

**Constraining neutrino oscillation
and interaction parameters
with the NOvA Near Detector
and Far Detector data
using Markov Chain Monte Carlo**

A dissertation submitted by

Michael Dolce

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Physics

Tufts University

August 2023

Adviser: Prof. Hugh Gallagher

Abstract

This thesis reports a constraint of the neutrino oscillation parameters Δm_{32}^2 , $\sin^2 \theta_{23}$, and δ_{CP} using the NuMI Off-Axis ν Appearance (NOvA) experiment's Near Detector (ND) data and Far Detector (FD) fake data set simultaneously. This thesis also reports a constraint on NOvA's systematic uncertainty model solely with its Near Detector data. The Hamiltonian Monte Carlo algorithm is used to estimate Bayesian Credible Intervals for the oscillation and interaction parameters. The 1σ Credible Intervals for $\sin^2 \theta_{23}$ are $(0.44, 0.512) \cup (0.536, 0.56)$, for Δm_{32}^2 ($2.41 \times 10^{-3} \text{ eV}^2$, $2.52 \times 10^{-3} \text{ eV}^2$), and for δ_{CP} $(0.74\pi, 1.1\pi) \cup (1.38\pi, 1.58\pi)$. The statistical power of the ND data constrains NOvA's interaction parameters, while the FD fake data constrains the oscillation parameters. This is the first analysis within NOvA to constrain the ND and FD prediction simultaneously, and to investigate the neutrino interaction modeling in the context of constraining the oscillation parameters. To constrain the ND data requires a sophisticated understanding of the neutrino interaction modeling and its uncertainties. The interested reader is advised to focus on Chapters 4 and 6, which discuss the ND selection, uncertainties, and ND-only fits to data. The reader interested in oscillation parameter constraints will find this in Chapter 7.

Acknowledgments

The old idiom “it takes a village” is certainly true to raise a PhD student. There are so many people that have helped me to reach this point that I must acknowledge.

First and foremost, I would like to thank my parents. From very early on I learned to be curious about all things, not just physics. Their constant support, encouragement, and listening are just a few things I will always be grateful for.

I would also like to thank my brother, who as an engineer, looks at physics in a slightly different way than I, and I always appreciate that. More importantly, his support and feedback – on all things, not just physics – are things I will always value.

I would like to thank my advisor Professor Hugh Gallagher for all of the support he has provided to me: professionally and personally. I have always appreciated your knowledge and expertise, but also your mentorship style to allow me to grow and develop as a researcher.

I want to thank Jeremy Wolcott who has helped guide me and kept me on the right path from the very beginning. I am extremely grateful for that, and for the physics, programming, and everything in between you have had the patience to teach me over the years.

I would like to thank my many NOvA collaborators that have helped to make this thesis a reality. I would like to specifically thank the members of the 3flavor group Louise Suter, Ryan Nichol, Zoya Vallari, and Erika Catano-Mur for their helpful comments over the last few years. Thank you to Professor Anthony Mann for his reminders about the “big picture”. I’d like to thank Professor Mayly Sanchez

for your input, guidance, and support to allow my work to be heard by others. I want to acknowledge Maria Martinez-Casales for her incredible work and help with my many questions. Much of our thesis projects were in lockstep. I've really enjoyed that, and to be able to work with you.

I want to thank all of the Tufts Physics & Astronomy graduate students, past and present. To name a few of the past, thank you Travis, Saqib, Josh, Katie, Alec for always being available for all kinds of questions. To those of the present, to name a few: Kenny, Anuja, Jack, Carter. You all have made the experience of being a graduate student that much more enjoyable. Best of luck on each of your journeys.

List of Figures

1.1	β decay spectrum	2
1.2	Charged-Current and Neutral Current Feynman diagrams	4
1.3	Oscillation probability as a function of energy and length	11
1.4	ν_μ (un)-oscillated spectrum	14
1.5	Neutrino mass ordering	15
1.6	NOvA Neutrino 2020 ν_μ spectra	17
1.7	NOvA Neutrino 2020 ν_e spectra	17
1.8	Δm_{32}^2 vs. $\sin^2 \theta_{23}$ 2D contour	18
1.9	NOvA δ_{CP} contour and exclusion	19
1.10	NOvA bi-probability plot	20
2.1	ν and $\bar{\nu}$ cross sections	23
2.2	Initial nucleon momentum	26
2.3	QE Feynman diagram	27
2.4	Effect of ZExp_AxialFFSyst2020_EV1 in Q^2	29
2.5	MEC Feynman diagram	30
2.6	Valenciá model 2p2h cross section in true (\vec{q} , q_0)	30
2.7	MEC tune on E_{had}^{vis}	31
2.8	MEC tune	32
2.9	RES Feynman diagram	34
2.10	DIS Feynman diagram	36
2.11	FSI cartoon illustration	37
2.12	hN FSI tune	39

2.13	hN FSI uncertainty	40
3.1	NOvA experiment map	43
3.2	Fermilab Accelerator Complex	45
3.3	NuMI beamline schematic	47
3.4	Horn focusing	48
3.5	Accumulated POT	48
3.6	NOvA Off-axis design	50
3.7	NOvA detectors rendering	51
3.8	Illustration of NOvA detector design	52
3.9	NOvA Near Detector	53
3.10	NOvA Far Detector	54
3.11	ND+FD Flux PC 00	57
4.1	Reconstructed ν_μ event	59
4.2	ND topological cuts	65
4.3	CVN $_\pi$ distributions for final state particles in FHC $\mu + \pi^\pm + X$ topology	66
4.4	ND 2D topologies for ν	69
4.5	ND 2D topologies for $\bar{\nu}$	70
4.6	$\mu + \pi^\pm + X$ FHC in Q^2 slices	72
4.7	$\mu + \pi^\pm + X$ FHC in W slices	73
4.8	NOvA ND MC FHC topological samples of 1D E_{had}^{vis} projections by truth interaction	74
4.9	NOvA ND MC RHC topological samples of 1D E_{had}^{vis} projections by truth interaction	75
4.10	NOvA ND MC FHC topological samples of 1D E_{had}^{vis} projections by final state particle	77
4.11	NOvA ND MC RHC topological samples of 1D E_{had}^{vis} projections by final state particle	78
4.12	Flow charts for ν_μ and ν_e reconstruction	79

4.13	FHC and RHC quantile cuts	80
4.14	FD FHC ν_μ MC	82
4.15	FD RHC ν_μ MC	83
4.16	Reconstructed ν_e event	84
4.17	FD ν_e and $\bar{\nu}_e$ MC	85
5.1	MCMC Typical Set	91
5.2	HMC 2-dimensional example	97
5.3	Marginalization vs. Profiling	99
6.1	Prior for uncertainties	103
6.2	MCMC fake data fit Pull Summary	106
6.3	Rep. Sample on E_{had}^{vis} projection of FHC topologies	108
6.4	Rep. Sample on E_{had}^{vis} projection of RHC topologies	109
6.5	χ^2 metric per ND topology from fake data fit	110
6.6	1D marginal distributions of NOvA uncertainties from fake data fit	112
6.7	2D marginal distributions from fake data fit	114
6.8	MEC degenerate parameter	115
6.9	Data-MC comparisons of ND FHC topologies projected onto E_{had}^{vis}	117
6.10	Data-MC comparisons of ND RHC topologies projected onto E_{had}^{vis}	118
6.11	2D data-MC ratio of FHC topologies	120
6.12	2D data-MC ratio of RHC topologies	121
6.13	Preliminary fit to NOvA data: χ^2 by Topology	122
6.14	Preliminary fit to NOvA data: Pulls Summary	123
6.15	ND FHC μ topology χ^2 systematic ranking	126
6.16	ND FHC $\mu + P$ topology χ^2 systematic ranking	127
6.17	ND FHC $\mu + \pi^\pm + X$ topology χ^2 systematic ranking	128
6.18	ND FHC $\mu + P + X$ topology χ^2 systematic ranking	129
6.19	ND FHC <code>EvElse</code> topology χ^2 systematic ranking	130
6.20	Ranking uncertainties by χ^2 averaged over all ND topologies	131

6.21	ND FHC topologies with variable binning and phase space cut applied	133
6.22	ND RHC topologies with variable binning and phase space cut applied	134
6.23	Over-prediction of low- Q^2 events in $\mu + \pi^\pm + X$	137
6.24	Resonant cross section scaling uncertainties	139
6.25	Validation of RES scale uncertainty on true W	140
6.26	Impact of RES Δ scale uncertainty on $\mu + \pi^\pm + X$ topology	141
6.27	Impact of DIS hadronization uncertainty on $\mu + \pi^\pm + X$ topology	144
6.28	ND Data Pulls Summary after development	148
6.29	ND FHC 2D topological data-nominal MC χ^2 distributions	150
6.30	ND FHC 2D topological data-Rep.Sample MC χ^2 distributions	151
6.31	ND RHC 2D topological data-nominal MC χ^2 distributions	152
6.32	ND RHC 2D topological data-Rep.Sample MC χ^2 distributions	153
6.33	χ^2 per ND topology after improvements to MCMC machinery	154
7.1	Bounded prior for δ_{CP}	161
7.2	ND FHC Rep. Sample prediction in the 2020 FD quantile cuts	163
7.3	ND RHC Rep. Sample prediction in the 2020 FD quantile cuts	164
7.4	ND <i>Reco</i> E_μ with and without <code>kCorrMuEScaleSyst2020</code>	165
7.5	1D marginal of <code>kCorrMuEScaleSyst2020</code> from ND fit	166
7.6	FHC FD ν_μ Residual Difference	168
7.7	RHC FD ν_μ Residual Difference	169
7.8	ND FHC & RHC inclusive ν_μ sample in E_ν	170
7.9	FD FHC & RHC ν_e distributions with ND Residual Difference	171
7.10	Pull Summary plot of ND+FD joint fit	173
7.11	2D marginal distributions of systematics from ND+FD Residual Difference joint fit	174
7.12	1D marginal of δ_{CP}	176
7.13	1D marginal of Δm_{32}^2	177
7.14	1D marginal of $\sin^2 \theta_{23}$	178
7.15	δ_{CP} vs. $ \Delta m_{32}^2 $ 2D marginal	179

7.16	δ_{CP} vs. $\sin^2 \theta_{23}$ 2D marginal	180
7.17	$ \Delta m_{32}^2 $ vs. $\sin^2 \theta_{23}$ 2D marginal	181
7.18	FHC FD ν_μ Representative Sample predictions	183
7.19	RHC FD ν_μ Representative Sample predictions	184
7.20	FD ν_e Representative Sample predictions	185
7.21	FD FHC ν_μ random MCMC sample predictions	186
7.22	FD RHC ν_μ random MCMC sample predictions	187
7.23	FD ν_e Random MCMC Sample Predictions	188
7.24	Data-Rep. Sample χ^2 per ND topology and FD sample from ND+FD joint fit	189
7.25	2D marginals of oscillation parameters & uncertainties	191
A.1	MEC tune on <i>Reco</i> $ \vec{q} $	199
A.2	NOvA ND MC FHC topological samples of 1D <i>Reco</i> $ \vec{q} $ projections by truth interaction	201
A.3	NOvA ND MC RHC topological samples of 1D <i>Reco</i> $ \vec{q} $ projections by truth interaction	202
A.4	NOvA ND MC FHC topological samples of 1D <i>Reco</i> $ \vec{q} $ projections by final state particle	203
A.5	NOvA ND MC RHC topological samples of 1D <i>Reco</i> $ \vec{q} $ projections by final state particle	204
A.6	Representative Sample predictions on <i>Reco</i> $ \vec{q} $ projection on FHC topologies from fake data fit	205
A.7	Representative Sample predictions on <i>Reco</i> $ \vec{q} $ projection on RHC topologies from fake data fit	206
A.8	Data-MC comparisons of ND FHC topologies projected onto <i>Reco</i> $ \vec{q} $	207
A.9	Data-MC comparisons of ND RHC topologies projected onto <i>Reco</i> $ \vec{q} $	208
A.10	ND RHC μ χ^2 systematic ranking	209
A.11	ND RHC $\mu + P$ χ^2 systematic ranking	210
A.12	ND RHC $\mu + \pi^\pm + X$ χ^2 systematic ranking	211

A.13 ND RHC $\mu + X$ χ^2 systematic ranking	212
A.14 ND RHC EvElse χ^2 systematic ranking	213
A.15 2D Data-Rep. Sample MC Ratio distributions in ND FHC Topologies	214
A.16 2D Data-Rep. Sample MC Ratio distributions in ND RHC Topologies	216
A.17 MCMC autocorrelations	217
A.18 1D Projections onto $Reco \vec{q} $ of the Rep. Sample from the ND-only fit and ND+FD joint fit in FHC ND topologies	218
A.19 1D Projections onto $Reco \vec{q} $ of the Rep. Sample from the ND-only fit and ND+FD joint fit in RHC ND topologies	219
A.20 1D posterior distributions of δ_{CP} and $\sin^2 \theta_{23}$ in NO and IO	220
A.21 2D posterior distributions of oscillation parameters in NO and IO	221

List of Tables

1.1	Fundamental forces	3
1.2	Fundamental fermions	4
1.3	Global results of oscillation parameters	8
2.1	GENIE models	22
2.2	MEC parameters	33
2.3	Adjustment of the parameters for NOvA's hN FSI central-value tune.	38
4.1	Prong-CVN scores	60
4.2	Near Detector topologies and neutrino models	62
4.3	Topology names	64
4.4	Topology event count	65
6.1	Resonant Δ decays	138
6.2	Impact of RES relative cross section scaling uncertainties on ND topologies	139
6.3	DIS hadronization uncertainty	143
6.4	MCMC Rep. Sample systematic Pull Summary from fit to ND data.	145
7.1	Oscillation parameters for FD fake data	172

Contents

1	Neutrino Oscillations	2
1.1	Neutrinos and the Standard Model	3
1.2	Neutrino Mixing	5
1.3	Two-Flavor Oscillations	8
1.4	Three Flavor Oscillations	12
1.5	Neutrino Mass Ordering	15
1.6	Constraint on δ_{CP}	16
1.7	Recent NOvA Results	16
1.8	Significance of Thesis Measurement	20
2	Neutrino Interactions	22
2.1	Kinematic Variables	24
2.2	Nuclear Model	25
2.3	Quasi-Elastic Scattering	27
2.3.1	Quasi-Elastic Uncertainties	28
2.4	2p2h/Meson Exchange Current	29
2.4.1	MEC model tuning in NOvA	30
2.4.2	Meson Exchange Current Uncertainties	33
2.5	Resonant Scattering	34
2.5.1	Resonant Uncertainties	35
2.6	Deep Inelastic Scattering	35
2.6.1	Deep Inelastic Scattering Uncertainties	36

2.7	Final State Interactions	36
2.7.1	FSI Tuning	38
2.7.2	Final State Interaction Uncertainties	40
2.8	Other Interactions	40
2.8.1	Other Uncertainties	41
2.9	Neutrino Simulation	41
3	The NOvA Experiment	42
3.1	The NuMI Neutrino Beam	44
3.1.1	Accelerator Complex	44
3.1.2	NuMI Beamline	46
3.1.3	Off-axis Design	49
3.2	NOvA Detectors	51
3.2.1	The Near Detector	51
3.2.2	The Far Detector	52
3.2.3	Simulating the Detectors	53
3.3	Additional Systematic Uncertainties	54
3.3.1	Neutron Uncertainty	55
3.3.2	Detector Uncertainties	55
3.3.3	Flux Uncertainties	56
4	Selection and Topological Samples	58
4.1	Event Reconstruction	58
4.1.1	Particle Identification	60
4.2	Energy Reconstruction	60
4.2.1	Muon Energy	60
4.2.2	Hadronic Energy	61
4.3	ND Topological Samples	61
4.3.1	CVN $_{\pi}$ Performance for FHC $\mu + \pi^{\pm} + X$ Topology	65
4.3.2	Variables for ND Selection: E_{had}^{vis} & $Reco \vec{q} $	67

4.4	Far Detector Selection	79
4.4.1	ν_μ Quantile Selection	79
4.4.2	ν_e Selection	83
5	Bayesian Inference and Markov Chain Monte Carlo	87
5.1	Bayes Theorem	87
5.1.1	Credible Intervals	89
5.1.2	Marginalization	90
5.2	Markov Chain Monte Carlo	90
5.2.1	Metropolis-Hastings	93
5.3	Hamiltonian Monte Carlo	94
5.4	Hamiltonian Monte Carlo 2-D Example	96
5.5	The “Representative Sample” (Max. \mathcal{LL})	98
5.6	Stan Software Package	100
6	Near Detector Fit	102
6.1	Prior Choice for Uncertainties	102
6.2	ND Fake Data Fits	104
6.2.1	Fake data fit to QE, RES, and FSI Parameters	104
6.2.2	Degeneracies in MEC Double Gaussian Parameters	114
6.3	Data & MC Comparisons	116
6.4	Initial Fitting Attempts to ND Data	122
6.4.1	Adjustments to MCMC Input	124
6.5	Newly Created Cross-Section Uncertainties	136
6.5.1	RES Relative Cross Section Scaling Systematics	136
6.5.2	Δ and Other Resonant Scaling Systematics	139
6.5.3	DIS Hadronization Systematics	142
6.6	The ND Residual Difference	145
6.7	Conclusions from ND Fits	155

7	Near + Far Detector Joint Fit	157
7.1	Evaluating the Residual Difference	158
7.2	Priors of Oscillation Parameters	159
7.2.1	Prior $\sin^2 \theta_{23}$	159
7.2.2	Prior Δm_{32}^2	160
7.2.3	Prior δ_{CP}	160
7.3	Creating the FD Fake Data with the ND Residual Difference	161
7.3.1	Application to ν_{μ}	162
7.3.2	Application to ν_e	170
7.4	ND + FD Residual Difference Fit Result	172
7.5	Future Improvements	192
7.5.1	ND Improvements	192
7.5.2	ND+FD Improvements	193
7.5.3	MCMC Improvements	194
7.6	Conclusions	195
A	Appendix	199
A.1	MEC Tune in NOvA	199
A.2	ND Topological Samples	200
A.2.1	ND MC 1D Projections onto <i>Reco</i> $ \vec{q} $	200
A.2.2	ND Representative Sample 1D Projections onto <i>Reco</i> $ \vec{q} $ from Fake Data Fit	200
A.2.3	ND Systematics Ranked by χ^2 in ND RHC Topologies	200
A.3	ND Fit to Data	210
A.4	ND+FD Joint Fit	215

Chapter 1

Neutrino Oscillations

The neutrino is perhaps one of the least understood particles in our universe. Its story originated with the Italian physicist Wolfgang Pauli in 1930, while studying beta decay from light nuclei ($N \rightarrow N' + e^- + p$), posited a particle to explain the source of the missing energy, momentum, and angular momentum in the beta decay. He famously stated, “I have hit upon a desperate remedy. . .there could exist in the nuclei electrically neutral particles”, naming the particle the “neutron” [1]. It was theorized this decay of a light nucleus would produce a proton and electron with equal and opposite momentum, conserving momentum, and the sum of the pair’s energy would obey energy conservation.

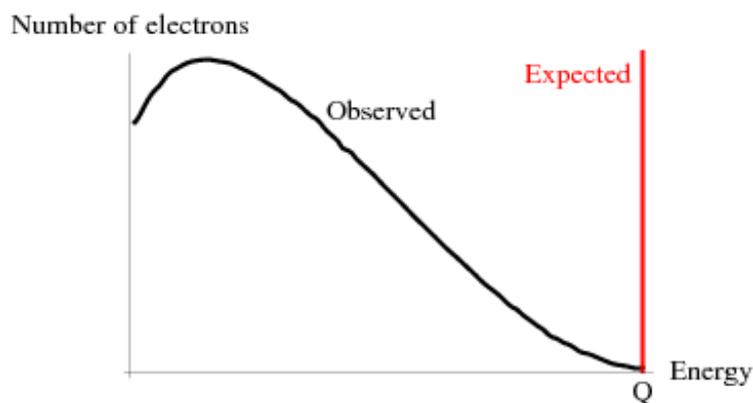


Figure 1.1: β decay spectrum. It was expected the energy of an electron in beta decay would be unique (red “Expected”). However, the electron was never observed at the “Expected” energy *and* instead exhibited a distribution of energies, “Observed”. These observations led Pauli to hypothesize the neutrally charge particle the neutrino. From [2].

However, an electron energy *spectrum* was observed – rather than a uniquely defined value – that observes no electrons at the expected electron energy (Fig. 1.1). Therefore, the electron *never* possesses the equal and opposite amount of energy from the proton, and the this missing energy must be present in the form of the “neutron” Pauli hypothesized; this neutrally charged particle – that cannot be directly detected – would carry away the remaining energy that was not observed in the electron energy spectrum. In 1932, the neutron as we know now was discovered, prompting Pauli’s “neutron” to be renamed by Italian physicist Enrico Fermi into “neutrino”, Italian for “little neutral one”, for its light mass and zero charge.

1.1 Neutrinos and the Standard Model

The Standard Model of particle physics is a theory of particle interactions encapsulating three of the four fundamental forces: electromagnetic, weak, and strong forces. Table 1.1 shows the relative strength of these forces and the associated (or theorized, in the case of gravity) boson, or force carrier.

Force	Strength	Boson	Spin	Mass/GeV
Strong	1	Gluon(g)	1	0
Electromagnetism	10^{-3}	Photon(γ)	1	0
Weak	10^{-8}	W boson(W^{\pm})	1	80.4
Weak	10^{-8}	Z boson(Z)	1	91.2
Gravity	10^{-37}	Graviton(?)(G)	2	0

Table 1.1: The four known forces of nature. The relative strengths are approximate indicative values for two fundamental particles at a distance of $1 \text{ fm} = 10^{-15} \text{ m}$ (roughly the radius of a proton). Adapted from Ref. [3]

The other component of the standard model are the Fermions, Table 1.2. Fermions can be divided into two sectors: the quark sector and lepton sector. The quark sector (right side of the table) shows the three generations of quarks. The strong interaction force is mediated by gluons. The theory is an $SU(3)$ gauge theory described by a color triplet and the six quark flavors. Color singlets are referred to as hadrons and have two subcategories: mesons and baryons. Mesons are pairs of a quark and anti-quark (such as a pion). Baryons are made of three quarks (such as a proton).

Gen.	Leptons			Quarks		
	Particle	Q	Mass/GeV	Particle	Q	Mass/GeV
1^{st}	electron (e^-)	-1	0.0005	down(d)	-1/3	0.003
	electron neutrino (ν_e)	0	$< 2^{-9}$	up(u)	+2/3	0.005
2^{nd}	muon (μ^-)	-1	0.0105	strange(s)	-1/3	0.1
	muon neutrino (ν_μ)	0	$< 0.17^{-3}$	charm(c)	+2/3	1.3
3^{rd}	tau (τ^-)	-1	1.7768	bottom(b)	-1/3	4.5
	tau neutrino (ν_τ)	0	$< 18.2^{-3}$	top(t)	+2/3	174

Table 1.2: The twelve fundamental fermions divided into quarks and leptons. The masses represent the current best experimental measurements. Adapted from [3] and [4].

The lepton sector is also divided into three generations. The charged leptons (e , μ , τ) have identical charge, -1 , and increase in mass with generation. Each charged lepton possesses a corresponding neutrino with the same lepton quantum number. This thesis focuses on the interactions of ν_μ and ν_e and the parameters that govern their oscillation.

Within the Standard Model, neutrinos only interact via the weak force. The weak force is mediated through one of two bosons, the W^\pm and Z^0 . In a Charged Current (CC) neutrino interaction, the weak force is mediated by the massive W^\pm boson, which exchanges a ± 1 charge. The Neutral Current (NC) interaction is predicted by the unification of the weak force with quantum electrodynamics. This interaction is mediated by a massive, neutral Z^0 boson. The tree-level Feynman diagrams can be seen in Fig. 1.2.

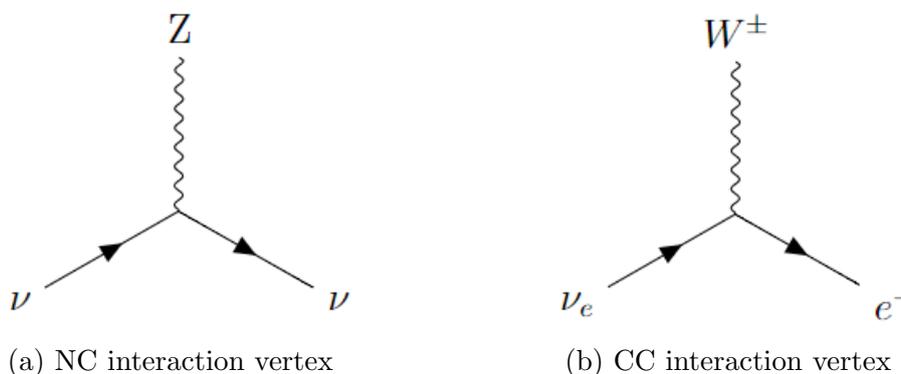


Figure 1.2: The neutral-current and charged-current weak interaction vertices (for the 1^{st} generation fermions).

The theory of leptogenesis is closely tied to the neutrino. In the theory, the Stan-

Standard Model contains both Dirac and Majorana neutrinos, and both are required to explain the matter-antimatter asymmetry in the early universe. Neutrino oscillation experiments seek to probe the theory by observing how matter particles differ from their antimatter partners. Neutrinos are the ideal tool because experiments such as NOvA can create and observe neutrinos as well as antineutrinos. By studying the different characteristics of these two particles provides insight into this question. In the context of NOvA, the objective is to constrain the parameter δ_{CP} , which describes the charge and parity violation between matter and antimatter in the neutrino sector, and will be discussed further in Sec. 1.6. However, NOvA's efforts will only clarify part of the theory; the remaining part requires observing heavy neutrinos at energies consistent with the early universe (or the Grand Unified Theory, GUT, scale). Nonetheless, NOvA will constrain CP-violation in right-handed light neutrinos. Understanding the characteristics of neutrinos will provide insight as to why there is matter at all in our universe.

1.2 Neutrino Mixing

The neutrino flavor eigenstates corresponding to the weak interaction – e , μ , τ – are composed of neutrino mass eigenstates via the propagation of the free particle Hamiltonian. Crucially, one state can be written as a linear combination of the other set of neutrino eigenstates, in matrix notation in Eq. 1.1, or with the bra-ket notation in Eq. 1.2. In Eq. 1.2 ν_α is the flavor eigenstate and ν_i is the mass eigenstate with mass m_i , where $i = 1, 2, 3$.

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{PMNS} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad (1.1)$$

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle. \quad (1.2)$$

Moreover, neutrinos are observed to oscillate, meaning there is a probability a neutrino of one flavor can be observed as a neutrino of another flavor after some time or long distances (typically in km). There are two fundamental concepts that govern neutrino oscillation: non-zero off-diagonal elements in the PMNS matrix U (commonly referred to as θ_{ij}), and non-degenerate mass eigenstates, $m_1 \neq m_2 \neq m_3$. We will see that neutrino oscillations depend on the mass-squared difference, e.g. $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$, and because $m_1 \neq m_2 \neq m_3$ there is some sinusoidal frequency to the oscillation probability. This is discussed further in Sections 1.3 and 1.4. This Section will focus on the off-diagonal elements of the PMNS matrix, or “mixing”.

The U is the unitary mixing matrix (i.e. $U^{-1} = U^\dagger$) known as the PMNS matrix, named after Pontecorvo-Maki-Nakagawa-Sakata who introduced the matrix and Pontecorvo who predicted the concept of oscillations [5] [6]. In the current understanding of the Standard Model, the PMNS matrix is a 3×3 matrix – relating to the three neutrino flavors – containing four parameters: three real, θ_{ij} , and one imaginary phase δ_{CP} ,

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13} e^{-i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{\frac{i\alpha_{21}}{2}} & 0 \\ 0 & 0 & e^{\frac{i\alpha_{31}}{2}} \end{pmatrix}. \quad (1.3)$$

where c_{ij} (s_{ij}) is $\cos \theta_{ij}$ ($\sin \theta_{ij}$) of the mixing angle. In Equation 1.3, the PMNS matrix is written as the product of multiple matrices. In this expression, each individual matrix corresponds to a neutrino mixing “sector”.

The first matrix (from the left) is referred to as the “atmospheric” sector, denoting the original source of detection: cosmic muon decays originating from the atmosphere. This sector – or matrix – contains the mixing angle parameter θ_{23} . Moreover, this sector can be accessed via accelerator experiments. Experiments

such as IceCube, Super-Kamiokande, and T2K all have produced measurements of the mixing angle and its partnering Δm_{32}^2 (the mass-squared difference term is explained in Sec. 1.5) [7]. Constraining θ_{23} is one of the primary physics goals of the NOvA experiment, and this thesis. Comparisons of previous accelerator-based oscillation experiments can be found in Fig. 1.8.

The second matrix refers to the “reactor” sector and the mixing angle θ_{13} , where nuclear reactors are used as the source of neutrinos. These types of experiments place detectors within proximity to the reactor, and observe oscillations – ν_e disappearance – over short baselines of $\mathcal{O}(100 \text{ m})$. For example the RENO, Double-Chooz and Daya Bay collaborations have produced leading constraints of this parameter [8]. In fact, the NOvA experiment was originally designed to constrain and measure this parameter. However, the profound success of the reactor experiments were able to produce an accurate measurement of θ_{13} that it is no longer a physics goal for NOvA; it is now a well-defined value used to determine θ_{23} . The reactor sector also contains the CP-violating phase δ_{CP} .

The third matrix is the “solar” sector, which is sensitive to θ_{12} . It can be observed from $\bar{\nu}_e$ disappearance – through inverse beta decay – over short ($\approx 200 \text{ km}$) distances. Experiments such as KamLAND have made strong measurements of this parameter proving it is non-zero [9].

The final matrix corresponds to the fate of the neutrino as a Majorana or Dirac particle. Neutrino oscillations are not sensitive to the Majorana phases, as it does not appear in the oscillations model.

In summary, there is a probability an initial neutrino, ν_α , will interact with matter as a different flavor, ν_β , which depends on the PMNS mixing angle values and the corresponding mass-squared differences. Many experiments exist to measure specific parameters of the matrix, and some have found great success. A summary of the parameters from a global fit (from 2019) can be found in Table 1.3. Furthermore, this fundamental concept of neutrino “mixing” is one of the two concepts that permits neutrino oscillations. The second is the non-degenerate mass eigenstates.

Parameter	value ($^\circ$)
θ_{12}	$33.82^{+0.78}_{-0.76}$
θ_{23}	$49.6^{+1.0}_{-1.2}$
θ_{13}	$8.61^{+0.13}_{-0.13}$
δ_{CP}	215^{+40}_{-29}

Table 1.3: Global results of oscillation parameter measurements from 2019. From [10].

To understand the impact of the mixing angle and the different mass eigenstate values, we will work through an example of neutrino mixing via a 2×2 PMNS matrix and calculate the subsequent probabilities for oscillation in this two-flavor neutrino paradigm. Then we will expand the mathematics of oscillations to the full three-flavor description using the complete 3×3 PMNS matrix. This is outlined in the following two Sections.

