).
- [83] J. V. Allaby *et al.*, CERN 70-12, Nuclear Physics Division (1970).
- [84] D. Dekkers *et al.*, Phys. Rev. **137**, B962 (1965).
- [85] T. Eichten *et al.*, Nucl. Phys. **B44**, 333 (1972).
- [86] I. A. Vorontsov *et al.*, ITEP-88-11 (1988).
- [87] R. A. Lundy *et al.*, Phys. Rev. Lett. **14**, 504 (1965).
- [88] G. J. Marmer *et al.*, Phys. Rev. **179**, 1294 (1969).
- [89] N. V. Mokhov and S. I. Striganov, *MARS15 overview*, Hadronic Shower Simulations Workshop, Batavia (2006).
- [90] V. V. Gachurin *et al.*, ITEP-59-1985.
- [91] B. M. Bobchenko *et al.*, Sov. J. Nucl. Phys. **30**, 805 (1979).
- [92] H. A. Tanaka, MiniBooNE Technical Note 186 (unpublished).

- [93] D. Ashery *et al.*, Phys. Rev. **C23**, 2173 (1981).
- [94] B. W. Allardyce *et al.*, Nucl. Phys. **A209**, 1 (1973).
- [95] R. J. Glauber, *Lectures in Theoretical Physics*, ed. W. E. Britten *et al.*, Interscience, New York (1959).
- [96] V. Franco, Phys. Rev. C **6**, 748 (1972).
- [97] L. Bugel (private communication).
- [98] G. P. Zeller, hep-ex/0312061 (2003).
- [99] N. J. Baker *et al.*, Phys. Rev. D **23**, 2499 (1981).
- [100] C. H. Llewellyn Smith, Phys. Rept. **3**, 261 (1972).
- [101] P. Lipari, Nucl. Phys. Proc. Suppl. **112**, 274 (2002).
- [102] G. F. Chew and S. Mandelstam, Phys. Rev. **119**, 467 (1960).
- [103] M. G. Olsson, E. T. Osypowski, and E. H. Monsay, Phys. Rev. D **17**, 2938 (1978).
- [104] H. Budd, A. Bodek, and J. Arrington, Nucl. Phys. Proc. Suppl. **139**, 90 (2005).
- [105] R. A. Smith and E. J. Moniz, Nucl. Phys. **B43**, 605 (1972).
- [106] O. Benhar, nucl-th/0307061 (2003).
- [107] D. C. Cox *et al.*, *A Guide to Neutrino Cross Sections at MiniBooNE* (unpublished).
- [108] D. Rein and L. M. Sehgal, Ann. Phys. **133**, 79 (1981).
- [109] W. Krenz *et al.*, Nucl. Phys. **B135**, 45 (1978).
- [110] H. Faissner *et al.*, Phys. Lett. **B125**, 230 (1983).
- [111] D. Rein and L. M. Sehgal, Nucl. Phys. **B223**, 29 (1983).
- [112] N. G. Kelkar, E. Oset, and P. Fernandez de Cordoba, Phys. Rev. **C55**, 1964 (1997).
- [113] S. K. Singh, M. Sajjad Athar, and S. Ahmad, Phys. Rev. Lett. **96**, 241801 (2006).
- [114] W. Reuter *et al.*, Phys. Rev. C **26**, 806 (1982).
- [115] T. Katori, R. Tayloe, and G. P. Zeller, MiniBooNE Technical Note 219 (unpublished).
- [116] G. P. Zeller (private communication).
- [117] V. G. Zinov, V. G. Konin, and A. I. Mukhin, Sov. JETP **19**, 1292 (1964).
- [118] T. Suzuki, D. F. Measday, and J. P. Roalsvig, Phys. Rev. **C35**, 2212 (1987).

- [119] C. Zeitnitz and T. A. Gabriel, Nucl. Instrum. Meth. **A349**, 106 (1994).
- [120] J. D. Jackson, *Classical Electrodynamics* 3rd ed., John Wiley & Sons, Inc. (1999).
- [121] J. Birks, *Theory and Practice of Scintillation Counting*, Pergamon Press (1964).
- [122] D. Ashery *et al.*, Phys. Rev. C **23**, 2173 (1981).
- [123] V. Papathanakos, MiniBooNE Technical Note 134 (unpublished).
- [124] R. C. C. Leite *et al.*, Phys. Rev. Lett. **14**, 7 (1965).
- [125] A. O. Bazarko *et al.*, MiniBooNE Technical Note 144 (unpublished).
- [126] B. C. Brown, S. Maza, and A. Pla-Dalmau, MiniBooNE Technical Note 145 (unpublished).
- [127] B. C. Brown *et al.*, IEEE Nucl. Sci. Symp. Conf. Rec. **1**, 652 (2004).
- [128] H.-O. Meyer, MiniBooNE Technical Note 90 (unpublished).
- [129] S. J. Brice *et al.*, Nucl. Instrum. Meth. **A562**, 97 (2006).
- [130] R. B. Patterson, MiniBooNE Meeting Archives, June 30, 2006, and July 27, 2006 (unpublished).
- [131] L. Lyons, *Statistics for Nuclear and Particle Physicists*, Cambridge University Press, Cambridge (1986).
- [132] I. Stancu, MiniBooNE Technical Note 50, §II.3 (unpublished).
- [133] MINUIT, CERN Program Library entry **D506**. Available at <http://wwwasdoc.web.cern.ch/wwwasdoc/minuit/minmain.html>.
- [134] R. Gran *et al.*, Phys. Rev. **D74**, 052002 (2006).
- [135] Aguilar-Arevalo, A. A. *et al.* (MiniBooNE collaboration), arXiv:0706.0926 (2007).
- [136] G. Cowan, *Statistical Data Analysis*, Oxford University Press, Oxford (1998).
- [137] Z. Djurcic *et al.*, MiniBooNE Technical Note 194 (unpublished).
- [138] H.-J. Yang, MiniBooNE Technical Note 203 (unpublished).
- [139] G. J. Feldman and R. D. Cousins, Phys. Rev. D **57**, 3873 (1998).
- [140] B. P. Roe *et al.*, Nucl. Instrum. Meth. **A543**, 577 (2005).
- [141] E. D. Church *et al.*, Phys. Rev. D **66**, 013001 (2002).
- [142] A. A. Aguilar-Arevalo *et al.*, Phys. Rev. Lett. **98**, 231801 (2007).

- [143] T. Kobilarcik, `DQ_beamline` documentation (unpublished).
- [144] AD829 Data Sheet. Rev. G (2004) available at http://www.analog.com/UploadedFiles/Data_Sheets/AD829.pdf.
- [145] HBOOK, CERN Program Library Long Writeup **Y250**. Available at http://wwwasdoc.web.cern.ch/wwwasdoc/hbook_html3/hboomain.html.
- [146] M. Maltoni and T. Schwetz, arXiv:0705.0107 (2007).
- [147] M. J. Berger *et al.*, XCOM: Photon Cross Section Database, version 1.3 (2007). Available at <http://physics.nist.gov/xcom>.
- [148] J. Ahrens, Nucl. Phys. **A446**, 229c (1985).