1.3 Two-Flavor Oscillations

In the two-flavor neutrino model, the “mixing” matrix, or PMNS matrix, is a 2×2 matrix with a single parameter θ that governs the mixing of flavors, and contains no phase factor

$$U_{PMNS} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \quad (1.4)$$

In this two-flavor paradigm, we will use two arbitrary flavors, α and β , and only have two mass eigenstates, ν_1 and ν_2 , where, again, $m_1 \neq m_2$. In addition, this derivation utilizes the plane wave approximation. The ν_α and ν_β flavors can be written as a linear combination of ν_1 and ν_2 given the 2×2 mixing matrix,

$$\begin{pmatrix} \nu_\alpha \\ \nu_\beta \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}. \quad (1.5)$$

A neutrino of arbitrary flavor ν_α , at some time t_0 , can be written in the bra-ket

notation as

$$|\nu_\alpha, t_0\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle, \quad (1.6)$$

where $U_{\alpha i}^*$ is the PMNS matrix and ν_i is the propagating neutrino mass eigenstate.

Next is to evolve $|\nu_\alpha, t_0\rangle$ in time. Starting at initial time t_0 , the state will evolve in time via the operator $U(t, t_0)$,

$$U(t, t_0) \equiv e^{-i\frac{\hat{H}(t-t_0)}{\hbar}}. \quad (1.7)$$

For time evolution, we require the free Hamiltonian

$$\hat{H}|\nu_i\rangle = E_i|\nu_i\rangle, \quad (1.8)$$

where E_i is the total energy of the particle. We utilize the near-massless approximation ($m_\nu \ll 1\text{eV}$) of neutrinos and Taylor expand about it,

$$H = E \approx pc \left(1 + \frac{m^2 c^4}{2p^2} \right). \quad (1.9)$$

Finally, with the time evolution operator, a neutrino of initial state $|\nu_\alpha, t_0\rangle$ evolves in time $t_0 \rightarrow t$,

$$\begin{aligned} |\nu_\alpha, t\rangle &= U(t, t_0) |\nu_\alpha, t_0\rangle \\ &= e^{-i\frac{\hat{H}(t-t_0)}{\hbar}} |\nu_\alpha, t_0\rangle \\ &= e^{-i\frac{\hat{H}(t-t_0)}{\hbar}} \sum_i U_{\alpha i}^* |\nu_i\rangle. \end{aligned} \quad (1.10)$$

We seek to observe a ν_β after some time, which is written as

$$|\nu_\beta\rangle = \cos\theta|\nu_2\rangle - \sin\theta|\nu_1\rangle. \quad (1.11)$$

The probability at time t the neutrino is observed as flavor β , we square and calculate

to get the amplitude,

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= |\langle \nu_\beta | \nu_\alpha, t \rangle|^2 \\ &= \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-iE_i(t-t_0)} \right|^2. \end{aligned} \quad (1.12)$$

We obtain the probability of ν_β as the observed neutrino – and by the conservation of probability – the probability of no oscillation (ν_α in the final state)

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta \sin^2 \left(1.27 \frac{\Delta m^2 L}{E} \right), \quad (1.13)$$

$$P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - P(\nu_\alpha \rightarrow \nu_\beta), \quad (1.14)$$

and rewrite t with L/c for the Hamiltonian in Eq. 1.9, if assuming natural units $\hbar = c = 1$. It is important to note the mass difference Δm^2 , as this plays a key role. If $\Delta m^2 = 0$ then there would be no oscillations at all; the probability would be 0. This is why the unique values of the mass eigenstates are so crucial to oscillations.

The $P(\nu_\alpha \rightarrow \nu_\beta)$ is known as the “appearance” probability – ν_β “appears” from ν_α . The $P(\nu_\alpha \rightarrow \nu_\alpha)$ is the “disappearance” or “survival” probability. Note how this probability is a sinusoidal function, and crucially, is dependent on two quantities: the neutrino energy, E , and the length of the neutrino travel, L . Fig. 1.3 shows the oscillation of probability as a function of these two quantities. As an experiment seeking to observe neutrino oscillation and measure the mixing parameters, maximizing the probability is critical to consider. The left plot shows the probability with a fixed energy and indicates the length where the probability for oscillation is largest, L_{osc} . This coincides with the $\sin^2 2\theta = 1$, i.e. $\theta = \pi/4$,

$$L_{osc} = \pi \frac{E_\nu (GeV)}{1.27 \Delta m^2 (eV^2)}. \quad (1.15)$$

The right plot shows the oscillation as a function of energy. In this case, E_{max}

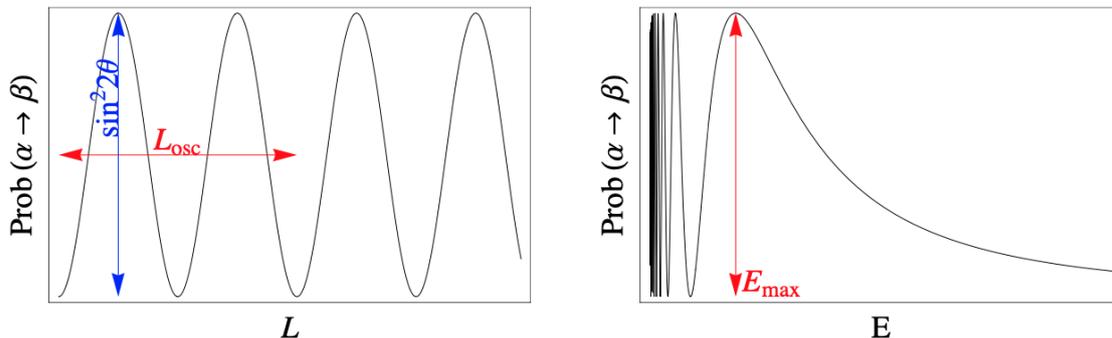


Figure 1.3: Oscillation probability as a function of energy (left) and length (right). To maximize the probability for a neutrino oscillation, there are ideal neutrino energies and lengths to observe the phenomenon.

indicates the ideal energy to maximize oscillations,

$$E_{max} = 1.27 \frac{\Delta m^2 (eV^2) L (km)}{\pi/2}. \quad (1.16)$$

E_{max} indicates the first oscillation maximum – the peaks are counted from right to left.

As neutrinos travel through matter, and not in a pure vacuum, it is important to consider how this matter impacts neutrino oscillations. When traveling through matter, the oscillation probabilities are altered through the MSW effect [11]. This occurs when a neutrino coherently scatters off an electron or nucleus within the Earth's crust. The equation of motion for the two flavor oscillation paradigm is written as

$$i \frac{d}{dt} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \frac{1}{2} \begin{pmatrix} -(\frac{\Delta m^2}{2E} \cos 2\theta - \sqrt{2} G_F N_e) & \frac{\Delta m^2}{2E} \sin 2\theta \\ \frac{\Delta m^2}{2E} \sin 2\theta & (\frac{\Delta m^2}{2E} \cos 2\theta - \sqrt{2} G_F N_e) \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} \quad (1.17)$$

where G_F is the Fermi constant and N_e is the electron number density of the matter. This $\sqrt{2} G_F N_e$ term is considered a potential. The effect is enhanced for neutrinos and is suppressed for antineutrinos, from Eq. 1.17. This is an important factor in long baseline oscillation experiments, such as NOvA or T2K. This effect is strongest when neutrinos travel through matter of variable density. However, because the

crust of the Earth is approximately constant, we will assume the MSW effect is negligible.

1.4 Three Flavor Oscillations

Next, we move to the complete picture of three flavor neutrino oscillation. This can be written simply as Eq. 1.18 with the full PMNS matrix (after omitting the Majorana/Dirac matrix).

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13} e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (1.18)$$

Carrying out the calculation, we obtain a single equation for the probability of oscillation from one flavor to another, $\alpha \rightarrow \beta$.

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= |\langle \nu_\beta | \nu_\alpha, t \rangle|^2 \\ &= \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-iE_i(t-t_0)} \right|^2 \\ &= \sum_{i,j} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{ij}^2 L}{2p}}, \end{aligned} \quad (1.19)$$

where we write the mass squared difference as

$$\Delta m_{ij}^2 = m_i^2 - m_j^2. \quad (1.20)$$

This mass difference quantity defines the oscillation frequency as a function of the two dependent variables L/E . Moreover, this Δm_{ij}^2 , specifically Δm_{32}^2 , is another parameter NOvA seeks to constrain and measure. In Eq. 1.21 and Eq. 1.22, we see the oscillations are driven by both Δm_{32}^2 and Δm_{31}^2 and the mixing matrix

parameters. To first order in α , the survival probabilities take a similar form of the two-flavor oscillation probability,

$$P(\nu_\mu \rightarrow \nu_\mu) \approx 1 - \cos^2 \theta_{13} \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m_{32}^2 L}{4E_\nu} \right) + \mathcal{O}(\alpha, \sin^2 \theta_{13}) \quad (1.21)$$

where $\alpha \equiv \Delta m^2 = \Delta m_{21}^2 / \Delta m_{32}^2$.

Next, the $P(\nu_\mu \rightarrow \nu_e)$ oscillations in the three flavor paradigm can be written to first order in α ,

$$P(\nu_\mu \rightarrow \nu_e) = \sin^2 \theta_{23} \sin 2\theta_{13} \frac{\sin^2 \Delta(1-A)}{(1-A)^2} + \alpha J \cos(\Delta \pm \delta_{CP}) \frac{\sin \Delta A \sin \Delta(1-A)}{A(1-A)} \quad (1.22)$$

where

$$J = \cos \theta_{13} \sin 2\theta_{13} \sin 2\theta_{12} \sin 2\theta_{23}, \quad (1.23)$$

$$A \equiv \pm 2\sqrt{2}G_F N_e E_\nu / \Delta m_{31}^2, \quad (1.24)$$

and

$$\Delta \equiv \Delta m_{13}^2 L / 4E_\nu. \quad (1.25)$$

It is important to note the second term of the $P(\nu_\mu \rightarrow \nu_e)$ probability is sensitive to CP violation, δ_{CP} . Here A describes the size of the matter effects as the neutrinos travel through the Earth, and alters the mixing angle θ_{13} . It becomes

$$\sin^2 2\theta_{13} = \frac{\sin^2 2\theta_{13}}{\sin^2 2\theta_{13} + (A - \cos 2\theta_{13}^2)}. \quad (1.26)$$

Two parameters, Δm_{32}^2 and $\sin^2 2\theta_{23}$, appear in both the appearance and disappearance probabilities. In the context of NOvA, these are the two oscillation parameters NOvA seeks to constrain. To further contextualize the significance of these two parameters within neutrino oscillations, Figure 1.4 shows the prediction of the ν_μ distribution before and after oscillations in reconstructed neutrino energy.

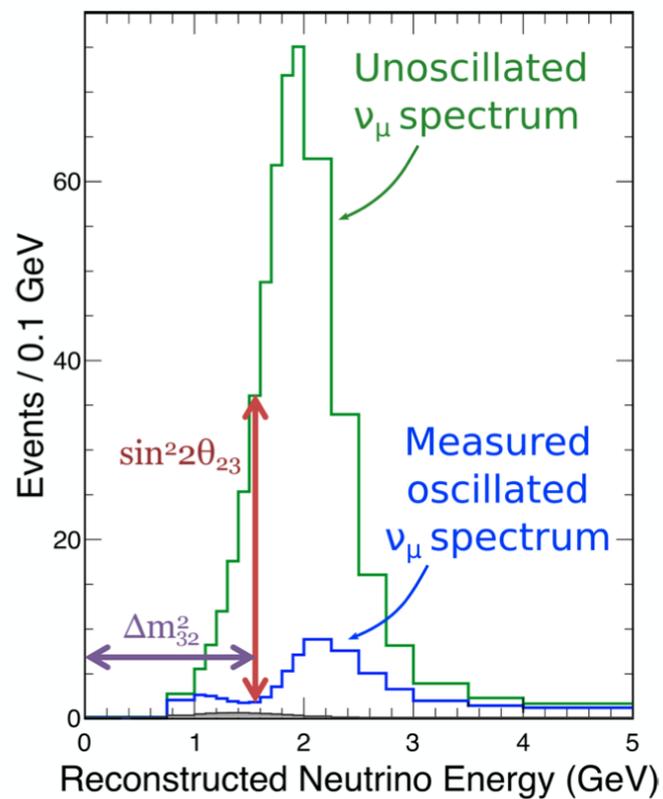


Figure 1.4: Reconstructed neutrino energy spectrum for oscillated ν_μ and unoscillated ν_μ neutrinos. The impact of $\sin^2 2\theta_{23}$ and Δm_{32}^2 sculpt the shape and normalization of the oscillated spectrum. Constraining these two parameters are NOvA's top scientific objectives. To demonstrate this effect, the two parameters are set to some arbitrary non-zero values. From [12].

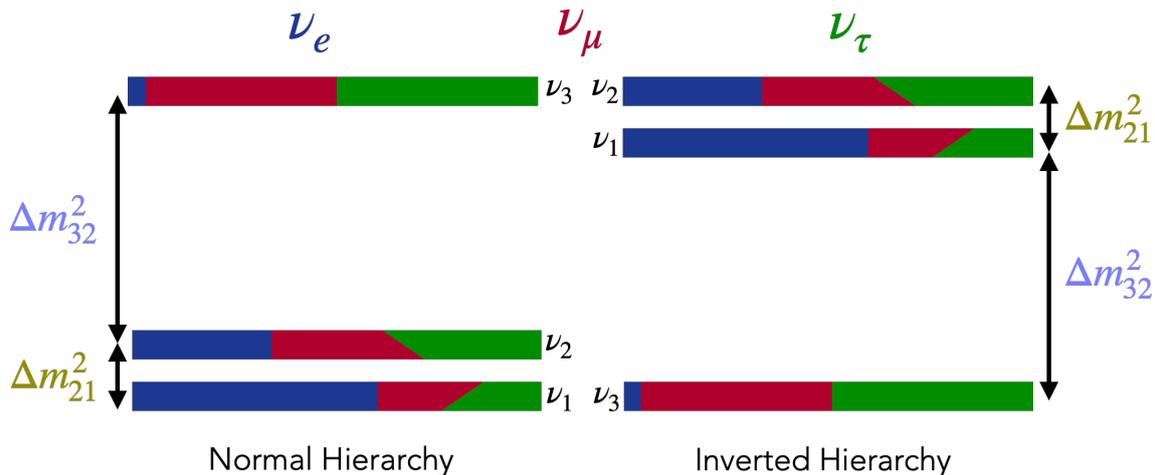


Figure 1.5: Diagram of the two possible neutrino mass orderings: normal (left) and inverted (right). In normal ordering m_3 is the heaviest mass state, while inverted ordering it is the lightest of the three. From [13].

In this figure, the blue, oscillated distribution represents the “survival” spectrum of ν_μ , where the missing neutrinos have predominately oscillated into ν_e (as well as a small fraction into ν_τ). This provides useful context into how the neutrino oscillation parameters – Δm_{32}^2 and $\sin^2 2\theta_{23}$ – are relevant to the neutrino prediction in NOvA. The value of Δm_{32}^2 determines the location of the “dip” in the neutrino energy – the point of maximal neutrino oscillation, while $\sin^2 2\theta_{23}$ determines the “amplitude” of the dip. Understanding these values as precisely as possible are crucial to understanding the phenomenology of neutrino oscillations.

1.5 Neutrino Mass Ordering

In addition to measuring the mixing parameter θ_{23} , the other main objective of the NOvA experiment is to determine the neutrino mass ordering. As we have discussed, there are three unique mass eigenstates ($m_1 \neq m_2 \neq m_3$). The masses are very small $m_i \leq 1eV$, but exact values are unknown. Again, it is the mass squared difference that appears in neutrino oscillations and can be constrained (Eq. 1.20). However, measurements of Δm_{ij}^2 indicate the mass squared difference between the eigenstates, and therefore we can learn the relative sizes of the masses eigenstates. This concept is illustrated in Fig. 1.5. There are two possible orientations. On the

left side of the figure, $m_3 > m_2$ which leads to $\Delta m_{32}^2 > 0$. This is **Normal** Ordering (also referred to as Normal Hierarchy). Conversely, if $m_3 < m_2$ then $\Delta m_{32}^2 < 0$. This is **Inverted** Ordering (Inverted Hierarchy). As a consequence, for neutrino oscillation experiments, there are only two independent mass measurements that can be made: Δm_{32}^2 and Δm_{21}^2 . As NOvA seeks to constrain the parameters of the atmospheric sector of the PMNS matrix, it cannot meaningfully constrain Δm_{21}^2 . NOvA's goal is to constrain the ordering for Δm_{32}^2 .

1.6 Constraint on δ_{CP}

One of NOvA's physics goals is to constrain the parameter δ_{CP} . This parameter quantifies the Charge-Parity (CP) symmetry violation present in neutrino oscillations. Maximal CP violation occurs at $\delta_{CP} = \pi/2$ and $3\pi/2$. To probe this parameter's value, NOvA divides its neutrino production between neutrinos and antineutrinos. Observing oscillations in neutrinos *and* antineutrinos helps to probe this matter-antimatter asymmetry more carefully. In recent NOvA results, maximal CP violation of $\delta_{CP} = 3/2\pi$ in the IO is excluded at 3.4σ and $\delta_{CP} = \pi/2$ in the NO is excluded at 2.4σ (see Figure 1.9a).

1.7 Recent NOvA Results

To summarize, the current, primary objectives of the NOvA experiment are to constrain/measure:

- θ_{23} (via ν_μ disappearance & ν_e appearance)
- Δm_{32}^2 (via ν_μ disappearance)
- δ_{CP} (via ν_e appearance).

These three items are the physics objectives of NOvA. These objectives are met from different physics; the octant (and value) of θ_{23} can be found from $P(\nu_\mu \rightarrow$

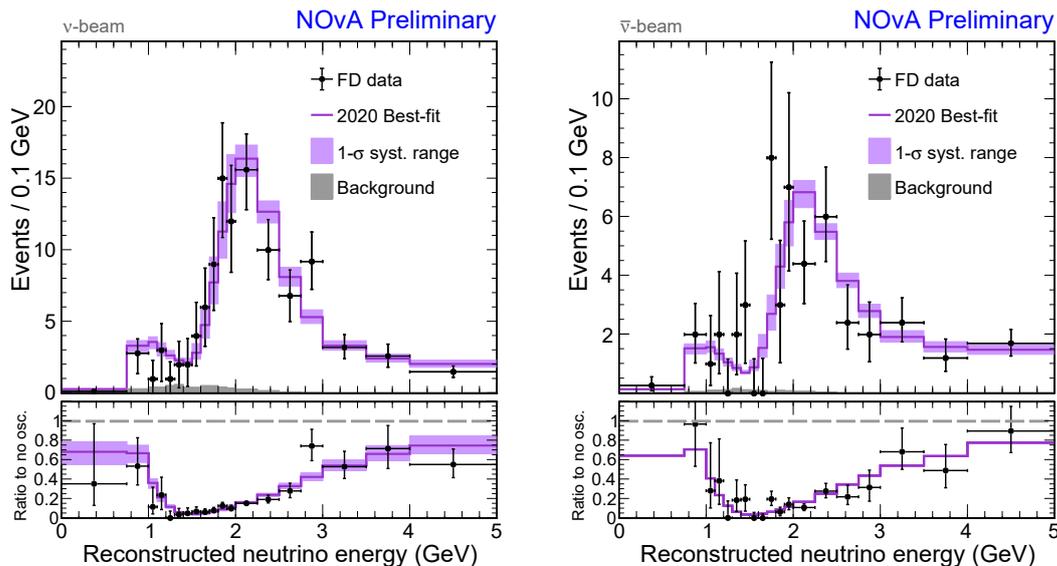


Figure 1.6: FD ν_μ spectra for neutrinos (left) and antineutrinos (right) in NOvA. From 2020 NEUTRINO conference presentation [14].

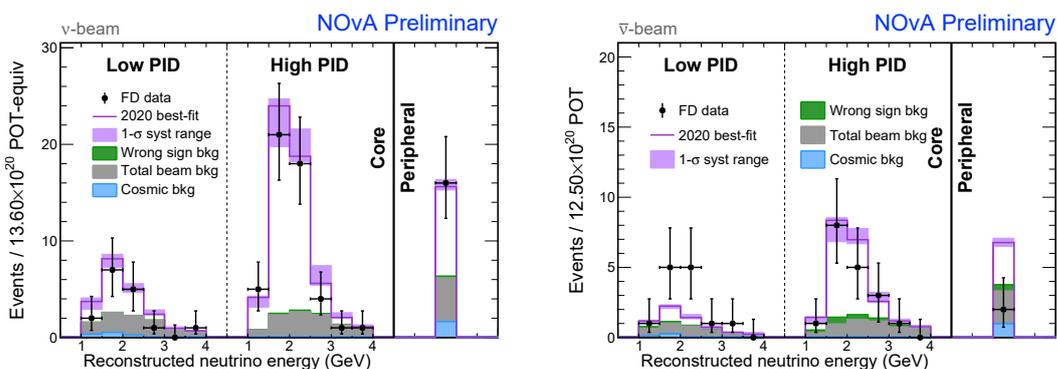


Figure 1.7: FD ν_e spectra for neutrinos (left) and antineutrinos (right) in NOvA. From 2020 NEUTRINO conference presentation [14].

ν_μ) disappearance; $\theta_{23} > 45^\circ$ is referred to as the “upper octant”, and $\theta_{23} < 45^\circ$ the “lower octant”. The mass ordering and δ_{CP} are determined from $P(\nu_\mu \rightarrow \nu_e)$ appearance. NOvA’s most recent analysis from 2020 shows the ν_μ spectra (Fig. 1.6) and ν_e spectra (Fig. 1.7) [14].

These plots summarize NOvA’s observations of neutrino interactions for ν_μ disappearance and ν_e appearance in both neutrinos and antineutrinos. These spectra will appear again in the ND+FD joint fit in Sec. 7.4 within this thesis. With these distributions of neutrino events, constraints on the oscillation parameters are made. Figure 1.8 shows NOvA’s contours for Δm_{32}^2 and $\sin^2 \theta_{23}$. This “NOvA and friends” plot (i.e. other oscillation experiments) contains 2σ contours for other neutrino oscil-

lation experiments in the Normal Ordering. NOvA's best fit point is the red circle at $\Delta m_{32}^2 = 2.41 \times 10^{-3} \text{ eV}^2$ and $\sin^2 \theta_{23} = 0.57$ and its 2σ contour in red. NOvA is one of the leading oscillation experiments that produces some of the strongest constraints on these parameters.

The current landscape of neutrino oscillations is clear from this plot too. The T2K experiment observes neutrino oscillations over a short-baseline experiment of 295 km (compared to NOvA's 810 km). However, T2K produces and observes neutrinos with the same L/E ratio from Eq. 1.14, and therefore observes neutrinos at a smaller energy. Thus, the separate measurements of the NOvA and T2K experiments strengthen the constraining power of the oscillation parameters.

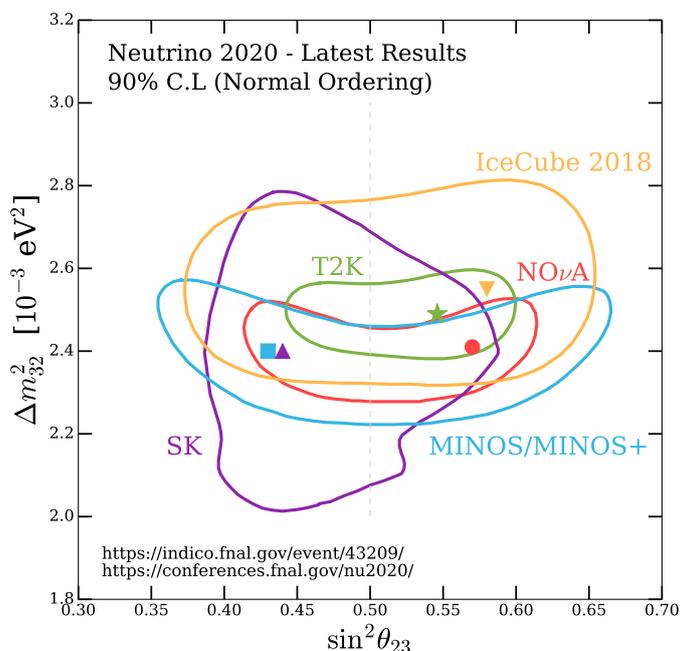
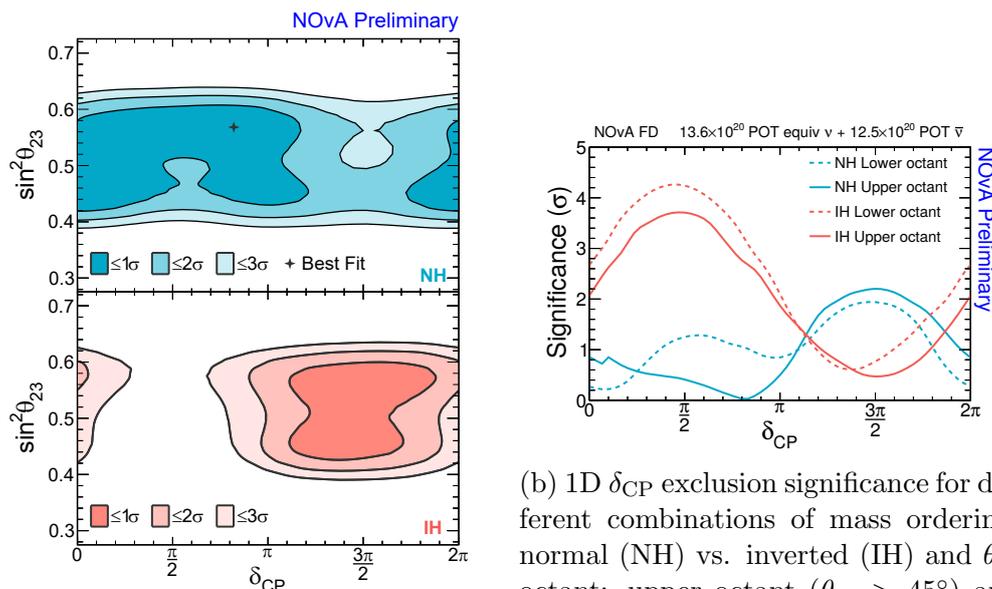


Figure 1.8: 2σ 2D contours of Δm_{32}^2 - $\sin^2 \theta_{23}$ in the Normal Ordering. Contours are from NOvA and other oscillation experiments. NOvA (in red) continues to produce one of the strongest constraints on Δm_{32}^2 and $\sin^2 \theta_{23}$.

Looking at NOvA's constraints further, Figure 1.9a shows the 2D contour of $\sin^2 \theta_{23}$ vs. δ_{CP} in both the Normal Ordering and Inverted Ordering. Again, NOvA's best fit to its data is in the Normal Ordering denoted by the black cross. Figure 1.9b shows the 1D exclusion significance for δ_{CP} in combinations of Normal and Inverted Ordering *and* upper and lower octant for θ_{23} . Fig. 1.9a & Fig. 1.9b together illustrate NOvA's constraining power of δ_{CP} especially in Inverted Ordering (IH).



(a) $\sin^2 \theta_{23}$ vs. δ_{CP} 2D contours. The lower octant ($\theta_{23} < 45^\circ$). NOvA most top plot shows the contour for Normal strongly constrains values of δ_{CP} in the Hierarchy (Ordering), bottom Inverted Inverted Ordering. Hierarchy. NOvA’s “Best Fit” point is in the Normal Ordering denoted by the cross.

Figure 1.9: NOvA 2020 results of oscillation parameters presented at NEUTRINO 2020 [14].

The “bi-probability” plot, Fig. 1.10, summarizes NOvA’s expected neutrino events to observe given different combinations of the oscillation parameters; there are four ovals to describe the possible permutations. For a single oval, the four different shaped points denote the different values of δ_{CP} . The blue and red ovals represent the mass ordering, while the the Upper Octant (UO) for θ_{23} is represented by the upper-most blue and red oval and Lower Octant (LO) is the lower-most red and blue oval. NOvA’s 2020 best fit analysis is the purple star and rests along the Normal Ordering, Upper Octant oval and is nearest to the $\delta_{CP} = \pi$ point. From this we learn that NOvA currently prefers Normal Ordering over Inverted Ordering, the Upper Octant, and also prefers a non-zero value of δ_{CP} . This is the current state of NOvA’s constraint on the oscillation parameters of the atmospheric sector of the PMNS matrix.

In the future, the next generation of oscillation experiments – DUNE and Hyper-Kamiokande – will collect more data and observe neutrino interactions with far more

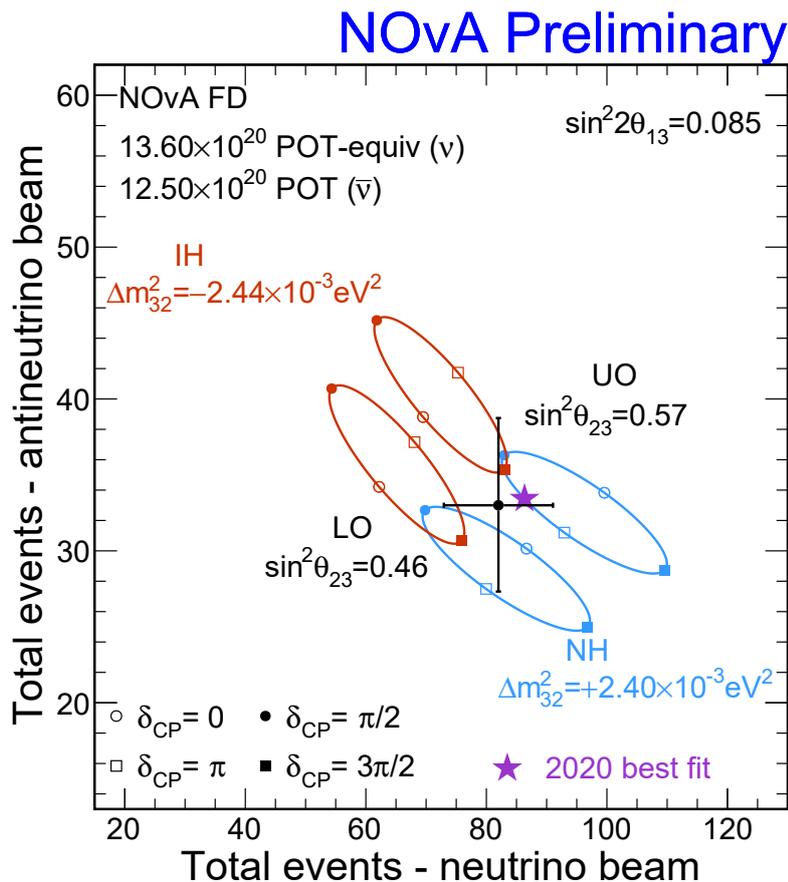


Figure 1.10: The “bi-probability” event plot for NOvA. Summarizes the number of expected neutrino events given different combinations of: δ_{CP} , mass ordering (NH vs. IH), and θ_{23} octant (upper octant – right diagonal, lower octant – left diagonal). NOvA’s 2020 best fit is marked with the purple star, and the NOvA data is marked by the black point with statistical and systematic errors [14].

sensitivity than the current generation of NOvA and T2K [15]. Improvements in detector technology, interaction reconstruction, neutrino interaction modeling, and increased production of neutrinos will make this possible. The expectation is these experiments will make contributions to make 5σ conclusions about the values of the oscillation parameters.

1.8 Significance of Thesis Measurement

This thesis reports the constraint of NOvA systematic and oscillation parameters (δ_{CP} , Δm_{32}^2 , and θ_{23}) from a simultaneous fit of the NOvA Near and Far Detector simulation to Near Detector real data and Far Detector fake data. The strong sta-

tistical power of the Near Detector data constrains the NOvA uncertainties, while the Far Detector fake data constrains the neutrino oscillation parameters. This measurement is the first of its kind; the NOvA Near and Far Detector simulation will be constrained simultaneously, as opposed to sequentially in the conventional “extrapolation” analysis, providing a novel means of interrogating the NOvA data. Furthermore, NOvA’s conventional “extrapolation” procedure of measuring the oscillation parameters uses Frequentist statistics, while this thesis performs the fit using Bayesian inference; having an alternative statistical methodology to measure the oscillation parameters will further reinforce both Frequentist and Bayesian analyses.

For the reader interested in solely neutrino oscillations, it is advised to skip to Section 7.4. However, the success of an oscillation parameter constrain requires a robust fit of the NOvA Near Detector simulation to the ND data, demanding a solid understanding of the underlying neutrino interaction models. To this end, the interested reader is advised to focus on Chapter 4 and Chapter 6 where the details of the ND data/MC selection are made, discussion of the deficiencies in the nominal NOvA MC, and the conclusions with respect to the interaction models and uncertainties are made from the MCMC sampling.

Chapter 2

Neutrino Interactions

The NOvA experiment is unique in that it observes nontrivial amounts of different scattering processes in the 2-3 GeV energy range. This chapter breaks down the underlying physics models present within the NOvA simulation, necessary to extract the oscillation parameters. First Section 2.1 describes commonly used kinematic variables. Section 2.2 discusses the underlying nuclear physics model, Sections 2.3 (Quasi-Elastic), 2.4 (2p2h scattering/Meson-Exchange Current), 2.5 (Resonant), and 2.6 (Deep Inelastic Scattering) discuss the four dominant neutrino interactions observed in NOvA. Section 2.7 discusses a physics process relevant to measuring neutrino oscillation parameter, Final State Interactions. Section 2.8 describes other models used in the NOvA simulation. Section 2.9 describes the simulation and neutrino interaction generator used for this analysis. A summary of the models used in the NOvA simulation is listed in Table 2.1.

Physics Content	Model Used
Nuclear Ground State	Local Fermi Gas (LFG) [16]
Quasi-Elastic Scattering	Valencia 1p1h, Z-Exp. Axial Form Factor [17]
2p2h Scattering	Valencia 2p2h [18]
Resonant Scattering	Berger-Seghal [19]
Deep Inelastic Scattering	Bodek-Yang [20]
Intranuclear Final State Interactions	hN, Semi-Classical Cascade [21]

Table 2.1: Important physical processes and the corresponding models that comprise NOvA's P Rod5.1 MC simulation.

The cross section ($[\sigma] = cm^2$) of a high energy particle describes the probability

of an interaction with matter. To successfully reconstruct neutrino interactions in the NOvA detector, it is crucial to accurately simulate the various types of neutrino interactions. There are four fundamental charged-current (CC) neutrino interaction processes that occur within NOvA: Quasi-Elastic Scattering (QE), 2p2h Scattering, Resonant (RES), and Deep Inelastic Scattering (DIS). The neutrino cross section on free isoscalar targets for neutrinos (left) and antineutrinos (right) are plotted as a function of the true neutrino energy in Figure 2.1. These plots only shows

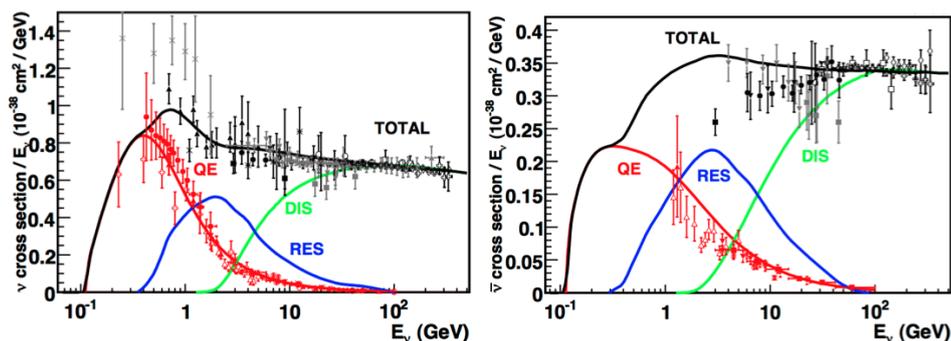


Figure 2.1: Neutrino charged-current (CC) cross section, σ / GeV , on free isoscalar targets for QE (red), RES (blue), and DIS (green) in neutrinos (left) and antineutrinos (right). The total cross section is represented in black. NOvA receives a narrow-energy band of neutrinos centered at approximately 2 GeV. From [22].

the cross section on free nucleons, so the 2p2h contribution is not shown. The black line and dots (prediction and data, respectively) represent the total neutrino cross section. Note the factor of ≈ 3 in relative size of neutrino cross section when compared to antineutrino. The red shows the predicted and observed cross sections for Quasi-Elastic (QE) interactions which grows until approximately 1 GeV and then plateaus. The blue shows the Resonant (RES) interactions centered at roughly 2 GeV; note how these two interactions are probable within this energy range. The Deep Inelastic Scattering (DIS), green, interactions occur at higher neutrino energies > 2 GeV. NOvA receives a beam of neutrinos centered tightly around 2 GeV, and receives most of its neutrinos within the energy range of 1.5 – 2.5 GeV. NOvA is sensitive to all three of these neutrino interaction types – and because NOvA detector is primarily made of carbon, the 2p2h cross section is significant too – therefore, NOvA observes large populations of events from all four of these CC

interaction types, illustrating the need for NOvA's underlying physics simulation to be as accurate as possible in all of these models. The following sections will discuss these models in more detail.

2.1 Kinematic Variables

In this discussion, it is important to consider the kinematics most commonly used when attempting to estimate the neutrino energy in NOvA. Moreover, these quantities are often useful to help characterize interactions from the four common interaction models. To start, the neutrino energy is estimated by

$$E_\nu = E_\mu + E_{had}, \quad (2.1)$$

where E_μ is the energy of the muon and E_{had} is the hadronic energy, which is all the energy observed in the NOvA detector not associated with the lepton. Each of these models will predict varying amounts of E_{had} that can more directly impact the accuracy of estimating the neutrino energy.

One common kinematic quantity is

$$Q^2 = -q^2 = |\vec{q}^2| - q_0^2, \quad (2.2)$$

where Q^2 is the negative of the dot product of the four-vector q ,

$$q = (q_0, \vec{q}). \quad (2.3)$$

The four-momentum transfer, q , is an extremely important quantity in neutrino interactions, which contains the energy transferred to the hadronic system, q_0 , and the momentum \vec{q}_3 transferred to the hadronic system.

This Q^2 can be calculated from NOvA observables,

$$Q^2 = 2E_\nu(E_\mu - p_\mu \cos \theta_\mu) - m_\mu^2, \quad (2.4)$$

where p_μ is the momentum of the muon, $\cos \theta_\mu$ is the angle the muon makes with the neutrino direction, and m_μ^2 is the mass of the muon. This quantity is commonly used when discussing nuclear and neutrino interaction models.

The quantity W is also useful; it is the invariant hadronic mass from the by-products of the interaction. This quantity is often helpful to separate QE interactions from RES or DIS interactions, as higher W is typically associated with RES and DIS interactions.

The Bjorken inelasticity,

$$y = E_{had}/E_\nu, \quad (2.5)$$

is the fraction of the hadronic energy with respect to the total (neutrino) energy. Values of $y \approx 1$ indicate the interaction is highly inelastic, where nearly all of the neutrino energy went into the hadronic system. A high y is frequently a DIS event. In NOvA, this quantity is often referred to as the hadronic “energy fraction”. These are all quantities that will be referenced later in the analysis.

2.2 Nuclear Model

The nuclear model describes the ground state density and momentum distributions of the nucleons. This is a highly nontrivial problem to model; the nucleons within a nucleus are not stationary and bound together from the strong force. When a nucleon is ejected, it is not simply a nucleon at rest. One must account for the exiting nucleon’s initial energy and momentum.

In NOvA’s prior simulation production campaigns, the Relativistic Fermi Gas (RFG) model has been used. This treats the nucleons within the nucleus as a potential well; if a nucleon were to be ejected, it would contain some initial momentum. This model has been useful for describing 2-particle 2-hole (2p2h) interactions (of which Meson Exchange Current interactions are a large contribution), given its high-momentum tail above 220 MeV of momentum (red, Figure 2.2). The high momentum tail was an ad-hoc correction to account for short range correlations (SRC)

of nucleons in the nucleus.

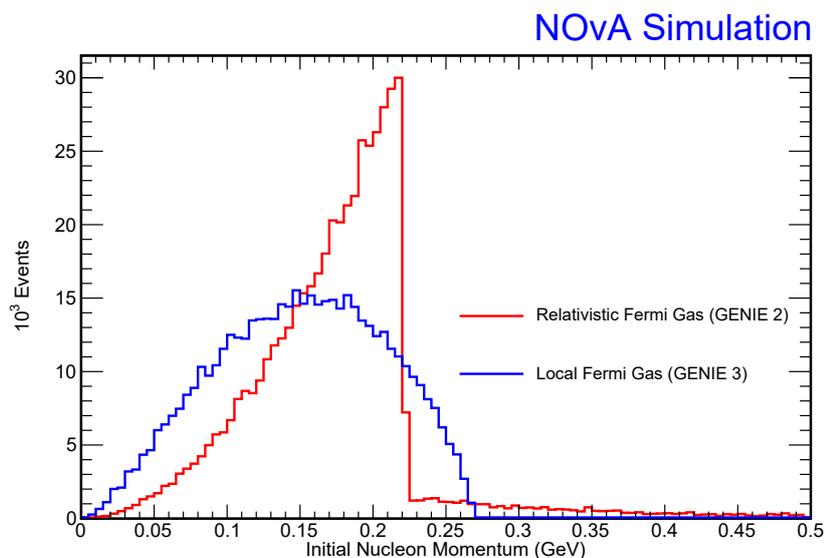


Figure 2.2: Two of the commonly used initial nucleon momentum distributions: the Nuclear Local Fermi Gas (LFG model, blue), and the Relativistic Fermi Gas (RFG, red). This analysis uses the LFG model to describe the nucleon momentum within the nucleus.

For this analysis, NOvA has moved to the Local Fermi Gas. As it is still a Fermi Gas model, many of the features remain unchanged. However, a significant change is the difference in the high-momentum tail (blue Fig. 2.2), which is predicted in the model and not an ad-hoc correction. In the LFG model, Random Phase Approximations (RPA) are added into the model to account for long-range correlations of nucleons within the medium. In addition, the LFG includes Pauli blocking and removal energy. It has become the preferred nuclear model of choice among neutrino oscillation experiments.

The removal energy is an important uncertainty in reconstructing the hadronic energy [23]. As stated, the nuclear medium is considered a potential well for a nucleon to escape. Therefore, the struck nucleon must have a minimum energy threshold to be knocked free. This energy is not accounted for within any kinematic variables and therefore will impact the accuracy the neutrino energy estimate. The impact of this uncertainty has been studied in NOvA [24].

2.3 Quasi-Elastic Scattering

Quasi-elastic (QE) scattering is a near-elastic interaction where an incoming neutrino (antineutrino), exchanges a virtual W^\pm boson with a bound neutron (proton) within the nucleus. This neutron (proton) is converted into a proton (neutron). Fig. 2.3 shows the Feynman diagram for this process. The struck nucleon contains some initial momentum (via the initial nucleon momentum distribution, Fig. 2.2) and exits the nucleus. Because this bound nucleon contains some initial momentum, the nuclear model is closely connected to QE interactions, thus the outgoing nucleon in this process will typically reflect the initial nucleon momentum distribution. A QE interaction is characterized by little energy/momentum being transferred into the hadronic system. This highlights why the nuclear medium effects are relevant for these types of interactions as nuclear medium effects non-trivially contribute to the small amount of observed hadronic energy. This hadronic component of the interaction in the detector is a small proton track for neutrinos. For antineutrinos, the neutron does not leave a track due to its neutral charge and only charged secondary particles from neutron activity are observed in the detector (generally some distance from the interaction vertex). Meanwhile, the outgoing muon (μ^+ , or μ^-

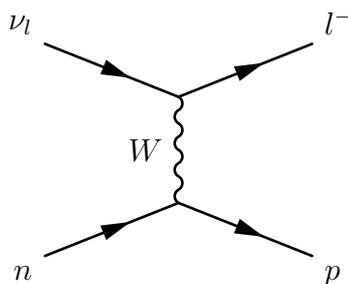


Figure 2.3: Quasi-Elastic neutrino scattering. An incoming neutrino exchanges a virtual W^+ with a bound nucleon, where an outgoing lepton (l^-) and p are observed.

for antineutrino) tends to deposit a long, clean track approximately 5 m long and generally in line with the neutrino beam direction.

The Valencía 1-particle, 1-hole model is employed in NOvA, which accounts for several nuclear effects: long range RPA correlations between nucleons, Final State Interaction (FSI) effects, and Coulomb potential corrections [23].

The free nucleon cross section depends primarily on two nucleon form factors which are a function of Q^2 . The vector form factor is fully determined in electromagnetic processes such as electron scattering, but the axial form factor must be measured in weak processes and therefore contains some uncertainty. A common parametrization is the dipole form

$$F_A(Q^2) = F_A(0) \left(1 + \frac{Q^2}{M_A^2}\right)^{-2}, \quad (2.6)$$

where $F_A(0)$ is the axial form function at $Q^2 = 0$, Q^2 is the kinematic variable Eq. 2.4, and M_A is the axial mass. Within QE interactions the value – and uncertainties – of the axial mass, M_A , plays a large role in shaping the prediction of these interactions. For this analysis, the axial current propagator’s mass is modeled as a Taylor expansion of the axial form factor in Q^2 and some auxiliary parameter t , t_0 , parametrized as $z(Q^2, t, t_0)$ [25]. This Z-expansion parametrization diverges from the historical dipole expansion uncertainty that has been used. The motivation for doing so provides a more conservative uncertainty on the axial mass as well as an increased flexibility in describing the shape of the axial form factor Q^2 dependence.

2.3.1 Quasi-Elastic Uncertainties

From Table 2.1, the neutrino physics model used is the València 1p1h model with the axial form factor expressed as a Z-expansion. The uncertainties associated with the Z-expansion are, `ZExp_AxialFFSyst2020_EV[1-4]`. These uncertainties are approximately 10–20% from the nominal prediction at $\pm 1\sigma$ in the region of $Q^2 \approx 1$ GeV [26]. This can be seen for the first eigenvector in Figure 2.4. In addition to these uncertainties, there is also a total QE normalization systematic that controls the total QE events by +20%/–15%, `ZNormCCQE`.

There are two uncertainties associated with the Random Phase Approximations within the nuclear model. These uncertainties impact QE interactions and are included here: RPA shape suppression and RPA shape enhancement. Again, the

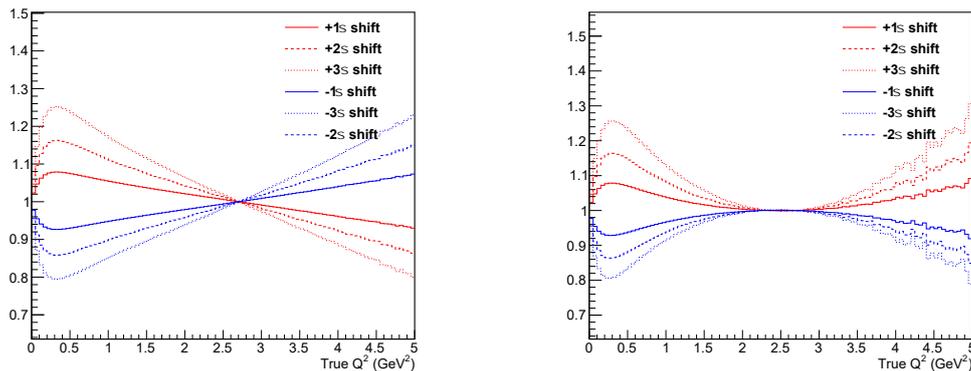


Figure 2.4: Effect of ZExp_AxialFFSyst2020_EV1 in Q^2 for neutrinos (left) and antineutrinos (right) in NOvA simulation. The majority of NOvA neutrino interactions occur at $Q^2 < 2 \text{ GeV}^2$.

Random Phase Approximation (RPA) in quasi-elastic interactions describes the random motion of the struck nucleon within the nucleus. This effect is evident in low Q^2 interactions, where nuclear effects – such as this one – are present. There are two systematics that account for this phenomenon: RPACCQESuppSyst2020 and RPACCQEEEnhSyst2020 that suppress and enhance this effect. The enhancement and suppression for these systematics at $\pm 1\sigma$ is approximately 10% in true q_0 and Q^2 .

2.4 2p2h/Meson Exchange Current

The 2-particle 2-hole scattering is another form of neutrino interaction. In terms of the physics modeling, Meson Exchange Current (MEC) interactions are subset of 2p2h scattering mechanisms. However, a crucial distinction must be made: within NOvA’s neutrino interaction generator GENIE, the **2p2h model is referred to as Meson Exchange Current interactions**. Therefore, henceforth 2p2h and MEC will be used synonymously, and will utilize the name MEC as investigations into the GENIE simulation will be performed in later chapters.

The Valencía 2p2h is the model used for this interaction type. MEC interactions describe an interaction where, aside from the outgoing lepton, sees two nucleons ejected – thus the “2 hole” left in the nucleus– from the nucleus via the virtual exchange of a pi meson among the two nucleons, Fig. 2.5.

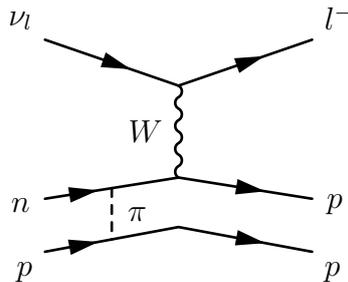


Figure 2.5: MEC interaction. A virtual pion is exchange among the two nucleons and are emitted from the nucleus.

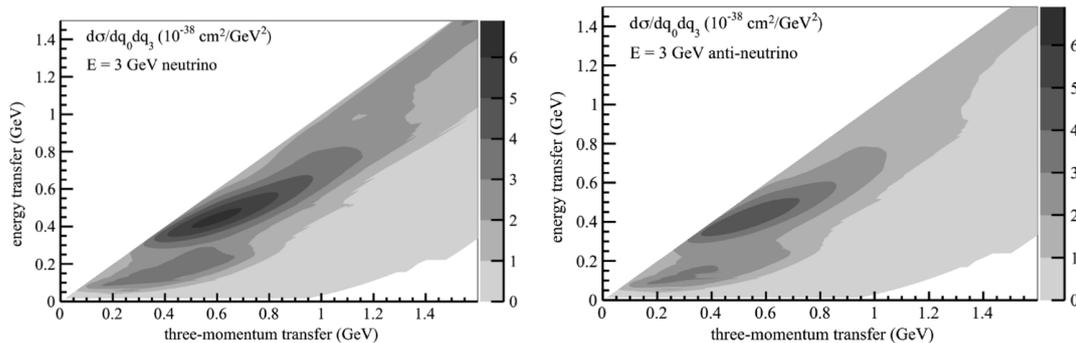


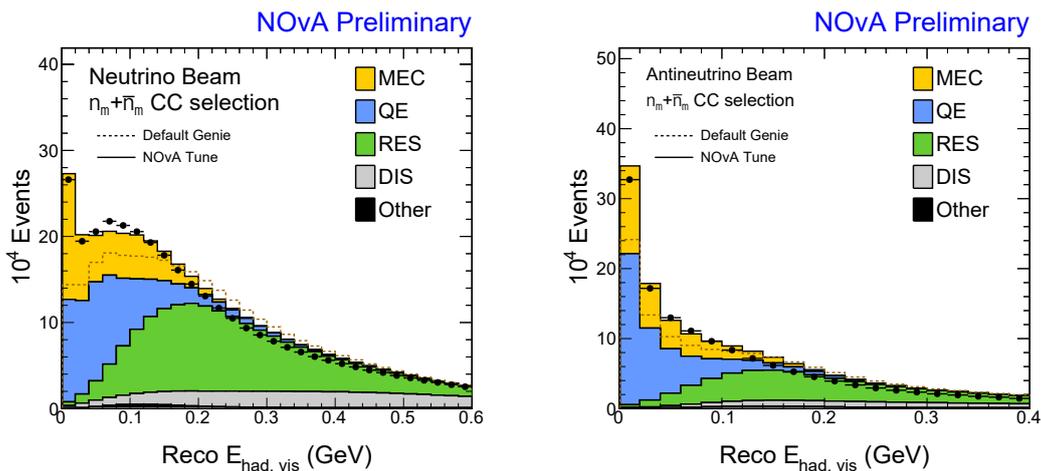
Figure 2.6: Valencía model 2p2h cross section in true $(|\vec{q}|, q_0)$ space for neutrinos (left) and antineutrinos (right) at 3 GeV. This plot is the cross section for all 2p2h processes, and not exclusively MEC.

The Valencía model is a semi-relativistic theory [18]. The model predicts two peaks in the true $(q_0, |\vec{q}|)$ space, as seen in Figure 2.6 of the plot of the differential cross section. The largest peak occurs at approximately 0.4 GeV of energy transfer, while the second peak occurs at 0.1 GeV of energy transfer. As this figure shows, the cross section is largest at lesser values of q_0 and \vec{q}_3 , therefore, this type of interaction is most commonly overlaps with QE interactions. Such an example would be a MEC neutrino interaction that produces a neutron and proton, because we do not observe the neutron, this interaction could be interpreted as a QE interaction. Thus, it is crucial to understand the QE and MEC simulation as precisely possible.

2.4.1 MEC model tuning in NOvA

NOvA's base prediction produces poor agreement with the ND data; NOvA tuning of the MEC model is performed to improve agreement with the data. Figure 2.7 shows the prediction for hadronic visible energy for neutrinos (left) and antineutrinos

(right). The poor agreement of the base prediction to the data is seen in the dotted line in, the “Default” GENIE prediction, while the black dots represent the data. The discrepancy is maximal at the lowest hadronic energies, in particular at the lowest bin of hadronic energy, where the disagreement between the data and MC is nearly a factor of 2. This discrepancy motivates the need for a tune of NOvA’s MEC model. The NOvA simulation is tuned to the NOvA inclusive ND data. The tune is



(a) NOvA ND data and MC predictions of neutrinos in variable E_{had}^{vis} . (b) NOvA ND data and MC predictions of antineutrinos in variable E_{had}^{vis} .

Figure 2.7: Hadronic visible energy in the NOvA ND for the data (black dots) and the MC prediction, broken down by true neutrino interaction. Before the MEC tuning, the MC predicts a small amount of MEC interactions – up to the dashed line (“Default GENIE”). After tuning in true $(|\vec{q}|, q_0)$, (“NOvA Tune”, solid), the MEC tuned prediction produces a stronger agreement with the data. Note low E_{had}^{vis} is the most discrepant region that motivates the MEC tune.

performed in the same two variables as the 2p2h Valencía model: true $(|\vec{q}|, q_0)$. This model, from Figure 2.6, resembles two peaks in the cross section – one at 0.1 GeV and the other at 0.4 GeV of q_0 . The MEC tuning is done in a manner that reflects the original model; a weighting function is made in true (\vec{q}_3, q_0) space for true MEC events, where events are scaled up or down based on the disagreement

$$\text{weight}_{MEC} = \text{baseline} + \text{gauss}_1(|\vec{q}|, q_0) + \text{gauss}_2(|\vec{q}|, q_0). \quad (2.7)$$

The weighted parameters are two ellipsoid shapes containing 13 parameters: 6 parameters for a 2-dimensional Gaussian distribution at ≈ 50 MeV of true q_0 (“Gaus-

sian 2”) and the other 2D Gaussian distribution of 6 parameters at ≈ 350 MeV of true q_0 (“Gaussian 1”), and 1 parameter to control the total normalization. These parameters are adjusted to agree with the data in E_{had}^{vis} and $Reco |\vec{q}|$ using a MINUIT χ^2 minimization technique [27]. Figure 2.8 shows the values of the weights produced from adjusting the 13 parameters in the MINUIT minimization. Note the weights are significantly larger for low values of true energy transfer, weighting up and improving the disagreement observed in the base prediction at low E_{had}^{vis} in Fig. 2.7. The plots of the tune in the other variable, $Reco |\vec{q}|$, can be seen in the Appendix A.1. The resulting MEC prediction (in yellow) is improved significantly

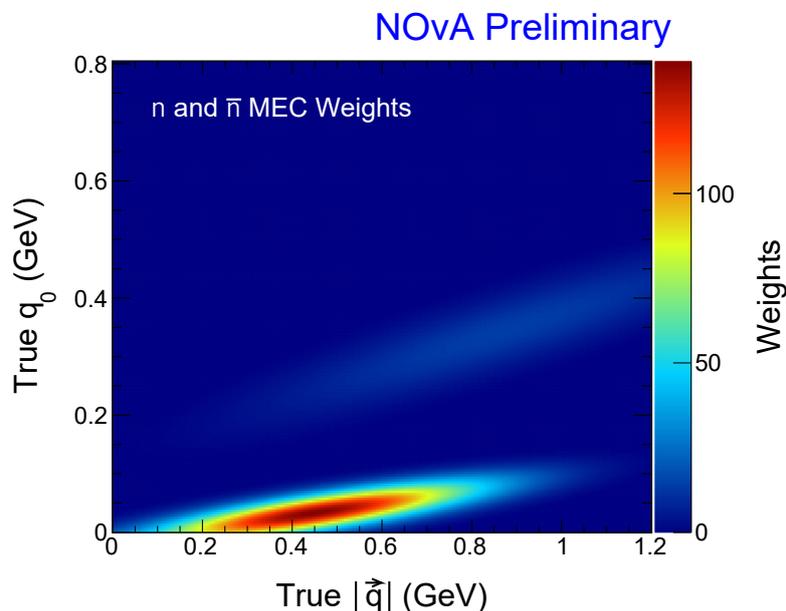


Figure 2.8: The tuning of the 2p2h/MEC contribution in NOvA. The χ^2 minimization produces large weights for low true q_0 events. The smaller-weighted Gaussian 1 is centered at (0.8 GeV, 0.3 GeV) in the true (\vec{q}, q_0) space. The larger weighted Gaussian 2 is centered at (0.4 GeV, 0.05 GeV). The enhancement of MEC events is clearly dominated at very low true energy transfer.

– from the dashed line to the solid black line – to the data in Fig. 2.7. The exact values of the thirteen parameters from the tuned MEC model can be found in Table 2.2.

Component	Parameter	Fitted value	Systematic Name
Gaussian 1	Normalization	15.03	..._Norm_1
	Mean q_0	0.34	..._MeanQ0_1
	Mean $ \vec{q} $	0.88	..._MeanQ3_1
	Sigma q_0	0.10	..._SigmaQ0_1
	Sigma $ \vec{q} $	0.39	..._SigmaQ3_1
	Correlation	0.90	..._Corr_2
Gaussian 2	Normalization	140	..._Norm_2
	Mean q_0	0.03	..._MeanQ0_2
	Mean $ \vec{q} $	0.45	..._MeanQ3_2
	Sigma q_0	0.04	..._SigmaQ0_2
	Sigma $ \vec{q} $	0.26	..._SigmaQ3_2
	Correlation	0.81	..._Corr_2
Base model	Normalization	-1.02	..._Baseline

Table 2.2: MEC weights parameterized as two, 2D Gaussian distributions. Note the large normalization value of the Gaussian 2 parameter associated with low hadronic energy. All systematics have `MECDoubleGaussEnhSystGSFProd5p1` prepended to the **Systematic Name**.

2.4.2 Meson Exchange Current Uncertainties

The uncertainty model for MEC interactions is complex. The 13 degrees of freedom from the MEC tune (Table 2.2) are also systematic uncertainties for the MEC model in NOvA’s simulation. Therefore, these parameters will be utilized in the Markov Chain Monte Carlo sampling as well to constrain the ND simulation to the data.

In addition, there are MEC uncertainties correspond to the neutrino energy dependence of the MEC cross section. These are referred to as `MECEnuShapeSyst2020Nu` and `MECEnuShapeSyst2020AntiNu`. This uncertainty uses a Landau distribution for the upper bound uncertainty with a most-probable value = 0.4 and $\sigma = 1$. The lower bound is determined by the function

$$f(E_\nu) = \frac{0.5}{1 + 2(E_\nu - 0.25)}. \quad (2.8)$$

The largest spread occurs at low $E_\nu < 1$ GeV and diminishes with increasing energy [26].

Another set of uncertainties controls the neutron-proton nucleon pair fraction.

These uncertainties are `MECInitStateNPFrac2020Nu` and `MECInitStateNPFrac2020AntiNu`. These systematics modulate the probability the neutrino (antineutrino) interacts with a neutron-neutron (proton-proton) or neutron-proton pairs within the nucleus. For neutrinos this fraction is

$$\frac{np}{np + nn} = 0.69 \begin{cases} +0.15\sigma \\ -0.05\sigma \end{cases}, \quad (2.9)$$

and for antineutrinos

$$\frac{np}{np + pp} = 0.66 \begin{cases} +0.15\sigma \\ -0.05\sigma. \end{cases} \quad (2.10)$$

Note the uncertainty is asymmetric in $+1\sigma$ and 1σ and for neutrinos and antineutrinos. More details can be found in the NOvA cross-section tuning 9 technical note [26].

2.5 Resonant Scattering

The Resonant (RES) neutrino model is the Berger-Seghal model [19]. In a RES interaction, the W boson exchanged has enough energy to create a nuclear resonance, which decays into a π and a nucleon. This model includes the production of many resonances, however, the most common resonance seen at NOvA's relevant neutrino energies is the Δ resonance, Figure 2.9.

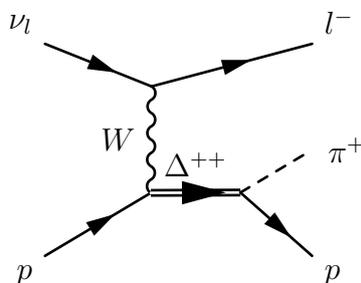


Figure 2.9: RES neutrino interaction. A Δ^{++} resonance is produced, and decays into a π^+ and proton.

Again, as seen in Figure 2.1, the RES model is a large component of the neutrino

interactions observed in NOvA's neutrino energy range.

2.5.1 Resonant Uncertainties

There are several RES uncertainties present within NOvA's uncertainty set. We have the two systematics `MaCCRES` and `MvCCRES`. These two modulate the value of axial and vector mass in the RES cross section. A $\pm 1\sigma$ shift corresponds to as large as 20% shift at true $q_0 \geq 0.8$ GeV. These two systematics are some of the leading sources of uncertainty for measuring the neutrino oscillation parameters.

There is an uncertainty `LowQ2RESSupp` that suppress RES neutrino interactions at low- Q^2 . This originates from results on the MINOS and MINERvA experiments which observed fewer RES events due to nuclear effects, where the initial momentum of the nucleon and the binding energy reduces the events that were observed in these experiments [28]. This systematic is a function of the form

$$1 - Ae^{-bQ^2}, \quad (2.11)$$

where $A = 0.391$ and $b = 12.9$ for neutrinos and $A = 0.429$ and $b = 20.9$ for antineutrinos. This uncertainty is valid for $[0\sigma, +1\sigma]$ only, where $+1\sigma$ suppresses RES interactions by a factor of 10% at $Q^2 = 0.1$ GeV².

The systematic `Theta_Delta2Npi` is also a one-sided systematic that describes the decay of resonances; it controls the direction of the Adler angle, which is the angle between the decaying Δ and p . This systematic produces more forward decays (i.e. in the momentum transfer direction) at $+1\sigma$ and at 0σ the angle between the two particles is fully azimuthal [29].

2.6 Deep Inelastic Scattering

Neutrino Deep Inelastic Scattering (DIS) processes are modeled via the Bodek-Yang formulation, Fig. 2.10 [20]. This model exchanges a W boson and has enough energy to interact with a single quark in a nucleon. The remaining quarks undergo

hadronization, where they interact via the strong force and form hadrons, such as pions.

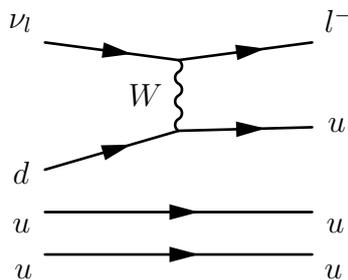


Figure 2.10: DIS Feynman diagram of neutrino interaction. The W boson has enough energy to strike a single quark within a nucleon. The remaining quarks undergo hadronization via the strong force and are observed as hadrons.

2.6.1 Deep Inelastic Scattering Uncertainties

There are several uncertainties associated with Deep Inelastic Scattering (DIS). There is an uncertainty associated with each final state pion channel, for neutral and charged current interactions and for struck nucleon of proton or neutron; these uncertainties have the name: DIS [v, vbar] [p, n] [N, C] C [0, 1, 2, 3] pi for neutrinos and antineutrinos, totaling 32 uncertainties. These adjust the normalization by 50% for DIS events above and below $W = 1.7$ GeV.

There is an additional systematic, FormZone2020, that controls the distance of the formation zone – the distance hadronization occurs within the nucleus. The uncertainty is 50% of its nominal value.

2.7 Final State Interactions

As the aforementioned models suggest (in particular the QE model), the nuclear medium plays an important role in reconstructing neutrino interactions, and therefore, measuring the neutrino oscillation parameters. Thus, modeling how hadrons interact with the nuclear medium as they exit the nucleus from a neutrino scatter is equally critical as the neutrino interaction model. This process of hadrons interacting within the nucleus as they exit is known as Final State Interactions (FSI). An

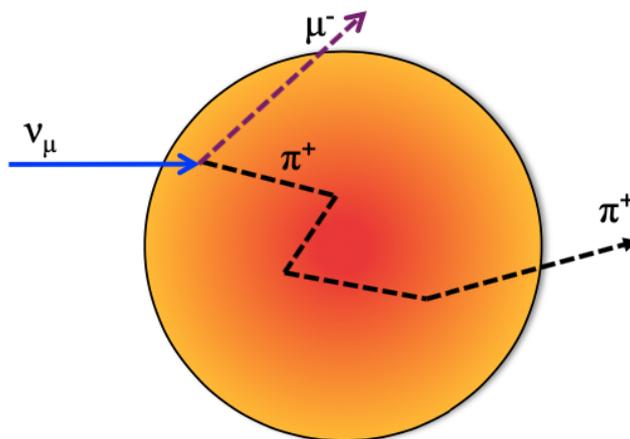


Figure 2.11: Illustration depicting the process of Final State Interactions. A struck nucleon produces a Δ^{++} resonance and quickly decays into a pion. The pion must first exit the nucleus before it is observed. As it traverses the nuclear medium, the pion undergoes scatters before it exits. These scatters alter the original pion's energy impacting the accuracy of neutrino energy estimate, and thereby the oscillation parameters. From [30].

illustration representing FSI is shown in Fig. 2.11. The π^+ scatters several times off of other nucleons, either gaining or losing kinetic energy before it exits the nucleus and is observed in a detector. Therefore the energy of the π^+ that is observed is not *explicitly* from the neutrino interaction, but now convolved with the multiple scatters, thus impacting the accuracy of the neutrino energy estimate [21].

The FSI model used in NOvA, is the hN intranuclear cascade Oset *et. al* model, which utilizes pion scattering quantum mechanical amplitudes and applies them to intranuclear pion scattering [31]. Therefore this hN model of FSI only applies to pions re-interacting within the nucleus.

There are three dominant channels of pion intranuclear interactions distinguished by the topological final state (assuming an initial state of a single pion produced in the nucleus):

Absorption (ABS) No pions are observed in the final state.

Charge exchange (CX) A single π exchanges charge by 1, and no other pions, are observed in the final state.

Quasi-elastic scattering (QE) A single pion is observed in the final state, with

Process	Parameter	Adjustment
Absorption	f_{ABS}	+40%
Charge Exchange	f_{CX}	-30%
Quasi-Elastic	f_{QE}	-10%
Mean-Free Path	f_{MFP}	-40%

Table 2.3: Adjustment of the parameters for NOvA’s hN FSI central-value tune.

same sign as the incoming pion.

The predictions of the hN model in these three categories – in addition to the cumulative total cross section – produce poor agreement to external pion scattering data. NOvA tuned the model in these three channels to improve agreement and construct uncertainties for the pion scattering processes [32]. It is important to note that NOvA’s hN FSI tune does not include any tuning of nucleons that re-interact within the nucleus; this model is a prescription for FSI with pions only.

2.7.1 FSI Tuning

The default hN FSI model poorly agrees with the external pion scattering data. The NOvA tune of the hN FSI model involves adjusting the three FSI physics processes: absorption (f_{ABS}), charge exchange (f_{CX}), and quasi-elastic scatter (f_{QE}) – and the pion mean free path (MFP) – to obtain the desired agreement. The result of the tuning can be seen in Figure 2.12 and the summary of the adjustments of these parameters is in Table 2.3. We see from this figure the need for tuning, especially in the absorption channel of pion scattering, where the default model (green) is under-predicted by nearly a factor of 2 near 200 MeV of π_{KE} . Moreover, we see the final, tuned prediction (red) produces stronger agreement in the the three channels ABS, CX, and QE, and the REAC channel, which is the sum of all pion scattering channels. With this new tune, we must construct uncertainties for the four parameters.

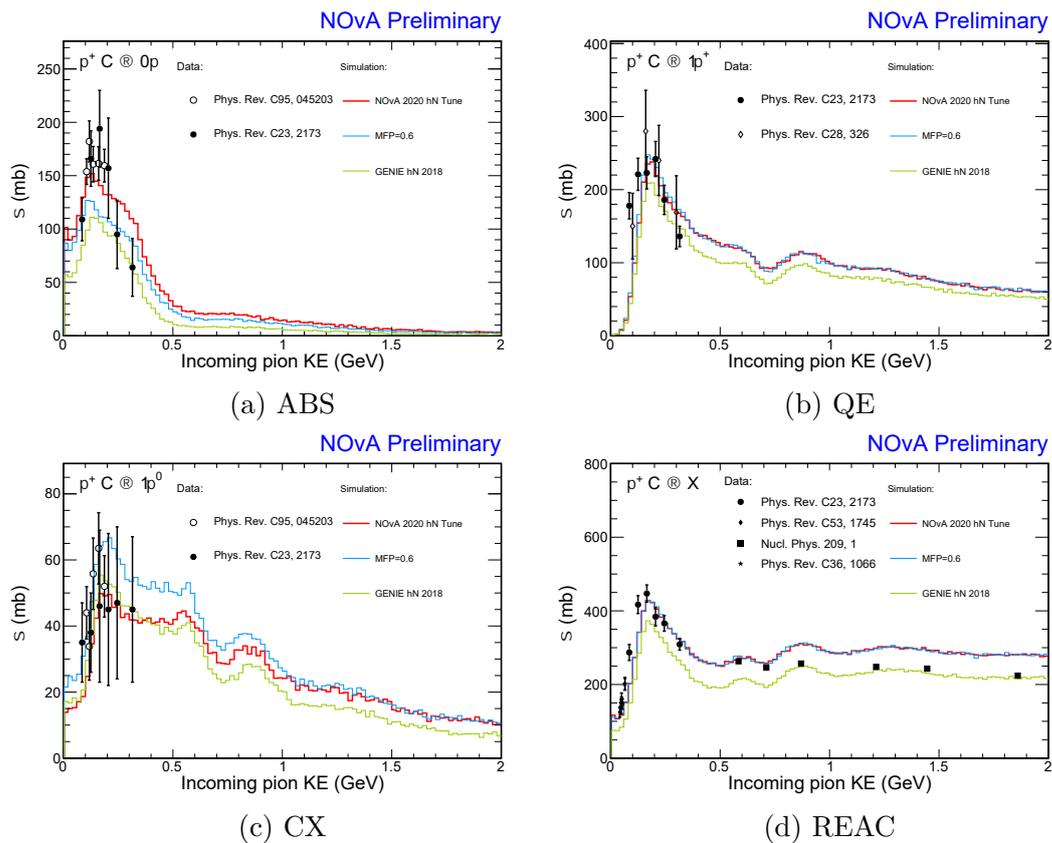


Figure 2.12: Comparison of $\pi^+ - {}^{12}\text{C}$ scattering data to predictions from nominal simulation (green), varying only the mean free path (blue), and the final central-value tune with parameters given in tab (red). 2.3. Each plot considers one scatterer channel. From [32].

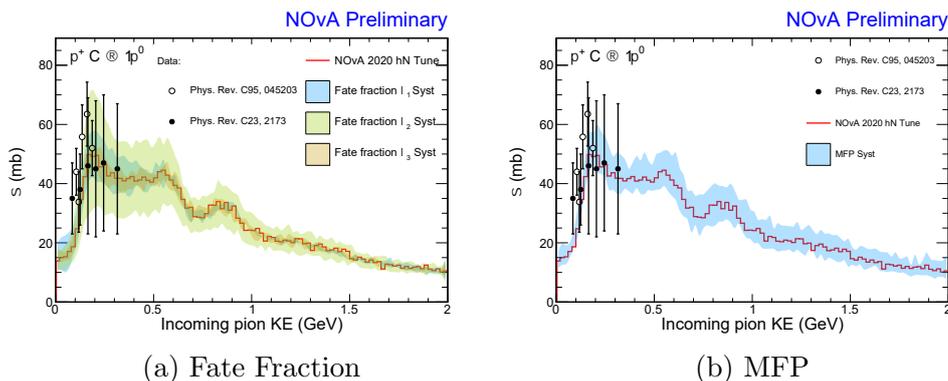


Figure 2.13: Error band of the fate fraction uncertainties (left) and MFP uncertainty (right) on the CX pion scattering channel. From [32].

2.7.2 Final State Interaction Uncertainties

There are four FSI uncertainties associated with NOvA's tuned hN model. The three uncertainties associated with the pion scattering processes are linear combinations of the three FSI process parameters: absorption (f_{ABS}), charge exchange (f_{CX}), and quasi-elastic scatter (f_{QE}). A correlation matrix of these three parameters is diagonalized to obtain its eigenvalues and eigenvectors. The uncertainty is computed by adding (subtracting) the product of the eigenvalue and the fate fraction linear combination to the central value tune for a $+1$ (-1) σ shift [32].

The last systematic uncertainty controls the mean-free path (MFP) of pions, seen in Figure 2.13. A $\pm 1\sigma$ shift alters the mean-free path by $\pm 33.3\%$ and scales linearly up to $\pm 3\sigma$. The error band can be shown for the CX pion scattering process in the blue band.

2.8 Other Interactions

There are coherent (COH) and neutral current (NC) scattering processes that are possible in NOvA. However, this analysis focuses only on CC neutrino interactions and therefore these interactions are not considered. These interactions also comprise a small fraction of the total number of neutrino events in NOvA, as the dominant interactions are CC.

2.8.1 Other Uncertainties

There are several uncertainties that are related to ν_e interactions, the signal in the Far Detector. One uncertainty is `2ndclasscurr`. This systematic is associated with secondary interactions of electron neutrinos in the detector; changes in this parameter, therefore, will impact the normalization of the appeared ν_e in the NOvA Far Detector.

Two other uncertainties are `radcorrnu` and `radcorrnuubar`. 2^{nd} class currents are additional terms in the QE hadronic current which would affect electron and muon neutrino CC cross sections slightly differently. All three of these uncertainties will be relevant for the joint ND+FD fitting only, where the ν_e sample is included.

There are uncertainties associated with coherent scattering and NC scattering, however, they are not considered as this analysis does not attempt to constrain NC nor coherent interactions.

2.9 Neutrino Simulation

The NOvA result at the Neutrino 2020 conference in Chicago refers to the Monte Carlo simulation known as Production 5 [14]. This analysis is performed with next iteration of NOvA’s production campaign of data and Monte Carlo simulation: “Production 5.1” or “Prod5.1”. Production 5.1 is a second iteration of the Production 5 campaign, with minor changes to modeling and reduction of systematic uncertainties. The NOvA simulation was produced with the neutrino interaction generator GENIE v3.02.06 [33]. This analysis uses the simulation produced from GENIE “Comprehensive Model Configuration” (CMC) or “tune” N18_10j_00_000, which describes the combination of physics models listed in this Chapter (Table 2.1).

Chapter 3

The NOvA Experiment

The NuMI Off-axis ν_e Appearance (NOvA) experiment is a long baseline neutrino oscillation experiment that consists of two detectors – one located in Fermi National Accelerator Laboratory (Fermilab), and the second detector located 810 km away near Ash River, Minnesota – to observe neutrino interactions (Figure 3.1). The primary physics goal of NOvA is to observe ν_e (ν_μ) appearance (disappearance) and measure $\sin^2 \theta_{23}$ of the PMNS matrix, improve the limits of the CP-violating phase, δ_{CP} , and constrain the mass ordering of the neutrino mass eigenstates Δm_{32}^2 .

To accomplish this goal NOvA utilizes Fermilab’s Neutrinos at the Main Injector (NuMI) beamline to produce a concentrated ν_μ beam. This beam of neutrinos passes through NOvA’s Near Detector (ND) underground at Fermilab, 1 km from the beam source. The beam continues to travel through the Earth where 810 km downstream the neutrinos pass through the Far Detector (FD) in Minnesota, on the border of Canada. The neutrinos travel through the Earth and when they arrive at the FD, there is a probability the beam of ν_μ will oscillate and interact with the FD material as a ν_e , where NOvA, as its name suggests, focuses on observing ν_e appearance in the FD.

The following sections describe the details of the NOvA experiment starting with the NuMI Beam, Section 3.1. Section 3.2 describes the construction of the Near and Far Detectors. Section 3.3 describes additional uncertainties in NOvA related to the detector response.

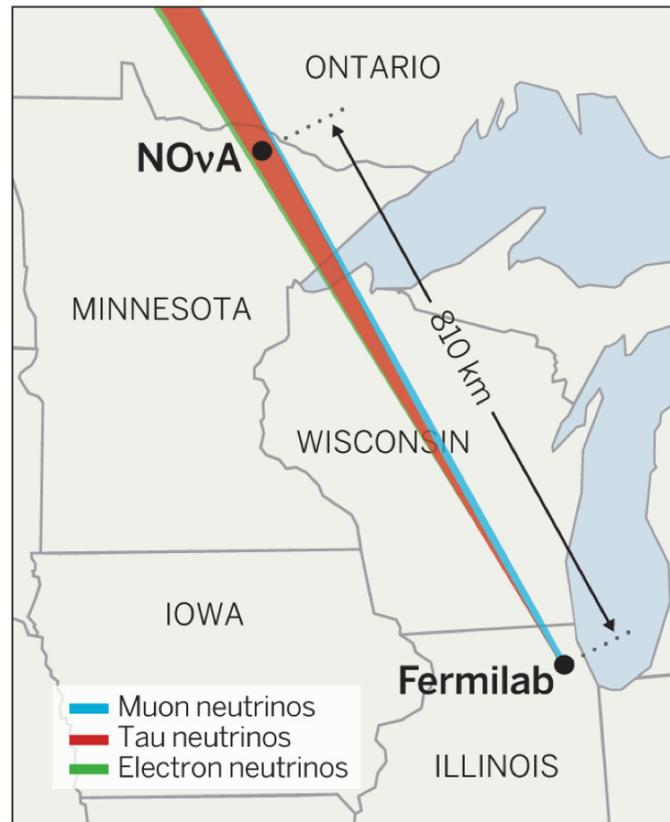


Figure 3.1: A layout of the NOvA experiment. Starting at Fermi National Accelerator Laboratory, a ν_μ beam is created and interacts with the NOvA Near Detector, also located at Fermilab. The beam propagates 810 km through the Earth where there is a probability ν_μ particles will interact and oscillate into ν_e and ν_τ with the NOvA Far Detector material in Minnesota. The energies from these interactions are reconstructed to approximate the energy of the interacting neutrinos in order to measure the neutrino oscillation parameters [34].

3.1 The NuMI Neutrino Beam

This section outlines the process for creating a high-energy neutrino beam for the NOvA experiment, starting from Fermilab’s Accelerator Complex and ending with a description of the NuMI beamline.

3.1.1 Accelerator Complex

NOvA’s source of neutrinos originates at Fermilab’s Accelerator Complex, Figure 3.2. Beginning at the Ion Source, or Pre-Accelerator, as H^- ions, they are accelerated to 750 keV through a series of radio frequency (RF) cavities, whereby an electric field is generated in the direction of the H^- beam. Following the Pre-Accelerator, the H^- ions travel to the linear accelerator, or Linac, where they are accelerated further to 400 MeV by additional RF cavities [35].

Next, the H^- ions reach the Booster Neutrino Beam (BNB). The BNB synchrotron accelerates the beam further to 8 GeV by way of additional RF cavities. At this stage, the H^- beam is stripped of its electrons by passing through a “stripping” carbon foil, becoming a proton beam. The BNB also “bunches” the protons, where alternating RF cavities are placed out of phase with each other, accelerating and decelerating parts of the beam to cluster the protons together – a “bunch”. The Booster collects 84 bunches, producing a “batch”; each batch is then delivered to the Recycler Ring in 1.6 μs intervals with each batch containing 4.3×10^{12} protons [35]. The Recycler Ring sits directly atop the Main Injector and merges 6 pairs of batches from the BNB – known as “slip stacking”. Once the 6 pairs of batches are merged, the beam is extracted into the Main Injector, where it is accelerated to 120 GeV. The result is 13.6×10^{13} protons on target (POT) at 400 kW in 10 μs spills at an average of 1.33 s are extracted to the Neutrinos at the Main Injector (NuMI) beamline.

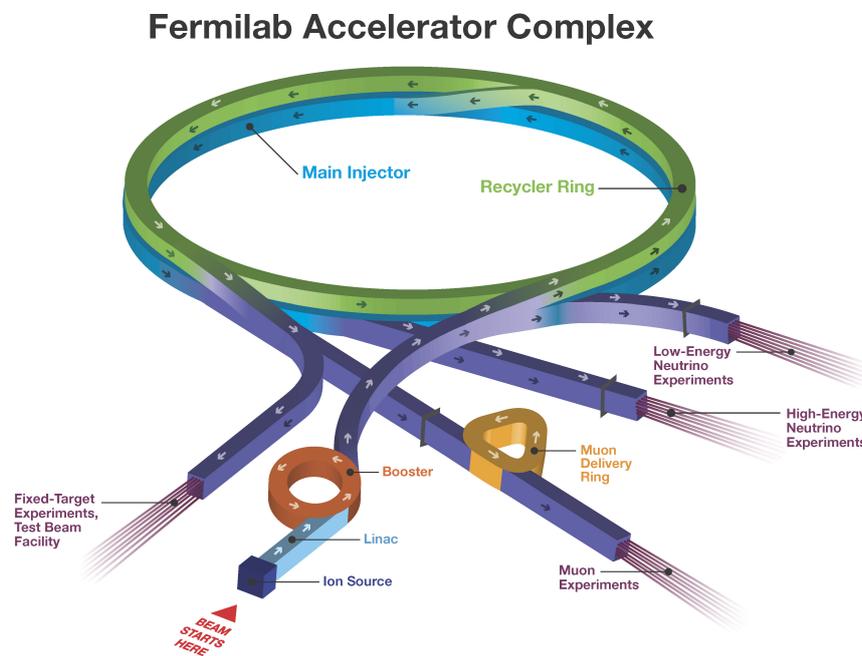


Figure 3.2: Fermilab's Accelerator Complex. The proton beam starts at the Ion Source, is accelerated to increasing energies at each subsequent stage, and ultimately 120 GeV protons are delivered to the NuMI beamline.

3.1.2 NuMI Beamline

After the proton beam is extracted from the Main Injector, the beam reaches the NuMI Beamline. It is here where the proton beam becomes a neutrino beam. The NuMI Beamline is a series of devices and equipment devoted to maximizing the production of neutrinos from the proton beam; a full diagram can be seen in Figure 3.3. From the left, protons strike a graphite target, interact with the internal nuclei, and produce secondary particles such as charged pions and kaons. The charged secondary particles are redirected with two toroidal shaped magnetic focusing horns. The magnetic field from these horns focus the secondary particles into the direction of the beam axis. Any undecayed secondary particles are stopped by the Hadron Absorber, while the muon and muon neutrinos pass through. Following the hadron absorber there are several Muon Monitors placed in the Earth to monitor the amount of muons to categorize the neutrino flux more accurately. The neutrinos, given their neutral charge and weak coupling to matter, travel through the absorber and the Earth towards the Near Detector. Arriving at the ND, some neutrinos will interact with the ND material and be observed, while the remaining neutrinos will travel through the Earth and propagate as mass states. For the few neutrinos that interact within Far Detector volume in Minnesota, ν_μ and ν_e will appear, dictated by the oscillation probability from Eqs. 1.21 and 1.22. NOvA attempts to reconstruct the particles from as many neutrinos interactions as possible to maximize sensitivity to the oscillation parameters.

Magnetic horns

To produce a neutrino beam or an antineutrino beam, the role of the magnetic focusing horns is crucial. Depending on the direction of the electric current running through them, the horns will focus either positive or negative charged particles yielding a neutrino or antineutrino beam, respectively. If the current running is in the forward direction, or downstream, towards the ND, the resulting magnetic will focus positively charged particles into the beam and push negatively charged

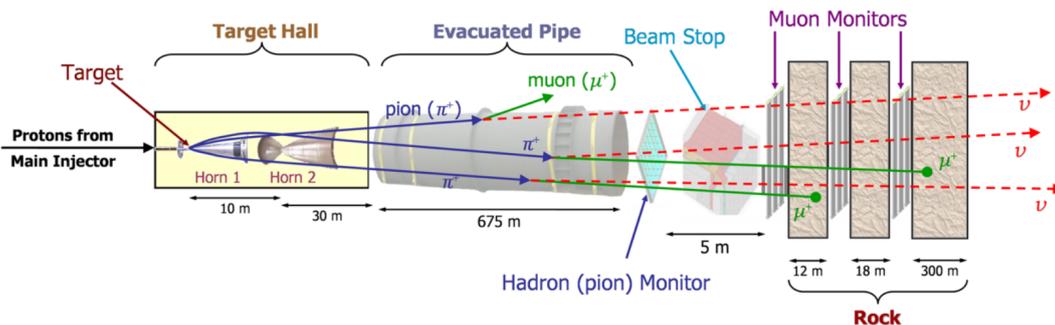


Figure 3.3: Side view of the NuMI Beamline. Protons from the Main Injector arrive from the left, impinge on a graphite target, and produce K^\pm and π^\pm . These particles are focused with the magnetic horns. The focused particles enter the decay pipe and decay to ν_μ and μ , with a small contamination of ν_e and e . At the end of the decay pipe a Hadron Absorber absorbs undecayed secondary particles. The muons pass through the Hadron Absorber and are absorbed by the Earth after passing several Muon Monitors to approximate the neutrino flux. Meanwhile, the neutrinos continue to the Near Detector and onto the Far Detector.

particles away from the beam. The result is a concentrated beam of predominately K^+ and π^+ particles, which will decay into

$$K^+ \rightarrow \mu^+ + \nu_\mu, \quad (3.1)$$

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \quad (3.2)$$

creating a beam of neutrinos. Thus, current running forward is synonymous with a neutrino beam; we frequently refer to a neutrino beam as the Forward Horn Current (FHC) configuration. Conversely, for electric current running backwards, or upstream – toward the Main Injector – the magnetic field will expel positive particles and focus negative particles, namely K^- and π^- , which predominately decay via the channel

$$K^- \rightarrow \mu^- + \bar{\nu}_\mu, \quad (3.3)$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu. \quad (3.4)$$

This is referred to as Reverse Horn Current (RHC), and produces an antineutrino

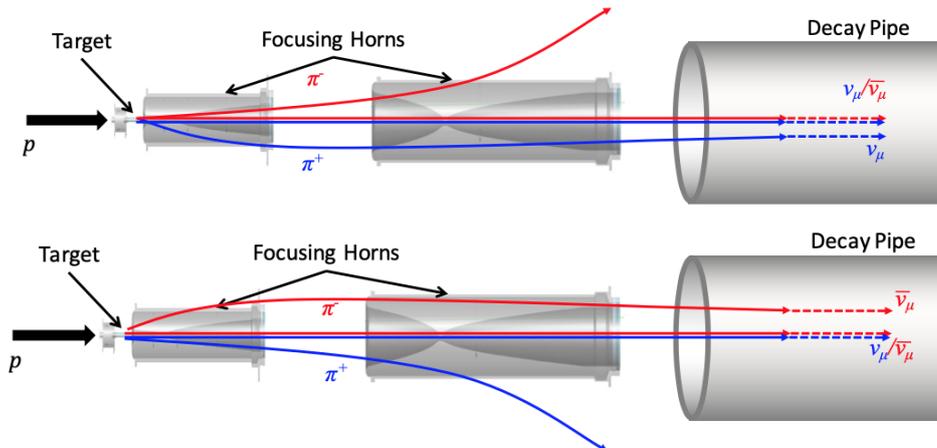


Figure 3.4: Diagram of the behavior of Forward Horn Current (FHC, top) and Reverse Horn Current (RHC, bottom) magnetic horn configurations. In the FHC (RHC) configuration, the positive (negative) particles are focused while the negative (positive) particles are pushed out of the magnetic horn; the result is a beam of neutrinos (antineutrinos).

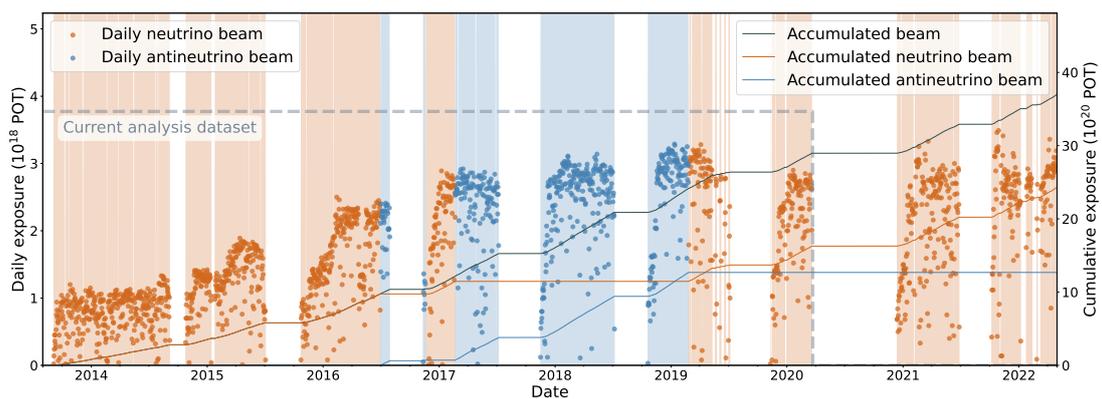


Figure 3.5: Total accumulated POT for neutrinos and antineutrinos. A neutrino or antineutrino beam is made from the crucial role of the magnetic focusing horns. NOvA's latest measurements analyze data up to 2020. From [36].

beam. Both the terminology FHC and RHC to describe the beam of neutrinos will be used commonly in the subsequent chapters.

A diagram of the neutrino and antineutrino beam running can be seen in Figure 3.4. NOvA runs in both of these configurations for extended periods of time (typically months). An overview of the POT and neutrino/antineutrino beam produced can be seen in Fig. 3.5 [36]. Note the current NOvA measurements has yet to incorporate data from beyond 2020.

3.1.3 Off-axis Design

When designing a long baseline neutrino oscillation experiment, there are several characteristics the experiment should possess. One is a long baseline, to ensure the probability for oscillation. The second is to maximize the size of the two detectors. Because neutrinos interact weakly and rarely, it is crucial to increase the target mass as much as possible to maximize the number of neutrino interactions within the detector volume. With these two key design choices, the ability to move the detectors becomes intractable; the distance between the two NOvA detectors is fixed at 810 km. Therefore, the neutrino energy is the only degree of freedom available to NOvA, and it is important to select an energy that maximizes neutrino oscillations (akin to E_{max} from the two-flavor example, Equation 1.16). However the secondary particles from the collisions of the 120 GeV protons to the graphite target are produced at well-defined energies. Therefore, NOvA prefers neutrinos with energies associated with higher oscillation probability. To do this, NOvA's detector utilizes an off-axis design.

The Near and Far Detectors are aligned off-axis to the neutrino beam. Equivalently, the beam axis and the Detectors subtend an angle, rather than sitting directly along the axis of the NuMI Beamline. This angle is 14 mrad. The consequence is NOvA receives neutrinos from the NuMI beamline that coincide with the first oscillation maximum for $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$. Figure 3.6 illustrates the utility of an off-axis design. The top plot is the ν_μ event rate without oscillations at 810 km – the lo-

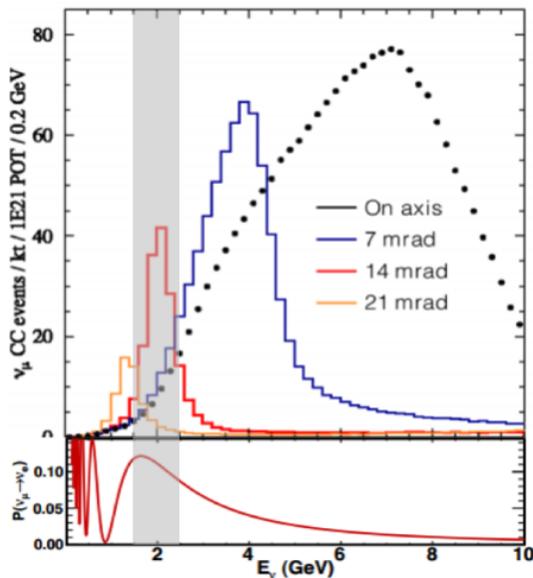


Figure 3.6: **Top:** ν_μ event rate at 810 km with no oscillations, as a function of E_ν . **Bottom:** $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$ as a function of E_ν at the same fixed distance 810 km. The grey band represents the energy range in which NOvA receives the NuMI beam. NOvA utilizes an off-axis design of 14 mrad indicating an ideal choice that coincides with the first oscillation amplitude of $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$ thereby maximizing the ν_e appearance in the FD.

cation of the NOvA FD – as a function of the neutrino energy. Each distribution represents increasing angles of the NOvA Near and Far Detectors with respect to the beam axis. The bottom plot is the $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$ distribution for oscillation at the same 810 km distance. When the NOvA detectors are on axis with the beam axis the event rate at the FD shifts towards higher energies with a peak at $E_\nu \approx 7$ GeV (black dots), while increasing off-axis angles shifts the ν_μ event rate to smaller energies at the cost of fewer ν_μ events. Analyzing this scan of angles, a 14 mrad option is ideal for NOvA’s physics goals; the peak of the ν_μ event rate coincides with the first oscillation amplitude of $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$ (the grey band), making it a useful design choice to maximize ν_e appearance for NOvA’s FD. Furthermore, this narrow band of neutrino energies provides an added advantage of an increased background rejection capability. Any particles appearing in the detectors beyond this energy range of the grey band can be disregarded as background and not a neutrino from the NuMI beam.

3.2 NO ν A Detectors

The NO ν A Near and Far Detectors are designed to observe CC ν_e and ν_μ interactions and to possess identical signal and background efficiencies and are functionally identical. The two detectors are composed of polyvinyl chloride (PVC) pipes filled with mineral oil – a liquid scintillator. The volume of the Near Detector is 290 tonnes, while the Far Detector contains 14 kton. Figure 3.7 illustrates the relative size.

Within each detector, groups of 32 PVC pipes are arranged together to form a plane and are aligned in alternating fashion – vertically and horizontally, in the xz plane, perpendicular to the beam direction – to allow for 3D reconstruction. Particles from interactions travels through the pipes and a flash of light is produced. The light is collected along a wavelength-shifting fiber in each pipe, where the ends of the fiber contain a pixel of an avalanche photo diode (APD) that converts the light into a digital signal (Fig. 3.8); the detectors consist of 63% active scintillating material [37]. Meanwhile, the energy of muon is reconstructed by range while the energy of all other particles are reconstructed from calorimetry.

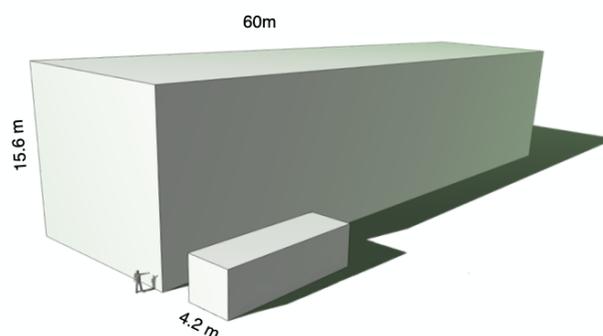


Figure 3.7: Graphic illustrating the relative size of the NO ν A Near and Far Detectors. The Near Detector consists of 290 tons and the Far Detector is 14 kton of total mass.

3.2.1 The Near Detector

The Near Detector (Fig. 3.9) is comprised of 130 tons of scintillating material, for a total mass of 290 tons. The cells within the detector are $3.9\text{ cm} \times 6.6\text{ cm} \times$

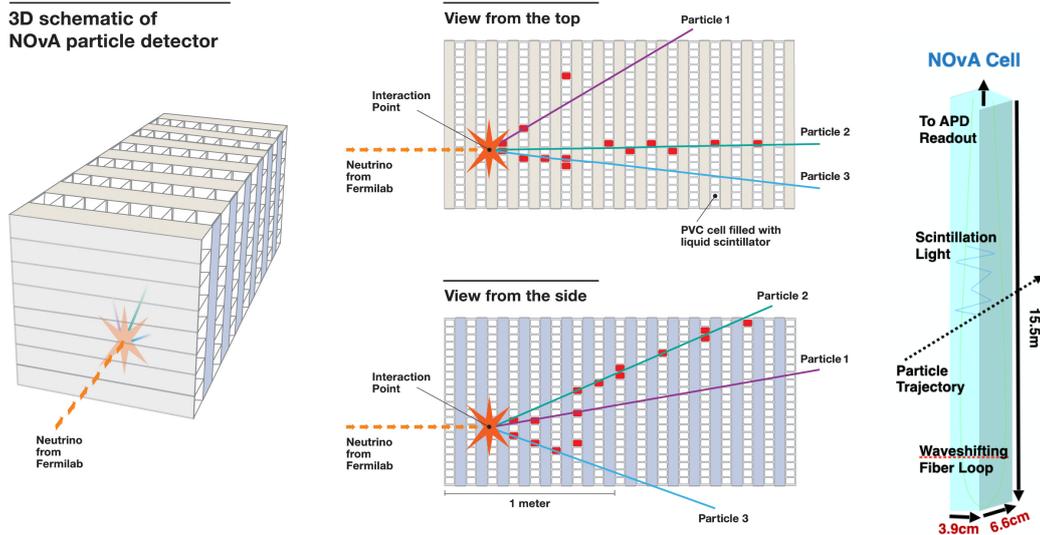


Figure 3.8: An illustration of the function of the NOvA Near and Far Detectors and an individual cell. Light is collected on the Fiber Loop and recorded onto the APD readout at the end of the cell, where it is then digitized into a single pixel. These cells are aligned in a cross-hatched fashion to allow for 3D reconstruction of a particle’s trajectory; both detectors are arranged in this manner. From [38].

3.9 m creating a total of 214 total planes in the ND. Before the start of the detector, 11 pairs of identically alternating vertical and horizontal PVC planes are separated by ten steel slabs of 10 cm, all perpendicular to the beam direction. This ”muon catcher” is designed to prevent muons from ranging out of the detector. These steel slabs allow muons to terminate and deposit all of their remaining energy within the detector to be observed.

Recall the ND is located 1 km downstream from the NuMI Beamline at Fermilab and receives the neutrino beam off-axis at 14 mrad. In addition, the ND is located 100 m underground, where cosmic ray background is not a concern in data collection.

3.2.2 The Far Detector

The NOvA Far Detector (3.10) is far larger than the ND; it is composed of 896 planes of PVC cells consisting of 344,064 channels and a total mass of 14 kton, Fig. 3.10. Just like the ND, the FD is oriented 14 mrad off-axis from the neutrino beam, and is 810 km away from the NuMI beamline in Minnesota. It sits on the surface of the Earth and stored in a housing unit that provides overhead protection from



Figure 3.9: The NOvA Near Detector located at Fermilab. The ND is housed underground and 1 km away from the NuMI beamline. From [38].

cosmic ray background via 1.2 m of concrete and 15 cm of barite. Despite this, cosmic rays provide a significant source of background when searching for neutrino oscillations in the NOvA data.

3.2.3 Simulating the Detectors

To reconstruct the energy of the particles, the calorimetric response must be well understood, specifically the detector’s ability to collect and record Cherenkov, scintillation light, and neutron response. To this end, the detector performance is simulated from the particle and material tracking software GEANT [39]. Because the hits of activity in the detector translate into energy that is used to estimate the neutrino, the accuracy of recording this information is crucial. For example, in the larger FD, the PVC cell pipes are very long, and any signal must travel the length of the cell to the APD. However, because of the length of the cells, attenuation is common where the light observed at the end of the APD is not the true response and why the light recorded is calibrated as a function of the distance it travels in the pipe.

In addition, NOvA’s neutron response is estimated from GEANT. As neutrons

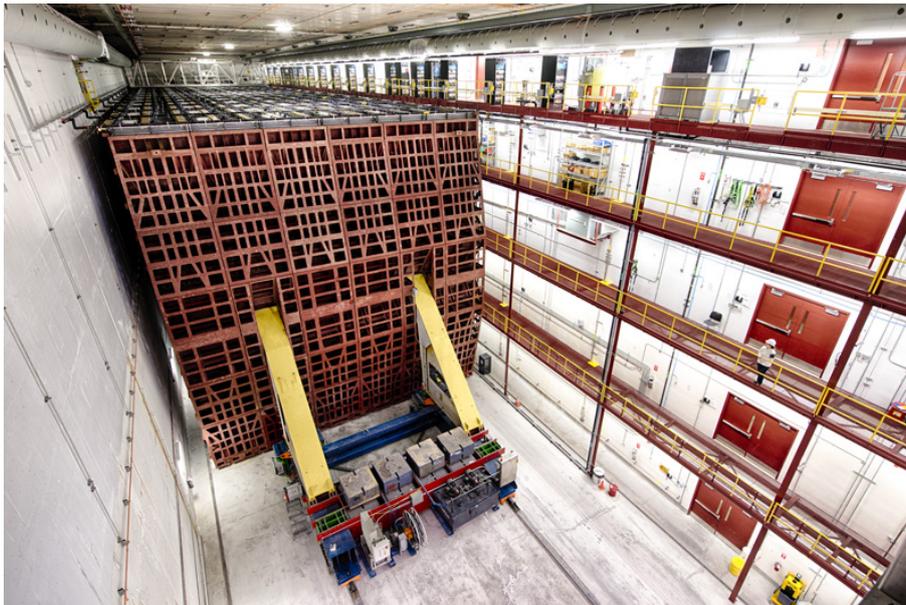


Figure 3.10: The NOvA Far Detector in Ash River, Minnesota. The mass of the detector is 14 ktons and sits on the surface of the Earth where it is exposed to cosmic rays. From [38].

are not observed in the NOvA detectors, GEANT is used to simulate how much energy is expected from the secondary particles, or daughters, of an energetic neutron.

3.3 Additional Systematic Uncertainties

In addition to the cross section uncertainties discussed in the previous chapter, there are uncertainties related to the accuracy of the neutrino flux that the NOvA detectors receive and the efficacy of the reconstruction from the detectors.

Thus far the the mathematics of neutrino oscillation, neutrino interactions, and design of the NOvA experiment have been outlined. Recall the objective of this thesis is to use Markov Chain Monte Carlo to constrain the neutrino interaction model uncertainties and the oscillation parameters. We have discussed neutrino oscillations, neutrino interactions, and the uncertainties associated with relevant neutrino interactions.

Following the discussion of the previous sections, this section introduces new uncertainties that are most closely associated with the functionality of the NuMI beamline and NOvA detectors: the neutron response uncertainty (Subsection 3.3.1),

detector efficiency (Subsection 3.3.2), and neutrino flux (Subsection 3.3.3).

3.3.1 Neutron Uncertainty

The `NeutronVisEPrimaries2018` uncertainty addresses the neutron response in the detectors. This is one of NOvA’s largest sources of uncertainties, as neutrons are invisible in the NOvA detectors, carry a significant amount of energy from the incident neutrino, and can significantly impact the shape of the muon disappearance spectrum and therefore the accuracy of the measurement of the oscillation parameters. This uncertainty focuses on the neutron propagation from GEANT [39]. Neutrons that are expected in the NOvA detector are on the order of hundreds of MeV, thus they are predicted via GEANT’s “Bertini Cascade” model [40]; this model predicts inelastic neutron activity from 20 MeV up to 10 GeV. In particular for NOvA, the simulation over-predicts neutrons in NOvA’s smallest observable calorimetric energy by approximately 30%, motivating the need for a conservative uncertainty [40].

3.3.2 Detector Uncertainties

There are five detector related uncertainties present in this analysis: `Calibration`, `CalibShape`, `Light_Level_ND`, `Light_Level_FD`, and `Cherenkov`.

NOvA reconstructs particle interactions with scintillation light, and therefore the ability to fully observe all of the light produced in the pipes is essential. This thesis uses uncertainties of the light collection, `Light_Level_ND` and `Light_Level_FD`, as uncorrelated between the two detectors, and moreover is the more conservative option. These uncertainties are $\pm 1\%$ shifts in the absolute calibration constants in GEANT.

The `Cherenkov` systematics adjusts the light-level response so that the muon response is kept constant. This systematic is identical in the ND and FD [41].

The `Calibration` energy scale systematic is 5% obtained from the data-MC discrepancy of candidate protons in the ND [42].

The `CalibShape` systematic addresses the different response rate at the edges of the detector. The discrepancy is observed in the data-MC difference between true and reconstructed energies of calorimetric particles and is taken as the uncertainty.

3.3.3 Flux Uncertainties

There are numerous uncertainties associated with a high-intensity neutrino beam. The uncertainties fall into two categories: beam transport and hadron production. Beam transport systematics are associated with the engineering of the beam – uncertainties such as the beam spot size, focusing horn positions in the target hall, and precise value of the magnetic field within the NuMI beamline. These uncertainties are small on the order of 5% and do not have a substantial impact on predicting the correct flux of neutrinos in the NOvA detectors [43].

The hadron transport systematics relate to the production of secondary particles – pions and kaons – from proton collisions at the target. NOvA uses the Package to Predict the FluX (PPFX), a package created to estimate the flux from hadron production models [44]. The uncertainties of hadron production are approximately 20%.

For past NOvA analyses and this thesis, a principal component analysis (PCA) is used to diagonalize the beam transport and flux uncertainties. This re-orientation of basis attempts to maximize the effect of all of the uncertainties while reducing the number of uncertainties, however, at the cost of redefining the physical interpretation of the new diagonalized uncertainties.

This analysis simultaneously constrains the NOvA Near and Far Detector MC, therefore, these PCA systematics must also account for shifts in the flux at the Near and Far Detector simultaneously (this is different from the conventional extrapolation NOvA flux PCA systematics [45]). This is done with the same software tools as the conventional PCA. The first 5 components, `ppfx_beam_nd_pc0[0-4]`, cover 97.5% of the variation in bins of true E_ν from 100 random universes [46]. Figure 3.11 shows the first of the five principal components, component 00. The figure

Chapter 4

Selection and Topological Samples

This chapter describes the event reconstruction algorithm in Section 4.1. The energy reconstruction algorithm is discussed in Section 4.2. Section 4.3 describes the selection process for the ND neutrino interactions. Section 4.4 describes the selection criteria for the FD ν_μ and ν_e samples.

4.1 Event Reconstruction

The event reconstruction starts with **raw hits**. This contains information of which pixels within the PVC pipes and plane (i.e. the location) the ADC charges deposited triggered, and is collected during 50 μs windows [47]. This information is recorded for both the data and MC.

Next the hits are grouped together via a **slicer** clustering algorithm, TDSlicer, that uses the spatial and temporal information of the hits. After “slicing”, NOvA obtains a neutrino event.

The next step utilizes the **Kalman-filter** algorithm to reconstruct particles with tracks that deposit their energy approximately uniformly such as muons (rather than via an electromagnetic shower) [48]. The reconstruction is performed in the 2D (yz , xz) views of the detector and then are later matched in 3D. A track can also be passed through the Reconstructed Muon Identification (ReMId) classifier and the track with the highest muon score is labeled the muon; this will be used for selecting

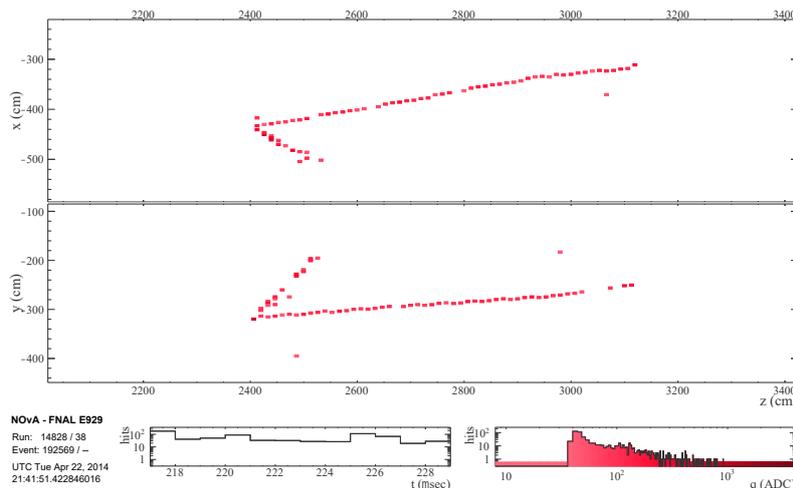


Figure 4.1: Reconstructed FD ν_μ CC event at $E_\nu = 2.75$ GeV. From the NOvA 2020 analysis [52].

FD candidate events.

The penultimate step is to identify the **vertex** from a slice – where the neutrino interaction took place. The vertexing occurs in two phases: first with a Multi-Hough transform algorithm [49] then an Elastic Arms algorithm. The Multi-Hough algorithm reads in pairs of pixels and determines if they are candidates for a particle track. The output of this algorithm is then fed into the Elastic Arms algorithm, which are then retraced to a single vertex to estimate the interaction site [50].

The last step is to identify **prongs**. A prong is a collection of hits associated with a single particle from a neutrino vertex, however the identity of the particle is not made. The construction of prongs is done in each 2D view via a possibilistic fuzzy means algorithm – hits can be assigned to noise and not a true track and there can be hits that are associated with multiple prongs [51]. At the end, these 2D prongs are joined together into 3D from the two 2D views. As we will see, we will utilize prong information to identify particles and divide the ND data and simulation into distinct topologies.

In the case of a ν_μ CC interaction, Figure 4.1 shows a candidate interaction reconstructed at 2.75 GeV in the NOvA FD data. Note the longer track would be identified as a μ and is easily recognized by its low energy deposition and long track. The proton can be identified by its shorter track, which deposits energy faster

per unit length. This is a common interaction in the NOvA detectors, in fact, this interaction will be one of the selection categories in the ND data and MC.

4.1.1 Particle Identification

At this stage the 3D prongs are identified into particles. We use a Convolutional Visual Network referred to as “prong-CVN” to identify all particles within a single neutrino interaction [53]. There are three primary particles this thesis must identify accurately: μ , p , and π . Table 4.1 summarizes the prong-CVN scores used in this analysis.

Particle	CVN cut
μ	0.5
p	0.5
π^\pm	0.7

Table 4.1: Table of the prong-CVN scores for the particles used in this analysis. For example, a prong is identified as a p if the p-CVN score is at least 0.5. Note a prong remains a prong (an unidentified particle) if no particle scores are met.

4.2 Energy Reconstruction

4.2.1 Muon Energy

Muon energy is calculated from the range of the muon in the NOvA detector. The constant and well-defined 2 MeV/cm of dE/dx energy deposition provides a method to estimate the total muon energy as it deposits all its energy in the detector. The muon energy resolution is 4% [54].

Muon Energy Scale Systematic

The muon energy is determined by range. NOvA’s uncertainty on the muon energy scale is the `kCorrMuEScaleSyst2020` systematic. This systematic is calculated in two parts: one element from GEANT & Bethe-Bloch equation parameters and

a second from the relative differences between the two detectors. These two parts sum in quadrature to approximately 1% [55].

4.2.2 Hadronic Energy

The hadronic energy is estimated from the calorimetric response of the detectors. The hadronic energy resolution is 0.23 GeV [54].

4.3 ND Topological Samples

To produce this fit, we must next identify the selection of NOvA ND ν_μ CC events. For the ND portion of this analysis, no ν_e events are analyzed. Consideration is made to determine samples of neutrino interactions via characteristics such as particle multiplicity and identity. Upon selecting samples, we then must determine variables to view the samples in. This section will discuss this process.

To determine an ideal neutrino interaction within the NOvA ND, one important criterion is to have a sufficient number of statistics; too few events will not capture a model's phase space sufficiently, nor produce meaningful results from MCMC sampling. Given the vast statistics of neutrino interactions at the ND, this is not a primary issue. Therefore it is likely any reasonable neutrino event selection will have sufficient stats.

Another consideration of determining candidate events is to have unique samples that are distinct and isolate as much of a physics model as possible; the ideal selection would consist of neutrino events that correspond to a single physics model. This is desired but in practice unachievable, as the interaction models overlap significantly (Fig. 2.1) and due to NOvA's detector reconstruction and resolution limitations.

We first apply the standard NOvA 2020 ND pre-selection cuts before we partition the ND data & MC. The first is a quality cut, `kNumuQuality`; the cut requires $E_\nu < 5$ GeV and a basic quality cut that there are sufficient hits in the slice. The second cut, `kNumu2020PID`, requires a muon ReMId score > 0.3 and NOvA's ν_μ

event pre-selection score > 0.8 . The `kNumuContainND2020` cut requires the particles are contained within the fiducial volume of the ND. All of these pre-selection cuts are contained within a single NOvA cut: `kNumu2020ND`.

We select 10 Near Detector topological categories we will use for the ND fit; there are 5 FHC topologies and 5 RHC topologies. These categories, along with the most common interaction type, can be seen in Table 4.2. Note the “Muon + (Proton)

Horn Current	Topology	Dominant Interaction Type
FHC + RHC	Muon	QE + MEC
FHC + RHC	Muon + Proton	MEC + RES
FHC + RHC	Muon + Pion + Prong	RES + DIS
(FHC) RHC	Muon + (Proton) + Prong	RES + DIS
FHC + RHC	Remaining	RES + DIS

Table 4.2: The ten ND topological categories and the most common interaction types associated with each one. The fourth category contains a proton for FHC ($\mu+p+X$) while the RHC category has no proton: $\mu+X$. There is modest separation of the interaction types by using these selected topological categories. The QE and MEC truth interactions are predominately observed in the Muon topology. The Muon + Proton topology consists mostly of MEC and RES events. The other three topologies are dominated by RES and DIS interactions, illustrating the difficulty to create a topology that is dominated by a single interaction type. The composition of interaction types for each topology is shown in the 1D projections of these topologies, Figs. 4.8 & 4.9.

+ Prong” topology contains the proton for the FHC topology and does not for the RHC topology. In addition, “Muon + Pion + Prong” and “Muon + (Proton) + Prong” categories require a prong – an unidentified particle track or shower.

To select these topologies from the NOvA ND data and simulation, additional prong-CVN cuts are applied and outlined below.

Muon

The Muon topology for both FHC and RHC, known as the shorthand μ :

- contains one 3D prong,
- the prong has CVN_μ score > 0.5 ,
- μ track length must be > 5 m.

In this topology, it is uncommon to observe events with energetic hadronic particles. This topology is predominately QE and MEC.

Muon + Proton

This topology also applies to both FHC and RHC beam running. Also known as the $\mu + P$ topology:

- contains two 3D prongs,
- one prong must have CVN_μ score > 0.5 ,
- the μ track length must be > 5 m,
- the other prong must have CVN_p score ≥ 0.5 .

This topology mostly contains MEC and RES interactions, as well as some QE.

Muon + Pion + Prong

The pion-producing sample, $\mu + \pi^\pm + X$, also applies to FHC and RHC:

- contains two or more 3D prongs,
- one prong must have CVN_μ score > 0.5 ,
- the μ track length must be > 5 m,
- one prong must be identified as a charged pion with CVN_π score > 0.7 ,
- there may be additional unidentified prongs, known as “ X ”; these must have CVN_π score < 0.5 to be considered X (a prong).

This topology is dominated by RES and DIS events.

Muon + (Proton +) Prong

The $\mu + (P+)X$ topology. For FHC, the topology requires a proton: $\mu + P + X$. The RHC topology does not: $\mu + X$. The $\mu + P + X$ topology has one added prong-CVN cut for protons, listed in parentheses:

- contains two or more 3D prongs,
- one prong must have CVN_μ score > 0.5 ,
- the μ track length must be > 5 m,
- (the other prong must have CVN_p score ≥ 0.5 .)
- there may be additional unidentified prongs, known as “ X ”; these must have CVN_π score < 0.5 to be considered X (a prong).

This topology is dominated by DIS events, as they are generally characterized by high particle multiplicity.

Remaining

This “Remaining” or `EvElse` category applies to both FHC and RHC:

- Any events that do not fall under the previous four topological samples.

It is a catch-all for neutrino events that do not fall into the other four categories, and it too is dominated by DIS events.

Going forward, the “Abbreviation” names in Table 4.3 will be used as a shorthand to refer to the topologies. Figure 4.2 also displays a summary of the selection requirements, including the prong-CVN cuts, for each of the topological samples.

Topology	Abbreviation
Muon	μ
Muon + Proton	$\mu + P$
Muon + Pion + Prong	$\mu + \pi^\pm + X$
Muon + (Proton) + Prong	$\mu + (P+)X$
Remaining	<code>EvElse</code>

Table 4.3: Table of the names of the particles and the commonly used shorthand name for each topology. Note the “Muon + (Proton +) X” category includes the proton for FHC mode and does RHC mode.

After creating the topological samples, we report the number of data and MC events in each topology in Table 4.4. We see there is an order of magnitude more events in the FHC and RHC μ topology than the other topologies, while the $\mu + P + X$ topology has the least events, given its more stringent cuts.

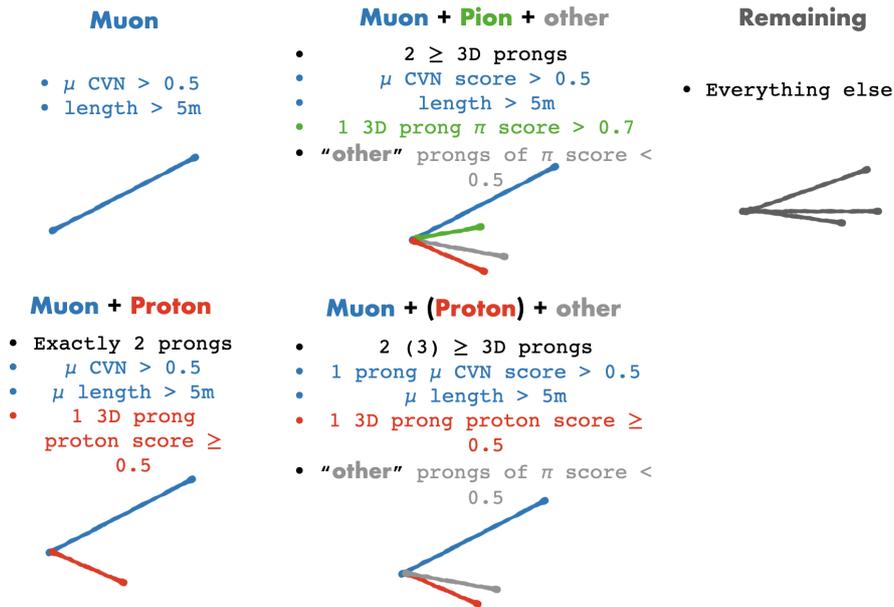


Figure 4.2: Cuts applied to interactions to separate the ND MC/data into the topological samples.

Topology	Data Events	MC Events
FHC μ	1,109,418	1,090,078
FHC $\mu + P$	486,409	515,242
FHC $\mu + \pi^\pm + X$	368,890	370,185
FHC $\mu + P + X$	102,958	139,022
FHC EvElse	676,479	752,449
RHC μ	714,840	725,667
RHC $\mu + P$	245,772	291,960
RHC $\mu + \pi^\pm + X$	114,078	123,575
RHC $\mu + X$	145,691	214,434
RHC EvElse	277,887	296,207

Table 4.4: Table of the data and MC events for each topology.

4.3.1 CVN_π Performance for FHC $\mu + \pi^\pm + X$ Topology

The efficiency of the prong-CVN in identifying particles is illustrated here in Figure 4.3. This Figure shows the CVN_π score distribution for all particles (top left), final state γ (top right), final state π^\pm (bottom left), and final state p (bottom right). The histograms in these figures are stacked histograms of the different true final state particle combination. Note how the dominant histogram colors are the orange: $\mu + nP + X$, and the purple: $\mu + 1\pi^\pm + X$. Also recall the CVN_π score cut to identify a π^\pm for this analysis is 0.7.

Looking at the "all particles" distribution (Fig. 4.3, top left), we see the data

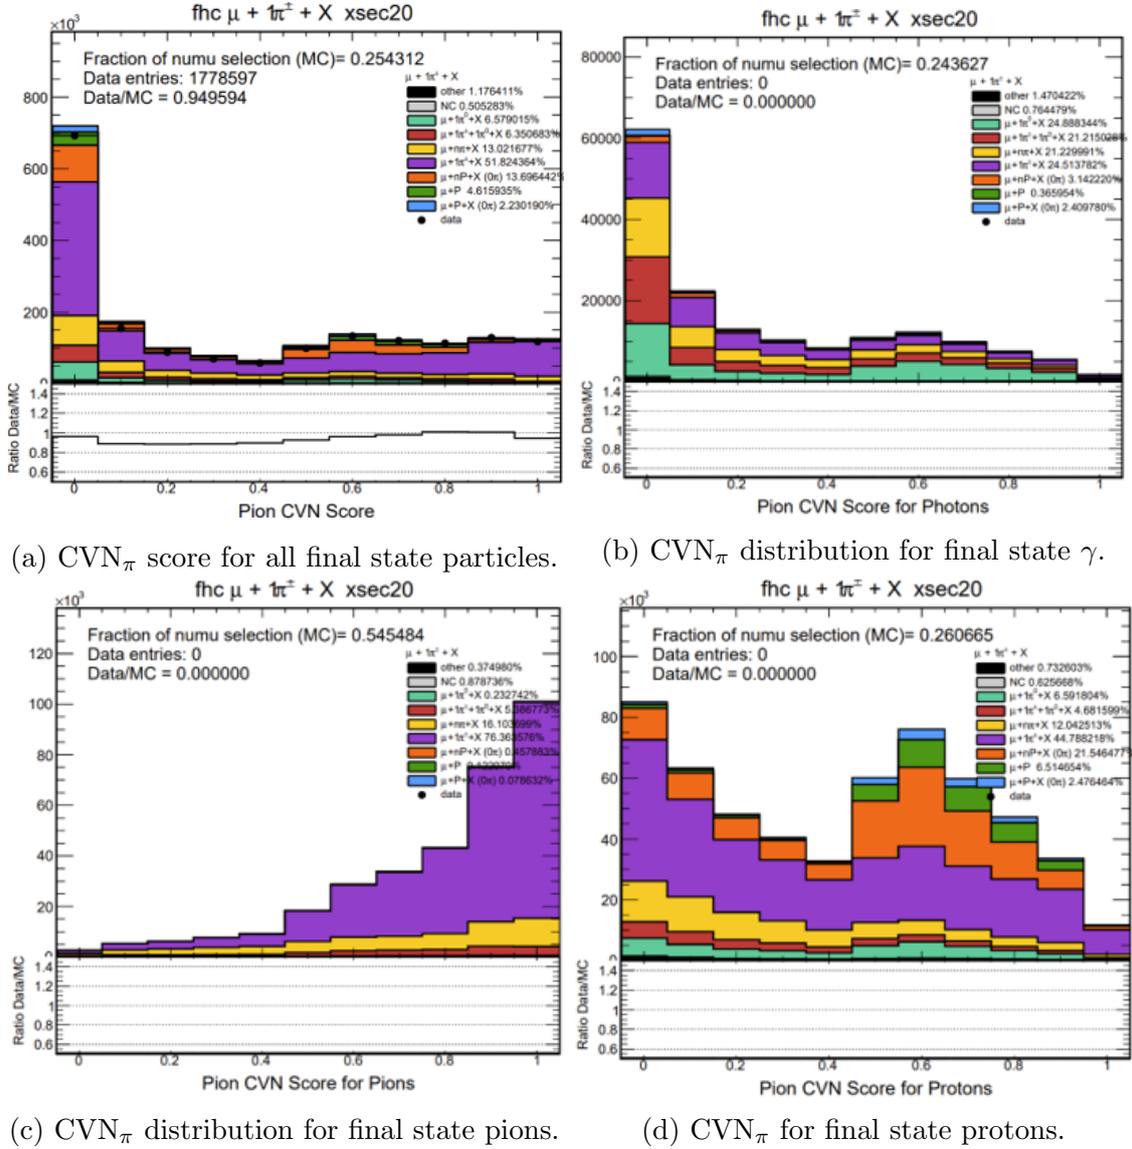


Figure 4.3: CVN $_{\pi}$ score distributions for all particles (top left), γ (top right), π^{\pm} (bottom left), and p (bottom right) broken down by the truth final state multiplicity. CVN $_{\pi}$ successfully distinguishes γ from π^{\pm} . However, not all protons can be distinguished from π^{\pm} . The orange histogram represents $\mu + nP + X$ which contain protons. Note how the bottom right plot contains a significant region of orange at CVN $_{\pi} > 0.7$ indicating protons are identified as pions. This highlights the difficulty in successfully identifying all particles in an interaction and will be revisited when using MCMC with the ND data. Note for the γ , π^{\pm} , and p distributions there is no data/MC ratio calculated.

and MC have a strong agreement at all values of CVN_π scores. This demonstrates CVN_π is successfully identifying π^\pm from the data. Moreover, the truth category for $CVN_\pi > 0.7$ is nearly entirely purple, $\mu + 1\pi^\pm + X$. This can be better seen in the π distribution (bottom left of Fig. 4.3). The distribution is nearly entirely purple, $\mu + 1\pi^\pm + X$, illustrating the CVN_π distribution is very effective at identifying π^\pm .

For the γ distribution (Fig. 4.3, top right) we see a small normalization (relative to the all particles normalization) and very little of any final state particle combination at $CVN_\pi > 0.7$, indicating γ are typically not misidentified as a π^\pm .

Lastly the p distribution (Fig. 4.3, bottom right), shows that there is a large portion of orange in the distribution, which contains protons: $\mu + nP + X$ and purple: $\mu + 1\pi^\pm + X$ at $CVN_\pi > 0.7$. This means that $CVN_\pi >$ is identifying protons as pions from the orange region. This demonstrates the difficulty in perfectly identifying all particles in an interaction. Therefore, these topological selections do not perfectly isolate neutrino interactions by particle identity and multiplicity, especially in the case of protons and pions.

4.3.2 Variables for ND Selection: E_{had}^{vis} & $Reco |\vec{q}|$

With these ten ND topological samples selected, we proceed to select the variables to observe them in. We opt for a two dimensional phase space: hadronic visible energy (E_{had}^{vis}) and reconstructed three-momentum transfer ($Reco |\vec{q}|$). The characteristics of these variables are outlined below.

$Reco |\vec{q}|$:

The reconstructed three-momentum transfer is calculated by

$$|\vec{q}| = \sqrt{Q^2 + (E_\nu - E_\mu)^2}, \quad (4.1)$$

where $Q^2 = 2E_\nu(E_\mu - p_\mu \cos \theta_\mu) - m_\mu^2$. Recalling Feynman diagram Fig 1.2, $Reco |\vec{q}|$ should contain necessary information of the muon kinematics from the interaction.

The resolution of this variable is 0.14 GeV, with a fractional resolution of 21%.

E_{had}^{vis} :

The E_{had}^{vis} variable is necessary to describe the hadronic system from an interaction, and is independent from the muon kinematics. We use E_{had}^{vis} as it is the sum of the non-leptonic energy deposited in the detector. This energy includes:

- kinetic energy of protons and charges pions.
- total energy of electrons, photons, neutral pions, and kaons.
- total energy minus nucleon mass of hyperons.
- total energy for anti-nucleons.

Notably, this does *not* include energy from neutrons. Again, the energy resolution of this variable is 0.23 GeV [54].

An important motivation for selecting these two variables is because of their strong connection to the neutrino interaction model; $(Reco |\vec{q}|, E_{had}^{vis})$ serve as a proxy for true (\vec{q}, q_0) . We have already seen, especially with the MEC model in Section 2.4, true (\vec{q}, q_0) are closely tied to the interaction model as it describes how much energy and momentum are divided into the leptonic and hadronic systems. The reconstructed variables will be provided into Markov Chain Monte Carlo sampling to learn as much about the true energy and momentum transfer as possible in an interaction. This will allow a more direct understanding of which interaction models are insufficient given what uncertainties are being stressed.

These 2-dimensional distributions – in each of the ten topological samples – can be seen in Figure 4.4 (FHC) and Figure 4.5 (RHC). These are the distributions that will be input to MCMC for sampling. The FHC Muon topology is peaked at low hadronic visible energy and reconstructed three-momentum transfer – approximately (0.25 GeV, 0.05 GeV). In the $\mu + P$ topology, we see the peak of the 2D distribution is centered slightly higher at (0.6 GeV, 0.15 GeV). The $\mu + \pi^\pm + X$ and $\mu + P + X$ topologies contain a broader distribution and populate mostly along the E_{had}^{vis}

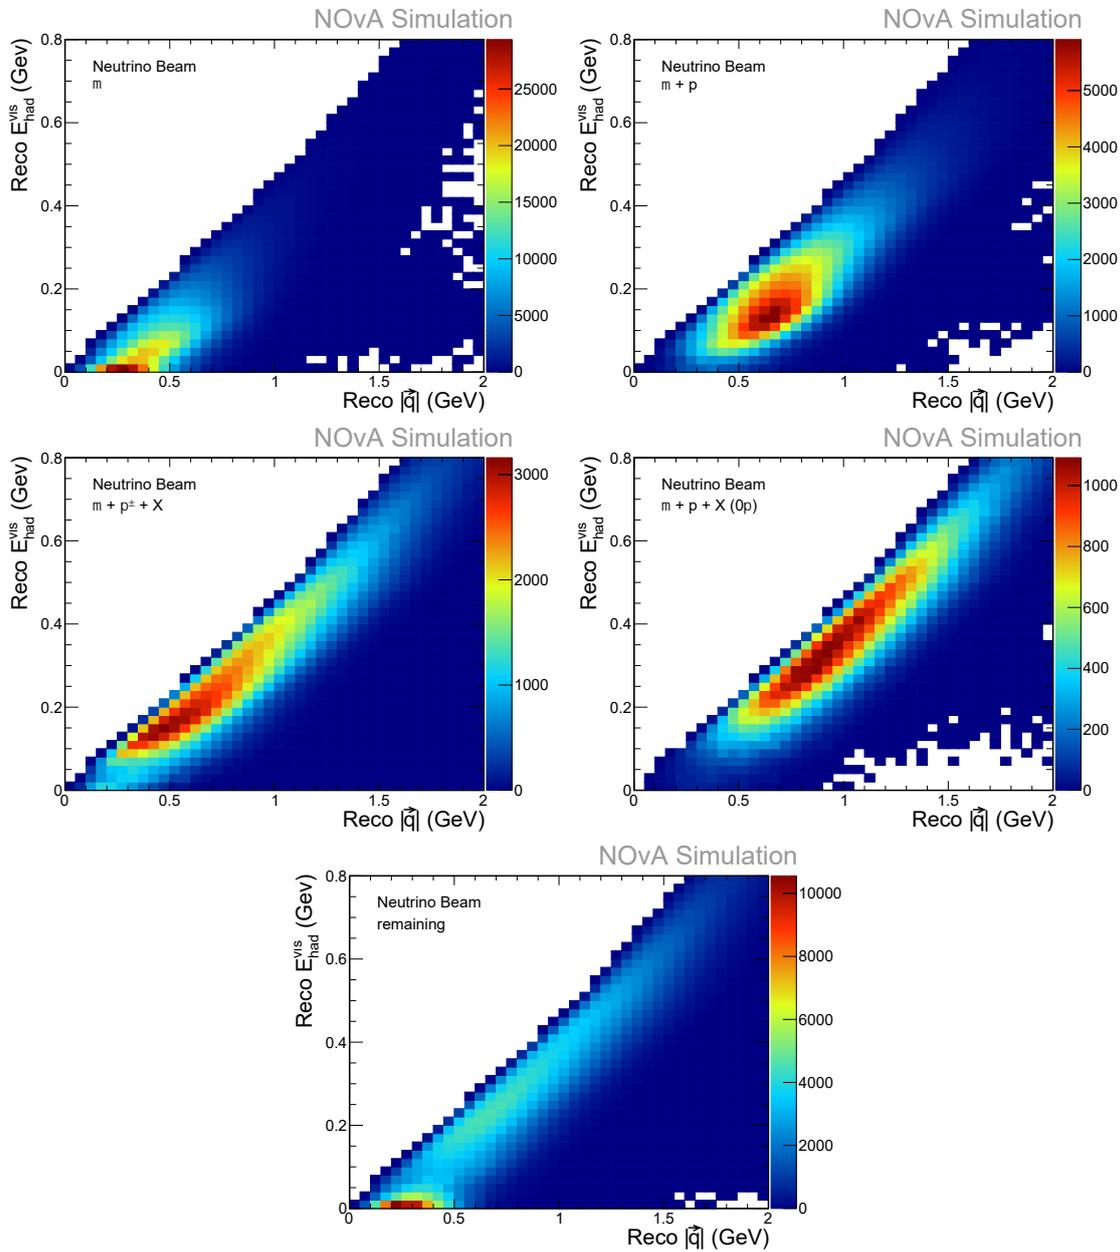


Figure 4.4: The 2-D distributions for the ND topologies in neutrino mode. The majority of events in all topologies are located at low E_{had}^{vis} and/or along the diagonal of E_{had}^{vis} and $Reco |\vec{q}|$. Note the different z-axis scales for the five plots.

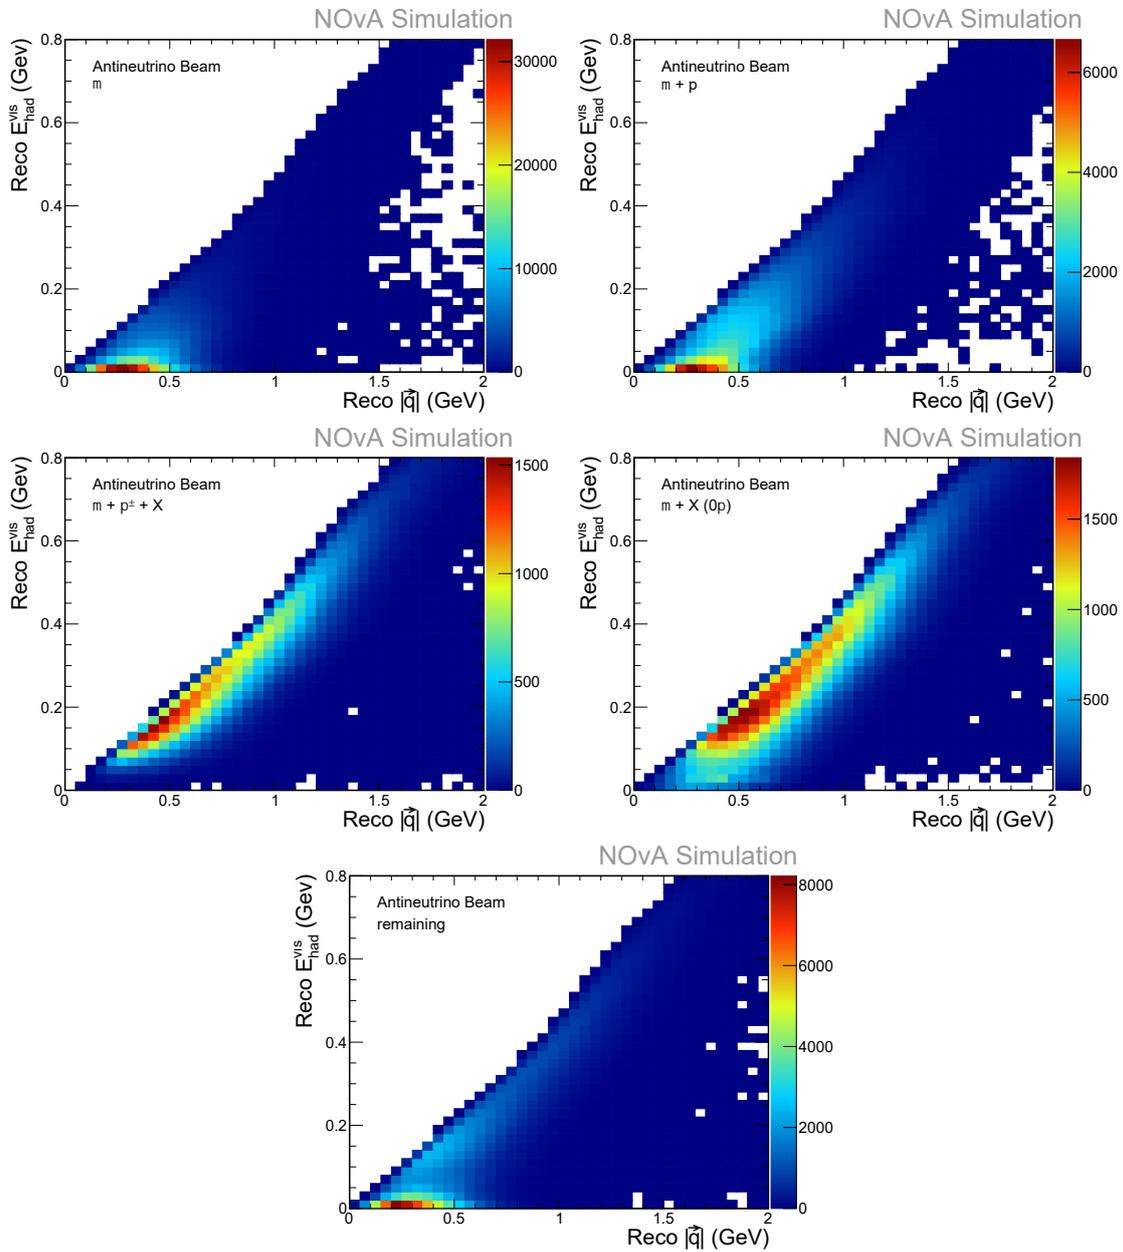


Figure 4.5: The 2-D distributions for the ND topologies in antineutrino mode. The majority of events in all topologies are located at low E_{had}^{vis} and along the diagonal of E_{had}^{vis} and $Reco |\vec{q}|$. Note the different z-axis scales for the five plots.

= $Reco |\vec{q}|$ diagonal, indicating there is likely more than one dominant neutrino interaction model in these samples. This observation holds for the **EvElse** topology as well, however these events are mostly centered at low E_{had}^{vis} and $Reco |\vec{q}|$.

The observations made for the FHC plots hold in the antineutrino (RHC) plots. However, the primary difference in the RHC plots (Fig. 4.5) is the Muon, $\mu + P$, and **EvElse** topologies are centered at lower E_{had}^{vis} than in FHC; this should be expected as the primary outgoing particle – particularly in a QE interaction – is a neutron, which has no observable scintillation in the NOvA detectors and neutron energy is not incorporated into the estimation of E_{had}^{vis} .

An appeal for using these two variables is the connection to the kinematic variable Q^2 , Eq. 2.4. With these 2D distributions, the $y = x$ diagonal in Figure 4.4 and Figure 4.5 represents the line of $Q^2 = 0$; events above this line are not permitted, while events of increasing Q^2 are located down and to the right of the distribution. This is seen in Figure 4.6 where CC RES events are shown in slices of Q^2 for the FHC $\mu + \pi^\pm + X$ topology (top $Q^2 < 0.5 \text{ GeV}^2$, bottom $0.5 \text{ GeV}^2 < Q^2 < 1.0 \text{ GeV}^2$). The two plots demonstrate this increasing Q^2 trend in the E_{had}^{vis} and $Reco |\vec{q}|$ phase space. This is a feature we will utilize later in the ND fit analysis.

Meanwhile, there is also a trend within this phase space for the hadronic invariant mass, W . Figure 4.7 shows CC RES events in the same FHC topology. As W increases, the population of events moves up and to the right in $(Reco |\vec{q}|, E_{had}^{vis})$ phase space. In addition, the RES $\Delta^+(1232)$ range ($W \approx 1.2 \text{ GeV}$) is very near the $Q^2 \approx 0$ diagonal, revealing that **NOvA simulates many low- Q^2 RES events**; this insight will be returned to again when MCMC constrains the ND MC. The other important conclusion is the higher regions of phase space can only be accessed with higher energy neutrinos, and therefore more inelastic scattering processes such as RES and DIS. These two figures help to illustrate the complex and overlapping physics processes that comprise the ND topologies.

Recall from Feynman diagram Fig. 1.2, in a CC neutrino interaction, the W boson transfers three momentum and energy to the hadronic system, therefore larger

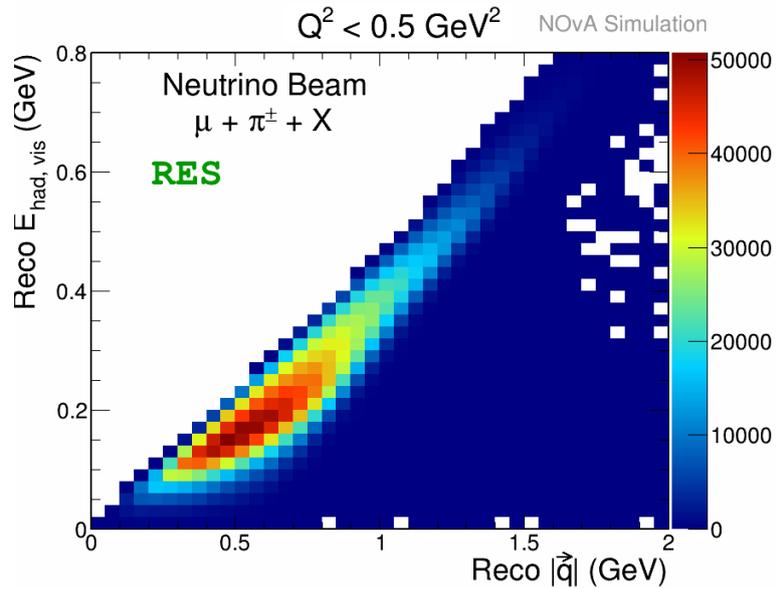
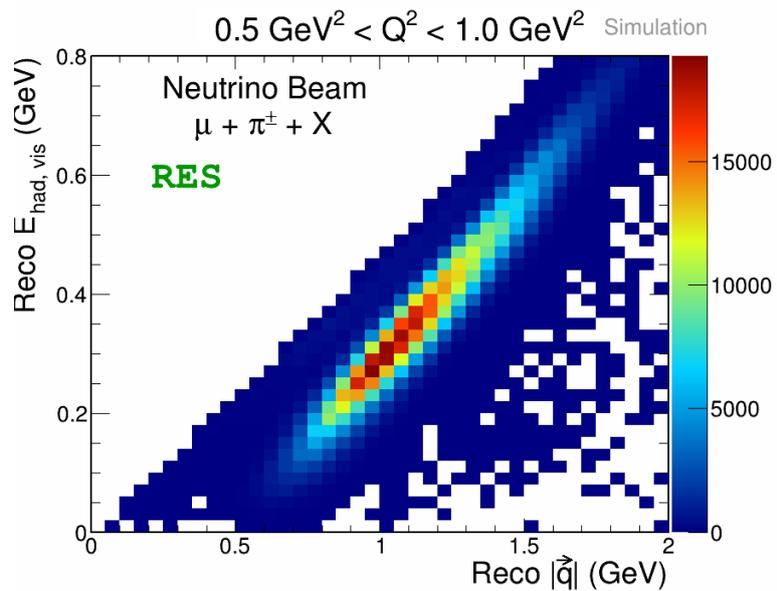
(a) RES events with $Q^2 < 0.5 \text{ GeV}^2$.(b) RES events with $0.5 < Q^2 < 1.0 \text{ GeV}^2$.

Figure 4.6: $\mu + \pi^\pm + X$ FHC topology in slices of Q^2 for true CC RES events. Note how low- Q^2 events ($Q^2 < 0.5 \text{ GeV}^2$) are populated along the $\text{Reco } |\vec{q}| - E_{\text{had}}^{\text{vis}}$ diagonal line of $Q^2 \approx 0$. Increasing values of Q^2 move the population of events further to the right, away from the diagonal.

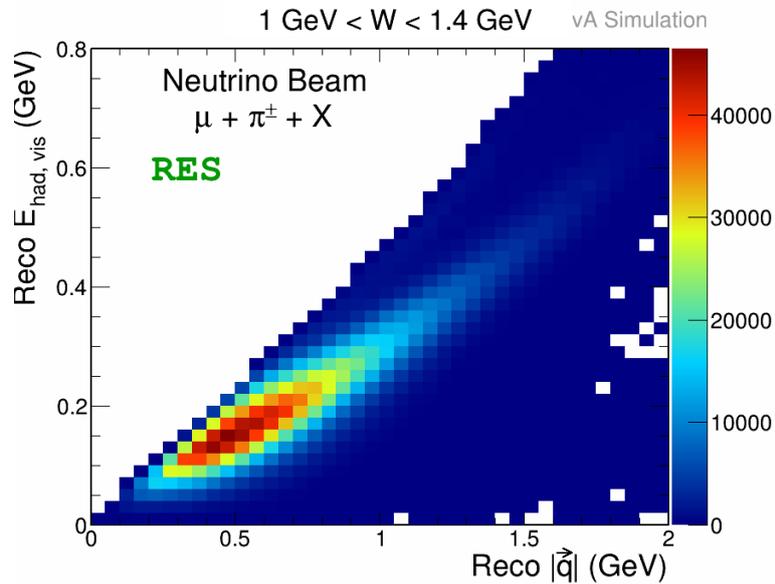
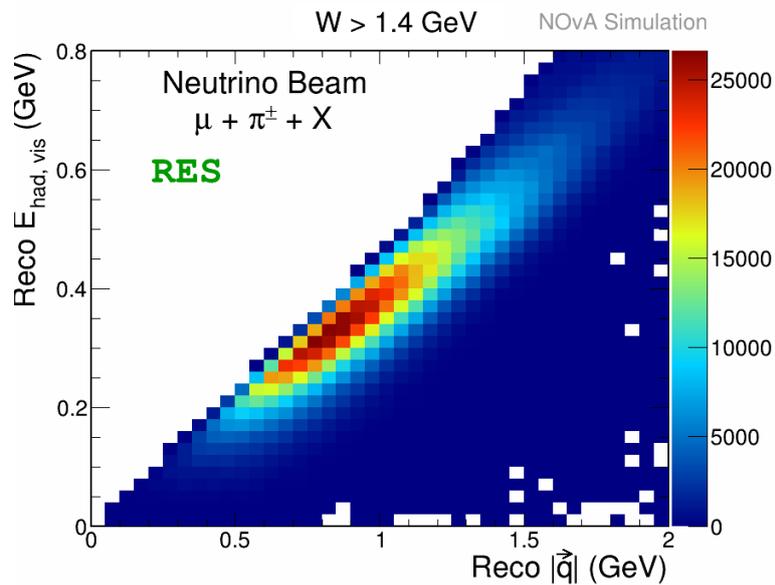
(a) RES events with $1. < W \leq 1.4$ GeV.(b) RES events with $W > 1.4$ GeV.

Figure 4.7: $\mu + \pi^\pm + X$ FHC topology in slices of W . The population of events move up and to the right in $(Reco |\vec{q}|, E_{had}^{vis})$ phase space. These events can only be accessed via larger E_ν values.

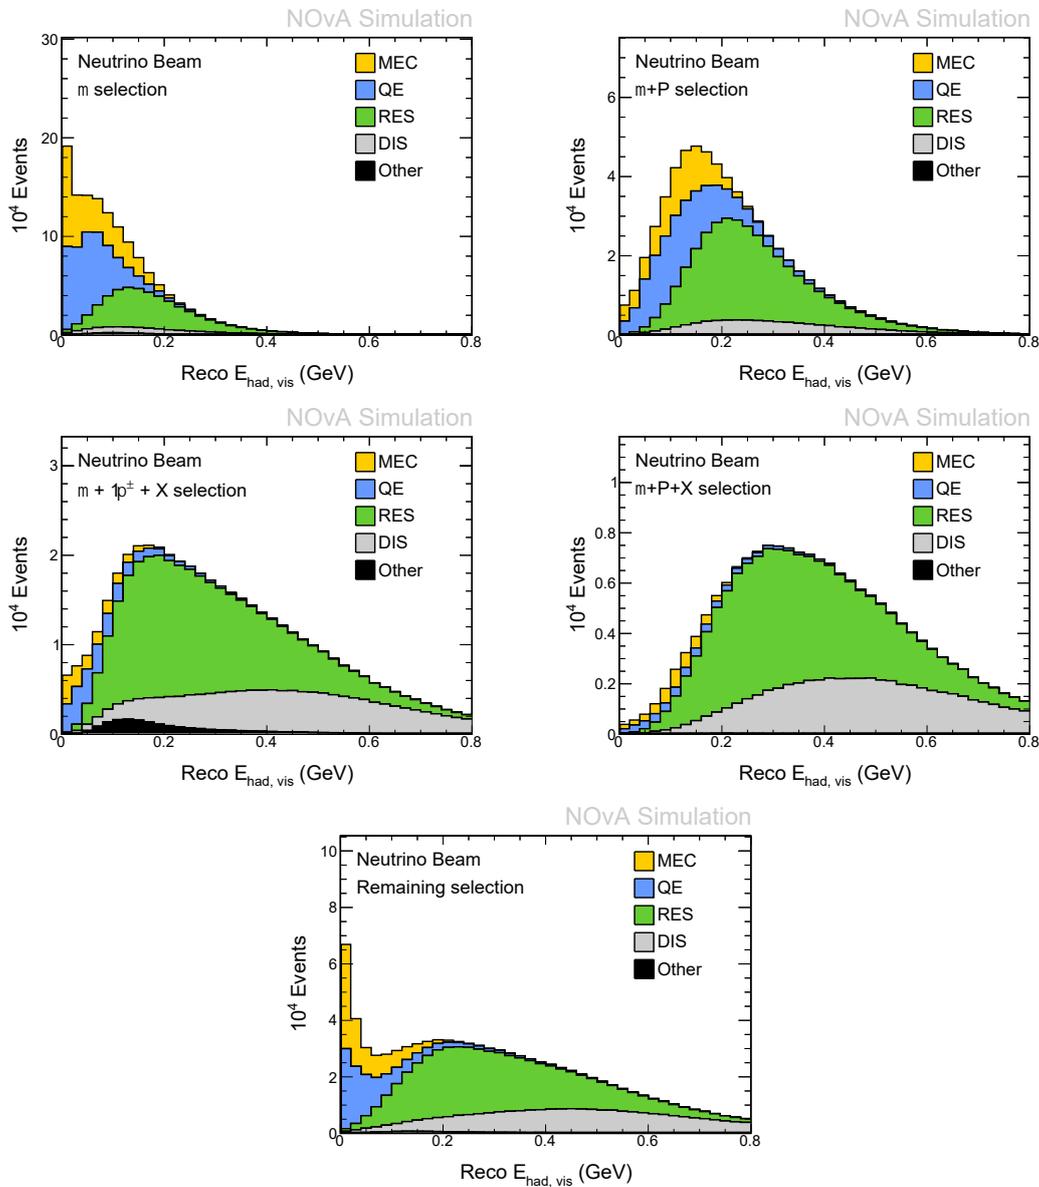


Figure 4.8: Projection of the 2D $(E_{had}^{vis}, Reco[\vec{q}])$ phase space onto the E_{had}^{vis} axis, projecting out $Reco[\vec{q}]$. These plots of the FHC topologies break down neutrino events by true interaction in the simulation, confirming the model breakdown from Table 4.2. The “Other” category is predominately COH scattering. See $\bar{\nu}$ plots in Fig. 4.9, and 1D FHC projections onto the $Reco[\vec{q}]$ variable are available in the Appendix A.2.

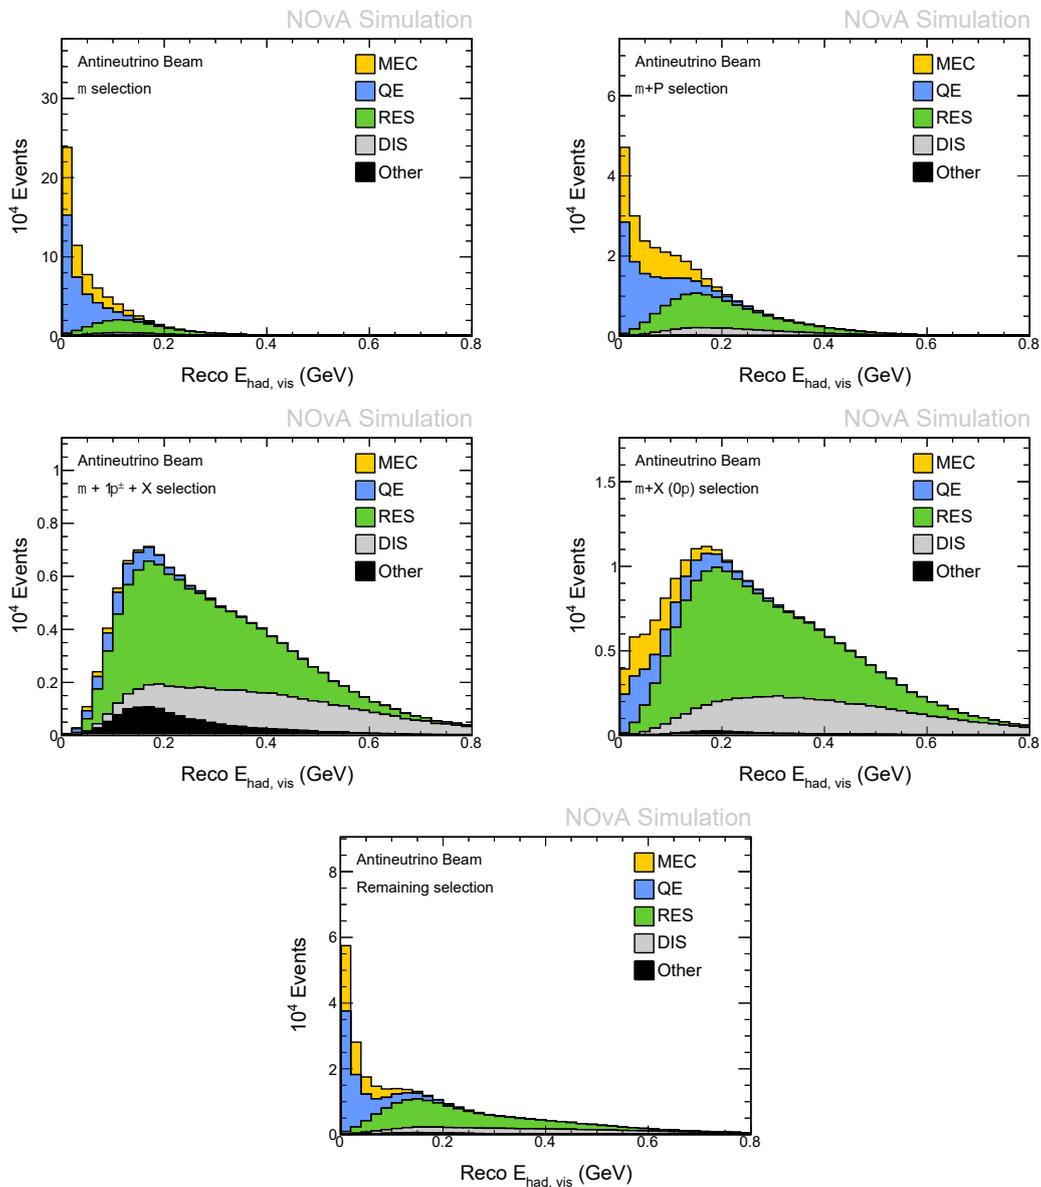


Figure 4.9: Projection of the 2D $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space onto the E_{had}^{vis} axis, projecting out $Reco |\vec{q}|$. These plots of the the RHC topologies break down neutrino events by true interaction in the simulation, confirming the model breakdown from Table 4.2. The “Other” category is predominately COH scattering. 1D RHC projections onto the $Reco |\vec{q}|$ variable are available in the Appendix A.3

values of E_{had}^{vis} and $Reco |\vec{q}|$ are likely to produce an inelastic interaction, as more momentum is transferred to the hadronic system (and less to the leptonic system). To further illustrate this point, we look at the 1D projections from these 2D distributions. Figure 4.8 (FHC) and Figure 4.9 (RHC) show the 2D distributions projected onto the E_{had}^{vis} axis (generally more intuitive than $Reco |\vec{q}|$) while also displaying the true interaction type from the NOvA simulation.

First, we recall Fig. 2.1 where we see nontrivial predictions for all four neutrino interaction processes in NOvA's 2 – 5 GeV energy range. Looking at Figs. 4.8 & 4.9 and the QE model, we see these events occur at low values of hadronic visible energy, across all topologies. Notably, the Muon topology is heavily dominated by QE and MEC (yellow) events with very little RES (green) and DIS (grey) contributions. Meanwhile, higher E_{had}^{vis} values across all topologies coincide with more inelastic neutrino interactions – namely RES and DIS interactions. These distributions highlight the challenge in producing a ν_μ selection that is pure with a single interaction type, and, conversely, highlights the overlapping of models in this energy range. However, the $\mu + \pi^\pm + X$ and $\mu + P + X$ topologies are moderately pure samples of RES and DIS interactions with little QE and MEC contribution. 1D distributions onto the $Reco |\vec{q}|$ variable can be seen in the Appendix A.2 and A.3. The summary of the dominant interaction types can also be seen in Table 4.2.

In addition, Fig. 4.10 (FHC) and Fig. 4.11 (RHC) shows the same ND topological samples in projections of E_{had}^{vis} , but showing the true final state particles (1D projections onto $Reco |\vec{q}|$ can be found in the Appendix A.4 and A.5). Understanding the breakdown of final state particle combinations is highly useful; for example, in the $\mu + \pi^\pm + X$ topology, it is highly pure in $\mu + \pi^\pm + X$ final state particles. Moreover, the challenge of prong-CVN misidentifying particles is another complexity of this analysis that must be considered when analyzing these topologies, and furthermore analyzing the results from MCMC sampling. It presents a challenge to assess NOvA's neutrino interaction modeling, and, by extension, NOvA's uncertainties of these models. This is a primary objective of this thesis, as we seek to

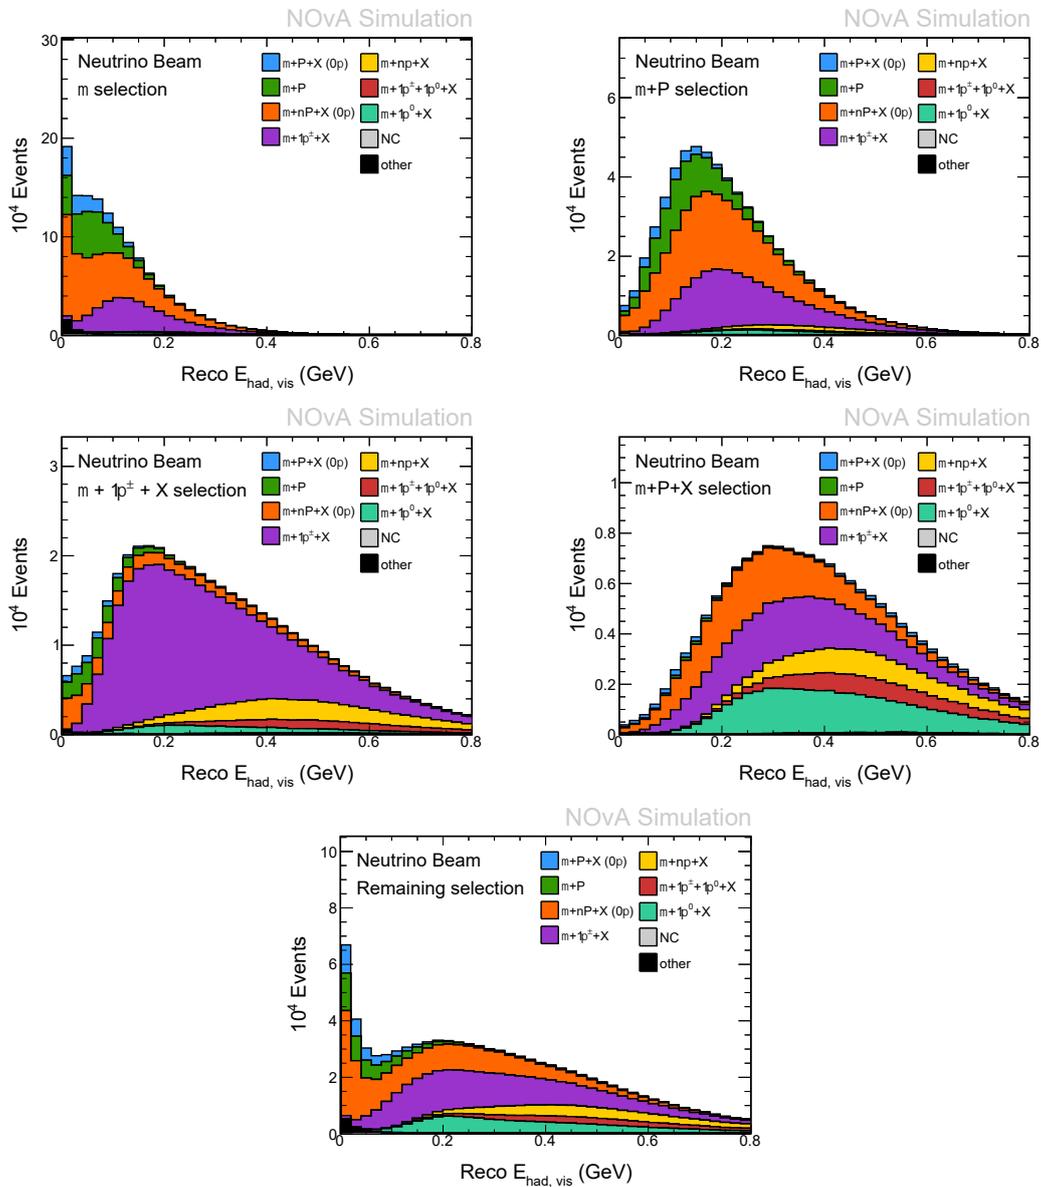


Figure 4.10: Projection of the 2D $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space onto the E_{had}^{vis} axis. These plots of the FHC topologies break down neutrino events by the final state particles from the interaction. 1D FHC projections onto the $Reco |\vec{q}|$ variable are available in the Appendix A.4.

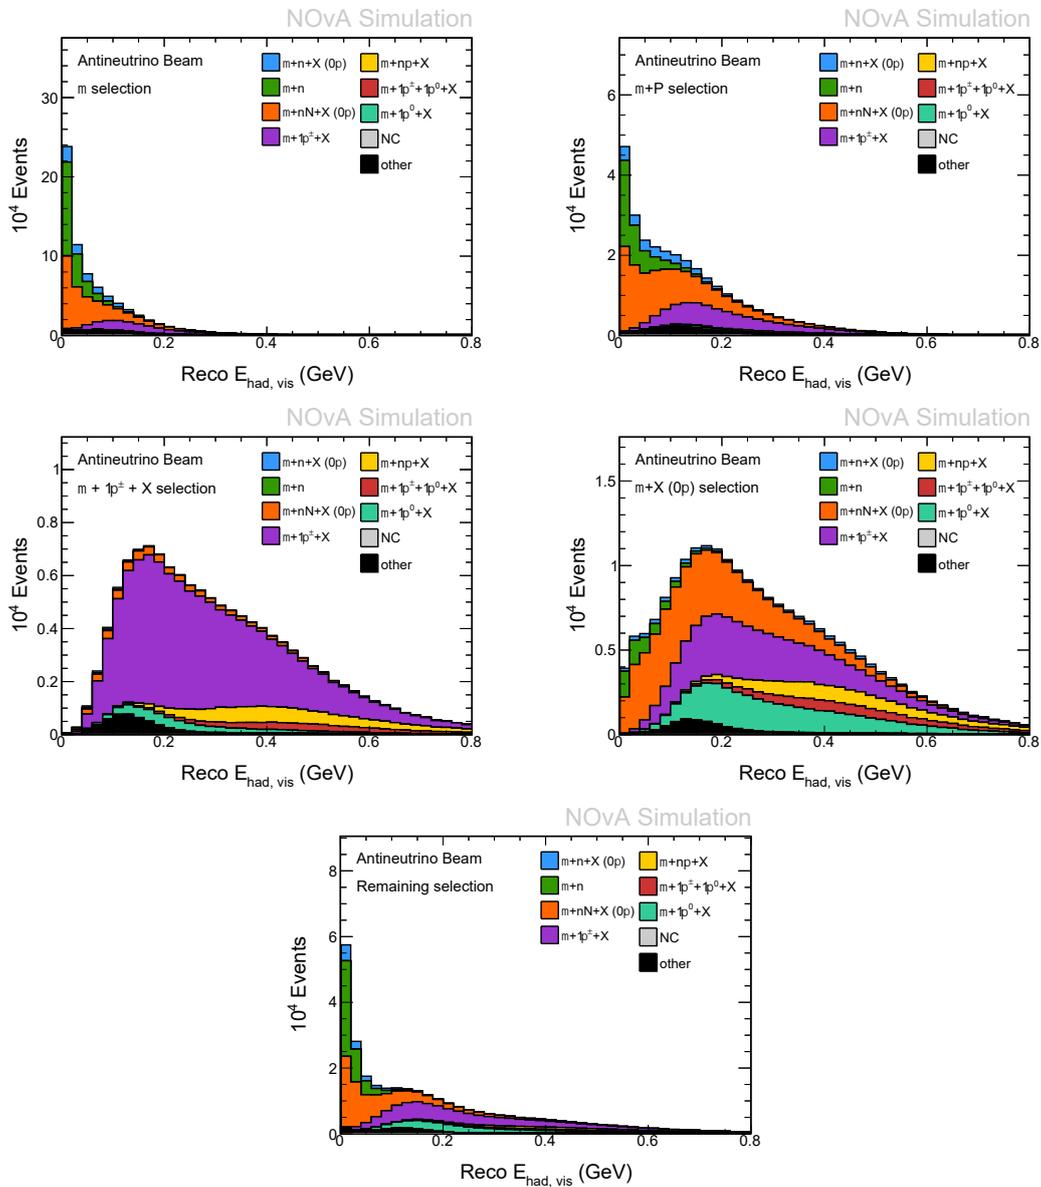


Figure 4.11: Projection of the 2D $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space onto the E_{had}^{vis} axis. These plots of the RHC topologies break down neutrino events by the final state particles from the interaction. 1D FHC projections onto the $Reco |\vec{q}|$ variable in the Appendix A.5.

understand where the interaction modeling can be improved within NOvA.

4.4 Far Detector Selection

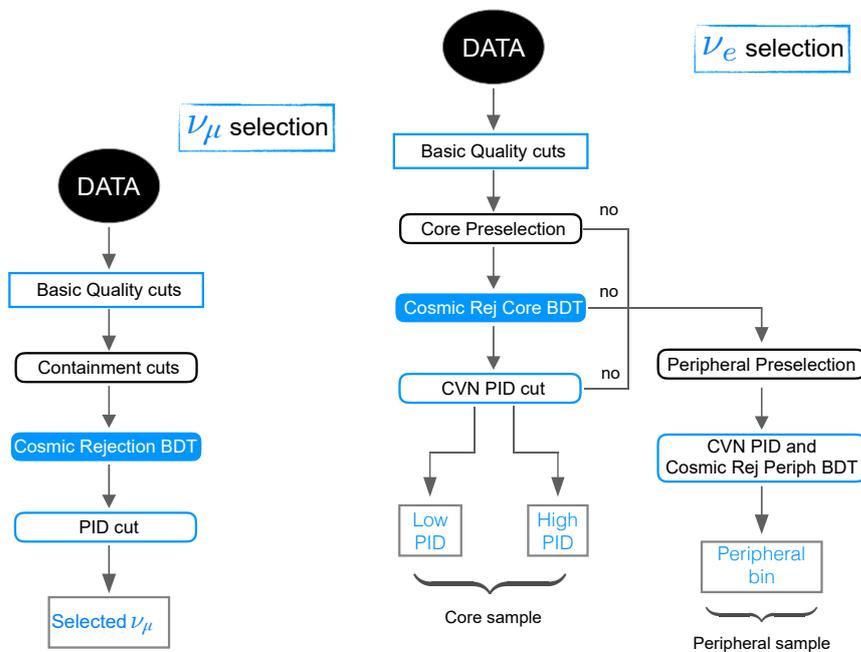


Figure 4.12: Flow charts for ν_μ and ν_e selection and reconstruction.

We must now select the FD neutrino interactions. A summary of the reconstruction and selection algorithm for ν_μ and ν_e is shown in Figure 4.12. We select ν_μ and ν_e events in the FD. The following sections describe the details of these processes. The ν_μ and ν_e events are selected from a different set of criteria. Subsection 4.4.1 outlines the selection process used for the FD ν_μ events, while Subsection 4.4.2 describes the process for the ν_e sample.

4.4.1 ν_μ Quantile Selection

One of NOvA's objectives is to measure ν_μ disappearance; the selection here is different from the ND topological selection. First, recall that to observe neutrino oscillations we must estimate E_ν (Eqs. 1.13 and 1.18) as oscillations are a function of L/E (where $L = 810$ km in NOvA).

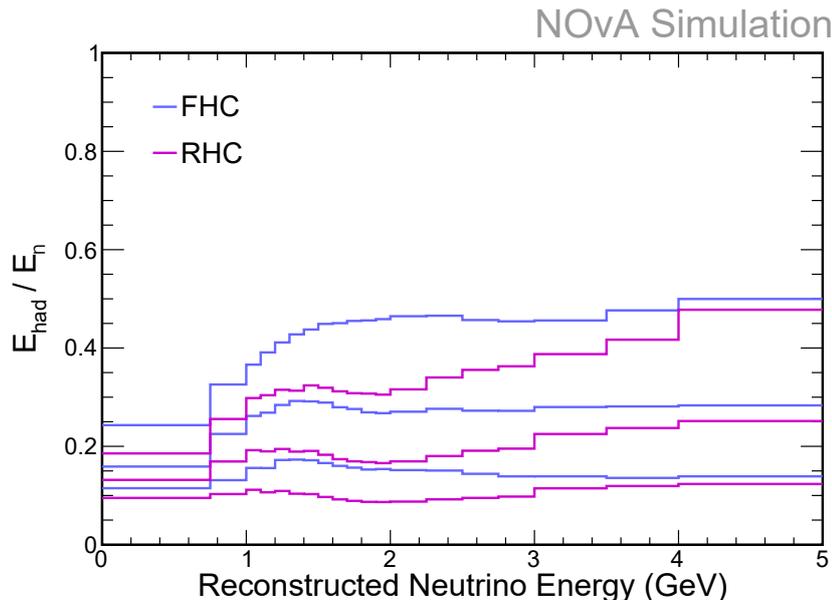


Figure 4.13: Distribution of the quantile cuts for the FD ν_μ events. In each bin of $Reco E_\nu$, the FHC (blue) and RHC (magenta) lines divide the total number of events into equal populations of events. The bottom-most blue & magenta line demarcates the first quantile, Q1. Quantile 1 contains the best energy resolution of 6%, containing the least amount of hadronic energy. The selection cuts used for this analysis are used from NOvA’s 2020 oscillation measurement analysis [56].

As stated, the characteristic feature of a ν_μ CC event is the long muon track in the detector. There are several preliminary selection cuts made to the FD ν_μ events observed in the NOvA detectors to estimate $Reco E_\nu$:

Basic Quality: identify a muon track with at least 20 hits.

Containment: ensure the muon track is located entirely within the fiducial volume (60 cm to the top, 12 cm to the bottom or west, 16 cm to the east, 18 cm to front or back).

Cosmic Rejection: reject cosmic ray particles.

PID: the output of CVN_μ must be greater than 0.5 (and for ReMID).

The next set of cuts are known as “quantile” cuts. The FD ν_μ events are grouped into four quantiles (also known as quartiles). These quantiles divide each bin of the inclusive FD unoscillated CC ν_μ energy distribution into four groups with an equal number of events. Specifically, this division of events in each bin is made using the

the ratio of hadronic energy fraction with respect to the total neutrino energy; this hadronic energy fraction, E_{frac}^{had} , is defined as

$$E_{frac}^{had} = E_{had}/E_{\nu}, \quad (4.2)$$

but also more generally referred to as the inelasticity variable y , Eq. 2.5. A plot of the quantile boundaries can be seen in Figure 4.13. The FHC (blue) and RHC (magenta) lines denote the 25% population of events within each bin of $Reco E_{\nu}$, as this process is done for both FHC and RHC separately. Quantile 1 contains the 25% of events in each bin of $Reco E_{\nu}$ with the least amount of hadronic energy fraction; this is the bottom-most line for each FHC (blue) and RHC (magenta) in Figure 4.13. For example, in the RHC (magenta) case, the boundaries of the first two quantiles in each of the bins of E_{ν} are made at roughly constant values of $E_{frac}^{had} \approx 0.1$ and 0.2 respectively. Meanwhile, Quantile 4 contains the most amount of hadronic energy – and the worst resolution – to measure the oscillation parameters; the top-most line. The resolution of the four quantiles from Quantile 1 to Quantile 4 is 6%, 8%, 10%, and 12%.

The result of these cuts is shown in Figures 4.14 (FHC) and 4.15 (RHC); the FD oscillated MC CC ν_{μ} predictions are plotted in E_{ν} in each of the four quantiles using the oscillation parameters from the NOvA 2020 measurement. The strength of these cuts is evident by examining the shape of the “dip” region in the oscillated ν_{μ} spectrum. Note how the oscillation “dip” at approximately 1.5 GeV is more pronounced in Quantile 1 for both FHC and RHC than for Quantile 4 demonstrating the superior resolution; Quantile 1 is most sensitive to ν_{μ} disappearance. The non-uniform binning of the FD ν_{μ} spectra is done to emphasize the effects of the oscillation dip region. The bin widths are small at 1-2 GeV to capture this effect, while the larger bins at higher energy are larger as NOvA does not observe ν_{μ} oscillations at higher energies. This is the same binning scheme used as in prior NOvA analyses [55].

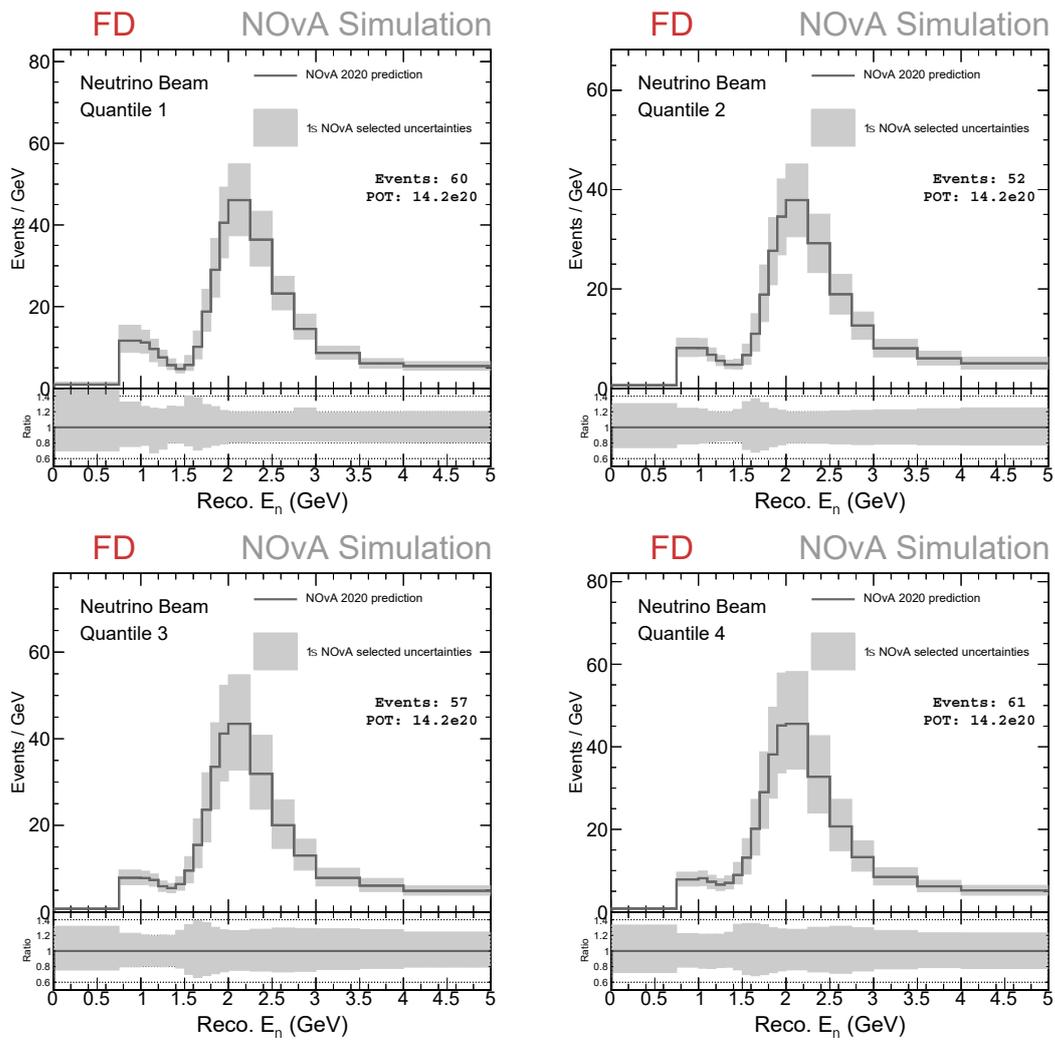


Figure 4.14: FD FHC ν_μ spectra divided into the quantiles for the Prod5.1 NOvA MC. Note the more pronounced dip region in Quantile 1. This prediction uses the NOvA 2020 best fit of the oscillation parameters. The grey bands show the 1σ uncertainty band for the uncertainties to be used in the MCMC sampling.

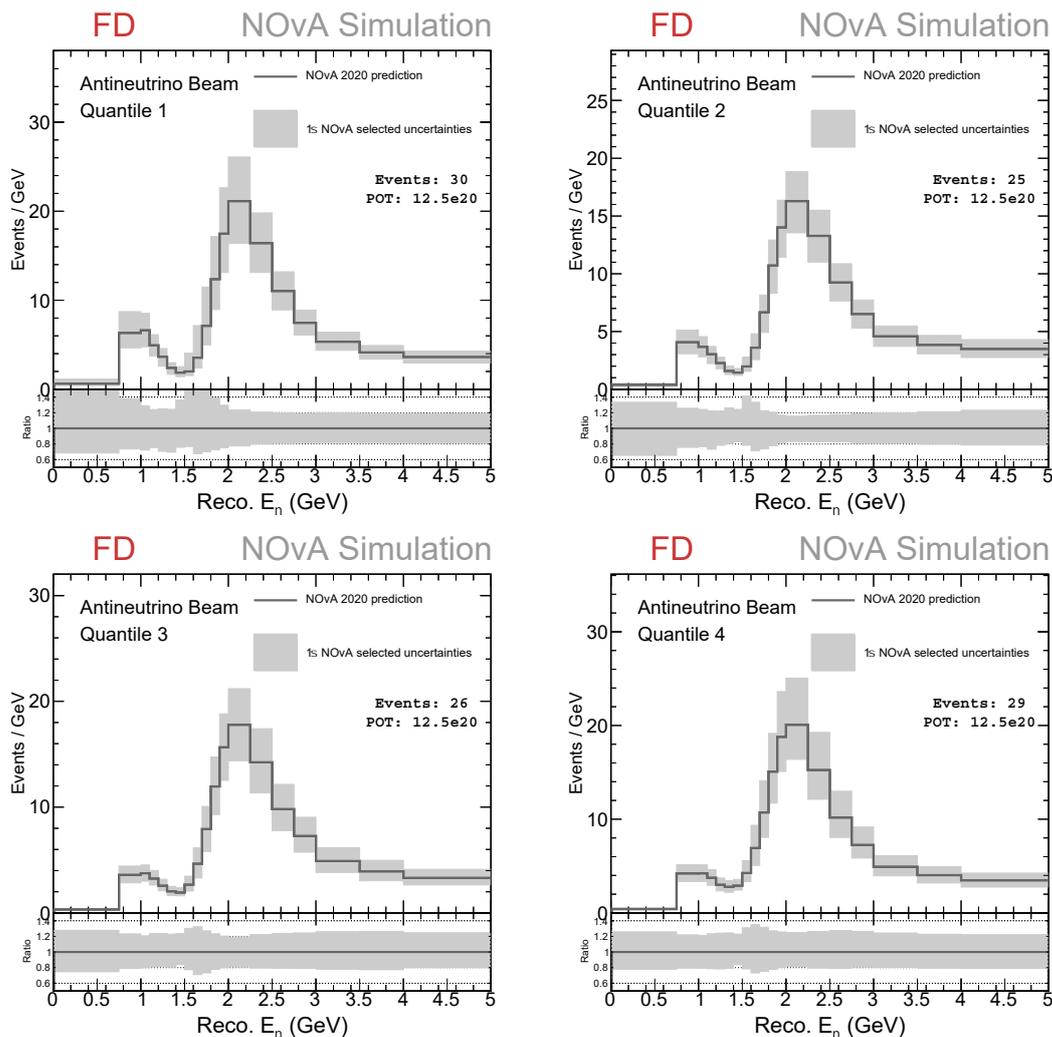


Figure 4.15: FD RHC ν_μ spectra divided into the quantiles for the Prod5.1 NOvA MC. Note the more pronounced dip region in Quantile 1. This prediction uses the NOvA 2020 best fit of the oscillation parameters. The grey bands show the 1σ uncertainty band for the uncertainties to be used in the MCMC sampling.

4.4.2 ν_e Selection

The ν_e selection contains its own set of cuts, from Figure 4.12. In order to maximize the possible number of appeared ν_e events, NOvA categorizes different classes of ν_e events: “Low PID”, “High PID”, and “Peripheral”. The logical flow of the cuts and how an event is classified is outlined below.

Basic Quality: removes substantial background or reconstruction failures.

Core Preselection: remove events that have hits are close to the edge of the detector (60 cm to the top, 12 cm to the bottom or west, 16 cm to the east, 18

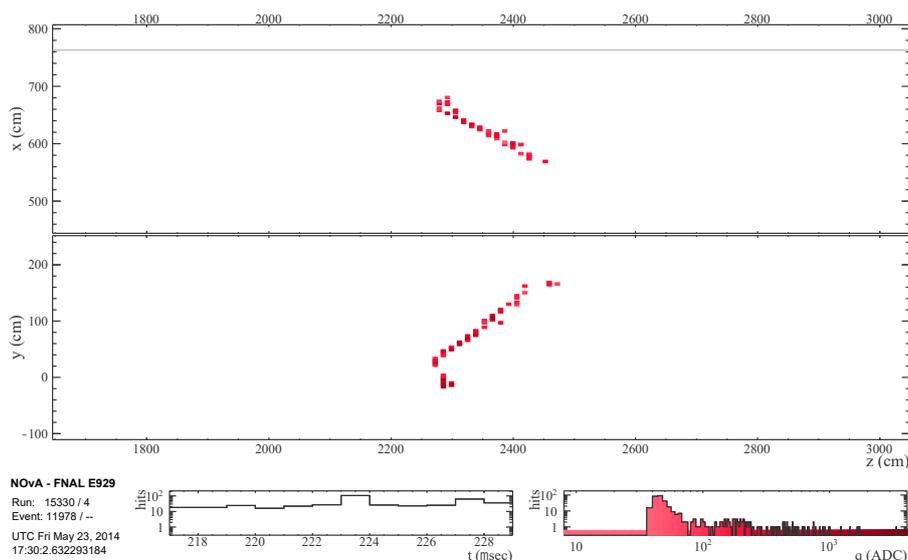


Figure 4.16: ν_e event at the NOvA FD with $Reco E_\nu = 2.29$ GeV from [57].

cm to front or back) – as they are likely background events. Require $Reco E_\nu$ events between 1-5 GeV.

Cosmic Rejection: Remove cosmic events that have tracks which are vertically aligned.

CVN PID: the output of the CVNe classifier is > 0.75 .

The Basic Quality cut removes any clear background or reconstruction failures from the sample. The Core Preselection cut removes the events that are on the periphery of the detector, which would potentially impact the $Reco E_\nu$ estimate. The next cut is the Cosmic Rejection which removes tracks that are vertically aligned (via p_T/p), which tend to be dominated by cosmogenic particles. The last cut is the CVN PID cut, where the CVN_e , the ν_e classifier to identify electromagnetic-like prongs, score must be larger than 0.75. At this stage if the candidate event passes this cut, it will either be grouped into the “Low PID” or “High PID” category depending on its score – known as the “Core” sample. Events that fall into the “High PID” region are the strongest ν_e candidate events. Events that do pass the Basic Quality cut but do *not* pass any of the remaining cuts are candidates to move into the “Peripheral” sample. The “Peripheral” sample is a single bin where $Reco E_\nu$ is not considered. For a ν_e candidate, the event pass an additional cut, BDT Cosmic Rejection.

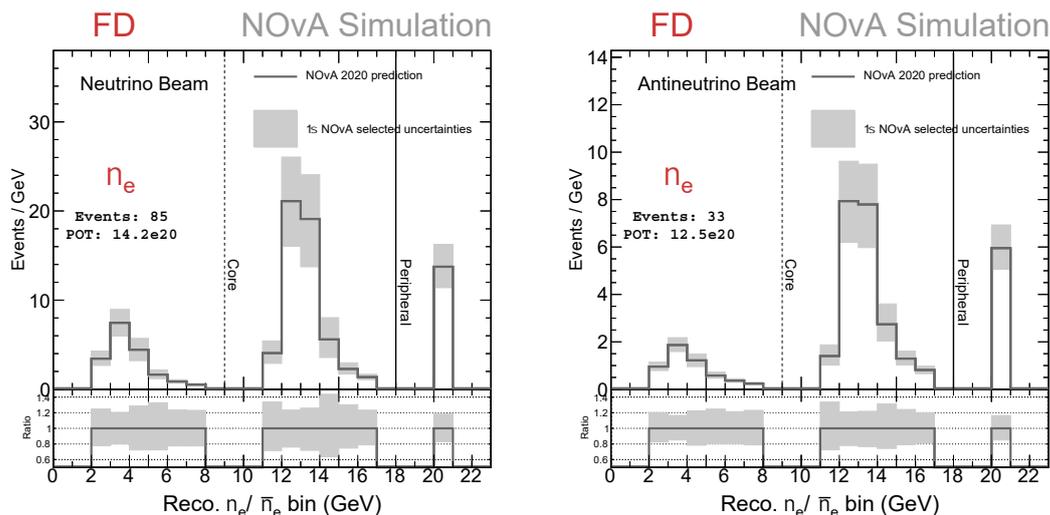


Figure 4.17: FD ν_e spectra divided for the Prod5.1 NOvA MC. Predictions are made using the NOvA 2020 oscillation analysis result. The grey band represents the 1σ error of systematic uncertainties to be used in MCMC sampling.

BDT Cosmic Rejection a Boosted Decision Tree classifier further eliminates cosmic background from events.

A Boosted Decision Tree (BDT) is used to further remove cosmic events from these Peripheral ν_e candidates. The BDT score is > 0.53 for the majority of events. These events are then merged together into the single bin, where $Reco E_\nu$ is not considered.

The result of the selection process is shown in Figure 4.17 where the FD ν_e sample is plotted using the NOvA 2020 oscillation measurements in FHC (left) and RHC (right). In each plot, the left most distribution is the Low PID region, the middle is the high PID region, and the right most bin is the Peripheral sample. Note the majority of ν_e events appear in the High PID region. The grey error band represents the 1σ error of uncertainties used in this analysis.

With a description of the event reconstruction, ND topological sample selection, FD ν_μ quantile samples, and FD ν_e sample selection, we can proceed to perform Markov Chain Monte Carlo sampling to constrain the NOvA MC to best agree with the NOvA data. We will begin with constraining the NOvA ND information only, first with ND fake data and then ultimately with the ND data. However, first we will summarize Bayesian inference and the MCMC sampling algorithm in the next chapter. In the following chapter the results of MCMC sampling on the ND

information only are analyzed.

Chapter 5

Bayesian Inference and Markov Chain Monte Carlo

Bayesian inference, proposed by Thomas Bayes (1701-1761) [58], infers the probability of true parameter values in a model from information – also known as data – provided by calculating a posterior probability. This is done with assumptions of the possible parameter values – the prior – and a prediction’s expected distributions given specific parameter values – the likelihood.

This chapter begins by discussing the details of Bayes Theorem in Section 5.1. The discussion continues in Section 5.2 with Markov Chain Monte Carlo (MCMC). Section 5.3 describes the specific algorithm used in this thesis to estimate the posterior probability – Hamiltonian Monte Carlo. The following Section, 5.4, describes a 2-dimensional example of Hamiltonian Monte Carlo for illustrative purposes. Section 5.5 discusses a metric we use to quantify agreement between the NOvA data and MC. The last Section, 5.6, briefly outlines the software Markov Chain Monte Carlo package, Stan, and how it is used in this analysis.

5.1 Bayes Theorem

Bayes Theorem is a simple equation written as

$$\mathcal{P}(\vec{\theta}|\text{data}) = \frac{\mathcal{P}(\text{data}|\vec{\theta}) \times \mathcal{P}(\vec{\theta})}{\mathcal{P}(\text{data})}. \quad (5.1)$$

The objective of Bayes Theorem is to produce probabilities for a given hypothesis. The equation is broken down into its elements below.

Posterior

The *posterior probability*, or $\mathcal{P}(\vec{\theta}|\text{data})$, is the quantity we seek to estimate. It is the probability for values of the parameters, $\vec{\theta}$, of the model in consideration are true given the data provided. Ultimately we seek to identify the values of $\vec{\theta}$ which maximize the posterior probability. These are the most probable values of the parameters.

Prior

The $\mathcal{P}(\vec{\theta})$ is referred to as the *prior probability*, also known as the “prior”. It represents the knowledge (or lack thereof) of the true values of the parameters, $\vec{\theta}$, we want to determine before we have analyzed the data. Before performing the experiment, there may be factors which can shape the prior. For example, an extreme value of a parameter that does not produce a physically meaningful prediction; as we will see, this is true for NOvA systematic uncertainties. Another example is a previous measurement made that would constrain an oscillation or interaction parameter(s).

Likelihood

The other term in the numerator is the *likelihood*, $\mathcal{P}(\text{data}|\vec{\theta})$, or \mathcal{LL} . The likelihood assesses the data; it evaluates if a given set of model parameters is consistent with the data. In our case, this metric is calculated via a Poisson statistic χ^2 (i.e. $e^{-\frac{\chi^2}{2}}$) between the data and prediction. The combination of the likelihood and prior are typically referred to as the “model”.

Evidence

The evidence, $\mathcal{P}(\text{data})$ in the denominator, is the probability for observing the data. Importantly, because the values of the data are fixed, the probability of observation is fixed, and, therefore, to include it would yield the absolute posterior probability. However, because the goal of this thesis is to estimate the parameters that maximize the posterior (oscillation parameters and NOvA’s uncertainty parameters), we can disregard this term as it is simply a normalization constant, and instead calculate the relative probability of the posterior. Therefore Bayes Theorem can be simplified slightly into

$$\mathcal{P}(\vec{\theta}|\text{data}) \propto \mathcal{P}(\text{data}|\vec{\theta}) \times \mathcal{P}(\vec{\theta}), \quad (5.2)$$

where the evidence term is a constant and omitted.

This equation can also be interpreted as drawing inferences from a hypothesis:

$$\text{prob}(\text{hypothesis} | \text{data}) = \frac{\text{prob}(\text{data} | \text{hypothesis}) \times \text{prob}(\text{hypothesis})}{\text{prob}(\text{data})} \quad (5.3)$$

It can be read as “the probability of a hypothesis” (the prior) multiplied by the “the probability the data is consistent with the hypothesis” is equal to the “probability the hypothesis is true.” Understanding the equation in this fashion can make Bayes Theorem more approachable. In the next Subsections, some of the features of Bayesian inference are described further.

5.1.1 Credible Intervals

One of the features of Bayesian inference is the interpretations one makes from its statistics; such as the Credible Interval. MCMC is a *sampler* of the posterior probability density, and does not calculate the posterior analytically. Therefore, the quantity returned from MCMC are *samples*. These samples must then be converted into a probability density, and is discussed in more detail Section 5.2. This analysis

does this by binning the samples into 1D or 2D histograms over one or two of the parameters, respectively (or implicitly marginalizing out all other parameters, see 5.1.2). With this histogram, the 1σ credible interval, for example, can be formed by selecting the highest probability bins within the parameter space that give a cumulative probability of 68%. This process is identical for 2σ and 3σ by selecting 95% and 99.8% of the highest probability bins, respectively.

5.1.2 Marginalization

A second feature of Bayesian inference is that of marginalization (Equation 5.4). One can integrate out undesired parameters, or “nuisance parameters”. Marginalizing away these parameters, $\vec{\delta}$, is not trivial and provides a means of analyzing only parameters that we are interested in for this analysis, $\vec{\theta}$. In fact, this is one of the advantages to using MCMC, given the difficulty in marginalizing or profiling.

$$\mathbb{P}(\text{data}|\vec{\theta}) = \int \mathbb{P}(\text{data}|\vec{\theta}, \vec{\delta}) d\delta \quad (5.4)$$

This is something implicitly done when constructing credible intervals from Subsection 5.1.1; by creating a credible interval for one or two parameters, for example, one is marginalizing over, or integrating out, the remaining parameters. This is an incredibly powerful feature of MCMC as it allows one to investigate relationships between specific parameters at no extra computational cost.

5.2 Markov Chain Monte Carlo

We move onto the discussion of the principles of Markov Chain Monte Carlo (MCMC). The posterior parameter space, whose dimension corresponds to the number of parameters of $\vec{\theta}$, can be considered a coordinate location space. As a result, we will use the new notation, \vec{q} . The posterior space becomes a probability density

function over this coordinate space

$$\mathbb{P}(\vec{q}) = \pi(\vec{q}) \times \mathcal{L}\mathcal{L}(\vec{q}). \quad (5.5)$$

To illustrate this further, the expectation value is the posterior probability integrated over its volume,

$$\mathbb{E}_\pi = \int_Q \pi(\vec{q}) \mathcal{L}\mathcal{L}(\vec{q}) d\vec{q}, \quad (5.6)$$

where $\pi(\vec{q})$ is the prior, $\mathcal{L}\mathcal{L}(\vec{q})$ is the likelihood, and is integrated over some parameter space, Q . To maximize the expectation, \mathbb{E}_π , we want to maximize the posterior, $\mathbb{P}(\vec{q})$, and volume, dq . Identifying the region where the posterior function alone is

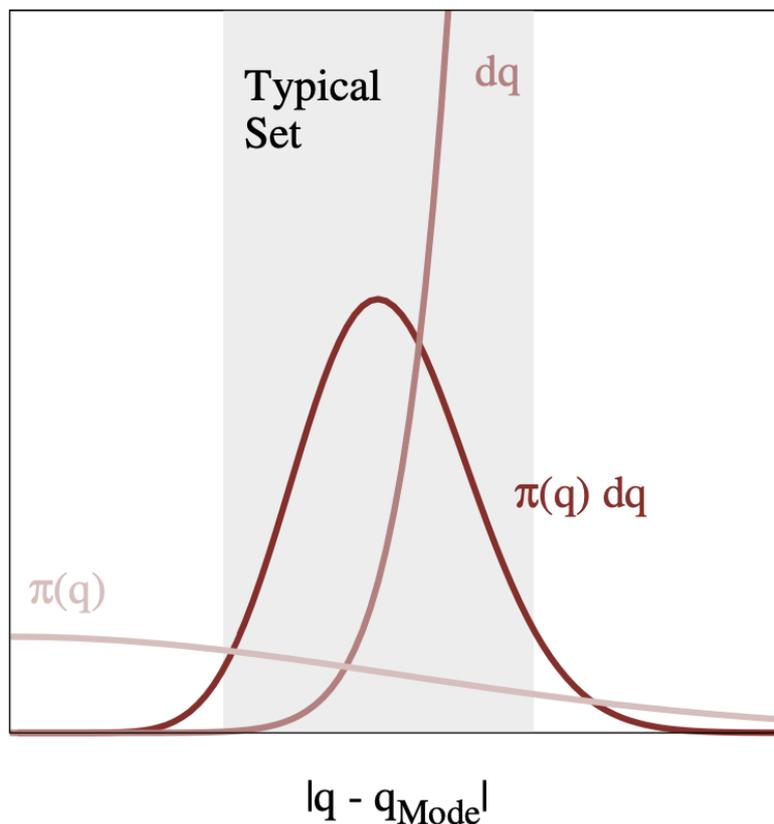


Figure 5.1: Representation of the typical set. The typical set is the region over which the *product* between the probability density function, $\mathbb{P}(\vec{q})$, and the volume element, dq , is maximized. Note in the figure the probability density function or posterior, $\mathbb{P}(\vec{q})$, is written as “ $\pi(q)$ ”. From [59].

large is insufficient; the *product* of the posterior and the volume, $\mathbb{P}(\vec{q})dq$ (equivalent

to “ $\pi(q)dq$ ” in the figure), maximizes the expectation, can be considered in Figure, 5.1. To help illustrate this further, in Figure 5.1, the grey shaded region of the illustration shows the region where the product between the posterior and volume element is maximal in the dark red distribution.

This high-probability region in the posterior is called the “typical set” – the region where the maximum expectation values exist within the posterior. To calculate the posterior within the volume of the typical set analytically can often be intractable when the number of parameters is large. An effective tool to approximate the posterior is to directly sample the posterior space proportional to its probability, thereby estimating the model’s parameter values. Markov Chain Monte Carlo (MCMC) performs this sampling to approximate the posterior and to do so more efficiently via a Markov Chain method.

The next step is to describe how a Markov Chain will explore the typical set. Within a Markov Chain are Markov Transitions,

$$\mathbb{P}(\vec{q}) = \int_Q \mathbb{T}(\vec{q}|\vec{q}') \mathbb{P}(\vec{q}') d\vec{q}', \quad (5.7)$$

where $\mathbb{T}(\vec{q}|\vec{q}')$ is the conditional probability to move to a new location \vec{q} and \vec{q}' is the starting location of the density function. We use the posterior (i.e. the probability density function) to determine the next location within the typical set to inform the chain how to move through the posterior space. One of the fundamental features of a Markov Chain is that it is reversible; the probability of moving between neighboring chain elements is the same from both directions. Running the Markov Chain for $N \rightarrow \infty$ transitions will reproduce exactly the expectation value within the typical set. However, infinite transitions requires an infinite amount of computing power, so we must estimate the expectation value with a finite number of transitions. We seek to reach the typical set as quickly and efficiently as possible. There are two common sampling algorithms: Metropolis-Hastings, a more common and general algorithm (which will be discussed briefly) and Hamiltonian Monte Carlo, a more computationally intensive algorithm that utilizes the geometry of the posterior space

(will be discussed in more detail) to estimate the typical set.

5.2.1 Metropolis-Hastings

The Metropolis-Hastings (MH) algorithm consists of two parts: a proposal and an acceptance criterion [60]. The proposal requires information of the current location in the posterior space and “proposes” a new location in the space; this proposal is typically sampled from a Gaussian distribution of N dimensions, Equation 5.8, where the dimension is equal to the the number of parameters in the posterior space and a step size matrix, Σ ,

$$\mathcal{N}(\vec{q}' | \vec{q}, \Sigma) = \frac{1}{\sqrt{2\pi^N \det(\Sigma)}} e^{\frac{1}{2}(\vec{q}-\vec{q}')\Sigma^{-1}(\vec{q}-\vec{q}')}. \quad (5.8)$$

The step sizes, Σ , set by the user, determine the distance of the probable jump distances the Markov chain can make between \vec{q}' and \vec{q} ; a smaller step size produces a more “fine-grained” exploration of the posterior at the cost of increased computational time. This N -dimensional Gaussian is used in calculating the acceptance (Equation 5.9),

$$A(\vec{q}', \vec{q}) = \min \left(1, \frac{\mathcal{N}(\vec{q} | \vec{q}', \Sigma) \mathbb{P}(\vec{q}')}{\mathcal{N}(\vec{q}' | \vec{q}, \Sigma) \mathbb{P}(\vec{q})} \right), \quad (5.9)$$

where the Gaussian (5.8) proposes steps and if the acceptance metric A is met, or exceeded, then the proposal is accepted and the algorithm moves to the new location in the posterior, \vec{q} , from \vec{q}' . If A is not met, a new proposal is sampled from the Gaussian. This process continues until the desired number of steps is reached.

The Metropolis-Hastings Algorithm is the most common Markov chain used within MCMC, however, with increasing parameter space MH can quickly become ineffective due to the number of possible directions (i.e. dimensions) the algorithm can move in. For this reason, Hamiltonian Monte Carlo presents an attractive alternative to MH exploration of the posterior space. Next we discuss the Hamiltonian Monte Carlo algorithm in more detail.

5.3 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (also referred to as HMC or HMCMC) is one particular algorithm to approximate the posterior and explore the parameter space by utilizing information of the posterior's geometry. We begin by treating the posterior space as a topological surface. The objective of MCMC remains the same – to sample the posterior by sampling from it proportionally to its probability. However, now considering the topological space and its intrinsic characteristics with its local minima and maxima – associated with less and more probable values, respectively – we introduce a potential and kinetic energy. We can then explore the space using kinematics, namely Hamilton's equations

$$\frac{d\vec{q}}{dt} = \frac{\partial H}{\partial \vec{p}}, \quad \frac{d\vec{p}}{dt} = -\frac{\partial H}{\partial \vec{q}}. \quad (5.10)$$

When applying Hamilton's equations to the posterior space, \vec{q} and \vec{p} represent the position and momentum, respectively, of the Markov chain's trajectory in the parameter space, while H is the Hamiltonian of the Markov chain. However, the momentum and Hamiltonian have not yet been defined. Starting with the momentum, if we introduce a corresponding auxiliary momentum within the posterior to describe the motion of the trajectory, every probability density can be located by specifying its location and momentum. Referred to as a joint probability of the posterior,

$$\mathbb{P}(\vec{q}, \vec{p}) = \mathbb{P}(\vec{p} \mid \vec{q}) \mathbb{P}(\vec{q}), \quad (5.11)$$

$\mathbb{P}(\vec{q}, \vec{p})$ can be written as a conditional probability where we can marginalize over the momentum, $\mathbb{P}(\vec{p} \mid \vec{q})$, and return to the probability density over only the parameters, $\mathbb{P}(\vec{q})$. We must also define the Hamiltonian in the parameter space. The

Hamiltonian, naturally, should contain information of the posterior space,

$$H = -\log \mathbb{P}(\vec{q}, \vec{p}) \quad (5.12)$$

$$= -\log \mathbb{P}(\vec{p} \mid \vec{q}) - \log \mathbb{P}(\vec{q})$$

$$H = T(\vec{q}, \vec{p}) + V(\vec{q}) \quad (5.13)$$

where $T(\vec{q}, \vec{p})$ is the kinetic energy term from the joint probability, and we use the logarithm to obtain a smoother posterior space. We use the Euclidean-Gaussian kinetic energy distribution

$$T(\vec{q}, \vec{p}) = \frac{1}{2} \vec{p}^T M^{-1} \vec{p} + \log |M| + \text{const.} \quad (5.14)$$

where M is referred to as the mass matrix because of the analogous role it plays in gravitational systems. Sometimes M^{-1} is a diagonal matrix where no parameter is correlated with one another. However, the more common use case is a dense mass matrix that has off-diagonal elements to quantify the correlations between any pair of parameters. One such choice of the mass matrix is to provide a correlation matrix to describe the off-diagonal terms of the matrix. The potential energy, $V(\vec{q})$, is the probability density over the parameters. Returning to Hamilton's equations,

$$\frac{d\vec{q}}{dt} = \frac{\partial H}{\partial \vec{p}} = \frac{\partial T}{\partial \vec{p}}, \quad (5.15)$$

$$\frac{d\vec{p}}{dt} = -\frac{\partial H}{\partial \vec{q}} = -\frac{\partial V}{\partial \vec{q}} \quad (5.16)$$

and observe the gradient of the potential energy, $\frac{\partial V}{\partial \vec{q}} = \frac{\partial \vec{p}}{\partial t}$, can be used to explore the parameter space more efficiently because the gradient of the joint probability will always point towards increasing probability. Thus, the Markov chain's exploration of the posterior is far more efficient than the MH algorithm; this is the critical concept in HMC – use of the potential to calculate a gradient to guide sampling. The second critical element of Hamilton's equations is the conservation of phase

space, which allows MCMC to explore the posterior probability distribution with conserved probability. In other words, the auxiliary momentum does not experience dissipative forces – like friction – acting in this system that would prevent the phase space from being conserved.

To ensure a trajectory does not simply oscillate around a maximum in the posterior, a condition is incorporated into the sampling referred to as the No-U-Turn Sampler (NUTS) condition [61]. Upon reaching the end of a gradient, a trajectory may turn back on itself, retracing its own path. When this condition occurs, the trajectory is terminated via the NUTS criteria.

There are several drawbacks to Hamiltonian Monte Carlo and Bayesian sampling in general. For example regions of high curvature in the posterior space can lead to divergences when calculating the gradients of the potential energy (Eq. 5.16), making the remaining samples difficult to interpret. In addition, if, for example, a parameter has two modes but are separated by in parameter space by zero probability, then it is unlikely an MCMC trajectory can move across this region of zero probability and sample the two modes sufficiently, resulting in an artificial bias, or even unable to sample the second mode at all. Another challenge to MCMC sampling is the lack of an effective means to locate the true maximum of the posterior, like in many probability problems; it is difficult to determine if a maximum is a local maximum or a global one.

5.4 Hamiltonian Monte Carlo 2-D Example

The full algorithm of HMC can be best understood with a 2-dimensional example – two parameters X and Y . If we consider the shape of the 2D posterior probability density in Figure 5.2, the grey contour lines surrounding this 2D “banana-shaped” posterior represent lines of equipotential from Hamilton’s equations. An object within this space will be influenced by this 2D potential; it will follow along the gradient of the potential, $\partial V/\partial \vec{q}$, and this quantity will direct the object towards less potential and therefore an increase in kinetic energy, or momentum. Therefore,

the object is more likely to traverse this region which corresponds to high probability. This is how the geometry of the space itself is used to explore the posterior more efficiently.

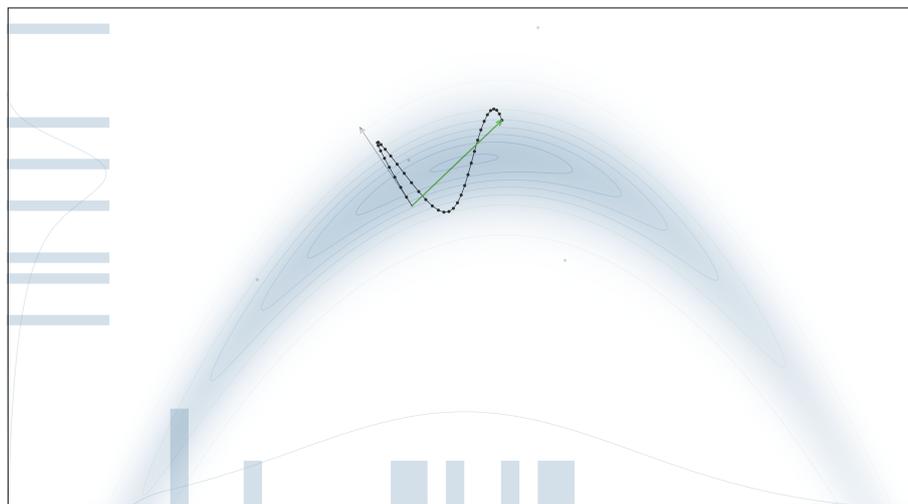


Figure 5.2: Hamiltonian Monte Carlo 2-dimensional example of a single MCMC chain exploring the posterior. The grey arrow indicates the initial “momentum” from Hamilton’s Equations. The final dot which the green arrow points to is the ending point of the chain, indicating the chain has achieved the No U-Turn Sampling (NUTS) criteria. From [62].

Looking further at Figure 5.2, the black dots show the trajectory of a MCMC sample under the influence of the potential (grey contour lines). The starting MCMC sample is the black dot with the grey arrow pointing from it. This grey arrow indicates the randomly selected *initial* momentum of the starting sample. In this case, the trajectory of the next several samples are nearly in-line with the initial momentum vector. Meanwhile, the black dot the green arrow points to marks the end of the chain where the NUTS criteria has been met. It is the end of the chain that marks where the MCMC sample is recorded. As noted, this location in the parameter space corresponds to specific values of the X and Y parameters, sampled proportional to their probability. Along the X and Y axes are 1D probability distributions of the individual parameters, obtained from marginalizing. This sample fills a bin within each of the 1D marginal distributions along the X and Y axes, indicated by the grey bars. Also along the axes are the true underlying 1D probability distributions drawn in light grey. In this Figure, there are only several bins

filled – samples recorded – in this parameter space. In the limit of infinite samples, the 1D marginalized histograms would perfectly reproduce the true 1D, light-grey probability distributions, also illustrating the need for high statistic samples.

5.5 The “Representative Sample” (Max. \mathcal{LL})

It is helpful to understand Bayes Theorem in the context of NOvA. Often in NOvA, and in the high-energy physics community, the χ^2 metric is useful to describe the goodness of a fit between data and MC. For the ND and ND+FD Bayesian fits in this analysis, we will report a χ^2 . We assess the likelihood via a test statistic using Poisson statistics

$$\mathcal{P}(\text{data}|\vec{\theta}) \equiv e^{-\chi^2} \equiv -2 \mathcal{LL} \quad (5.17)$$

where \mathcal{LL} is the log-likelihood of the data-MC agreement. In this statistic, the χ^2 takes the form of a χ^2 function in the limit of large event rates,

$$\chi^2 = \sum_i -2 \left(-\lambda_i + x_i + x_i \log \left(\frac{\lambda_i}{x_i} \right) \right), \quad (5.18)$$

where λ_i and x_i represent the expected (simulated) and observed (data) event rate in the i -th bin [4]. In the case of the ND-only MCMC sampling, these are bins of the 2D distribution in $(Reco | \vec{q}, E_{had}^{vis})$ phase space, while in the FD these are bins of $Reco E_\nu$.

To maintain consistency within the community, we report a χ^2 by introducing the “Representative Sample” in this analysis. The “Representative Sample”, **or maximum \mathcal{LL}** , is the sample that minimizes the χ^2 function. However, it is crucial to note that Bayesian samples does not produce a “best fit”. Instead, the “best fit” from Bayesian sampling would be the maximum of the posterior, which is the Representative Sample (max \mathcal{LL}). However, if one were to marginalize the posterior space, the max \mathcal{LL} sample might not be the same as the densest region of MCMC samples. This statement is true when marginalizing any number of dimensions of

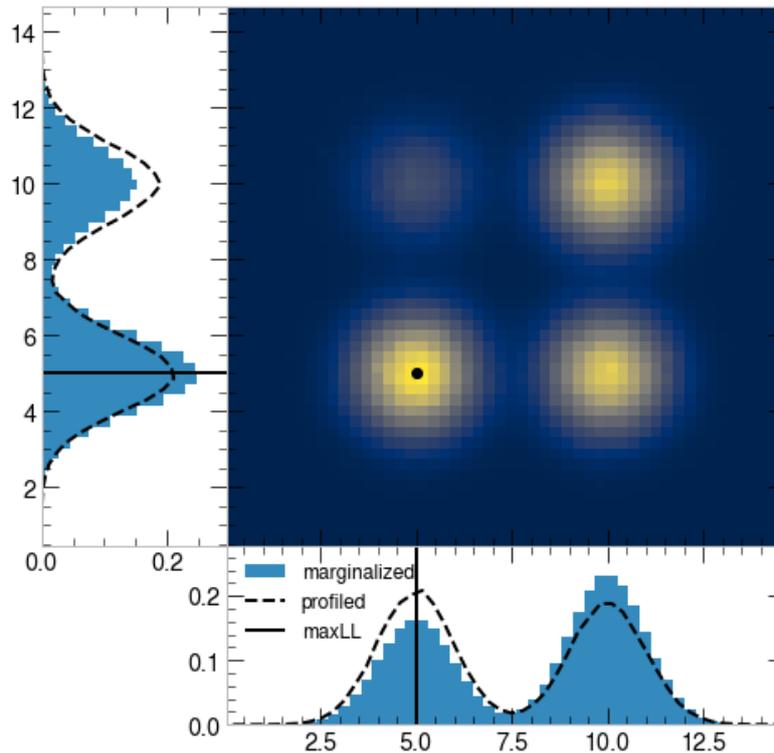


Figure 5.3: A 2D example of marginalizing vs. profiling. The black dot represents the maximum \mathcal{LL} point at $(5, 5)$. Profiling over the y variable (dashed line) yields $x = 5.0$ as the most probable value, while marginalizing over the y variable (blue histogram) produces $x = 10.0$ as the most probable value. This inconsistency between the max \mathcal{LL} ($x = 5.0$) and most probable marginalized value ($x = 10.0$) shows the “Representative Sample” (max \mathcal{LL}) may not always be the same as the highest probability point after marginalizing over parameters.

the posterior space; the maximum \mathcal{LL} over all parameters may not necessarily be equal to the marginalized maximum, as we see in a 2-dimensional example in Figure 5.3.

In Figure 5.3 there are four modes in the parameter space. The densest and most probable region is the bottom left quadrant, the two right-most dense regions are of equal probability. The black dot represents the maximum \mathcal{LL} . If one were to marginalize over the y variable, we see the Representative Sample is at $x = 5.0$, but the highest probability x value is at $x = 10.0$. This difference in the max \mathcal{LL} and the maximum marginalized probability concisely illustrate the subtlety in interpreting these two quantities; after marginalizing over any number of parameters, the highest probability point(s) from the marginal distribution may not be identical to the Representative Sample value, which is the maximum over the whole posterior

space. This will be important throughout this analysis.

Moreover, the Representative Sample resembles a *type* of “best fit”, but is not equivalent to a Frequentist “best fit”. This is also shown in Figure 5.3. This is because a Frequentist “best fit” is obtained via profiling, shown as a dashed line in the Figure. In this method, if we want to profile out the y variable, we select the y value with the maximum \mathcal{LL} , rather than average over the variable. In this way, we obtain a different probability distribution for the x variable; one where the profiled distribution, or “best fit”, in the x variable is consistent with the max \mathcal{LL} black dot. Despite the incongruent meanings of “best fit” and “Rep. Sample”, we still seek to report a χ^2 metric. The “Rep. Sample” is still a useful – but approximate – means of quantifying the strength of the agreement between (fake) data and MC from MCMC sampling; thus we will continue to report its value.

5.6 Stan Software Package

This analysis uses a third party software tool Stan, a Markov Chain Monte Carlo sampler [63]. Stan contains an interface for multiple different high-level programming languages, such as python, R, and C++. This analysis utilizes the C++ interface where it is adjoined to the NOvASoft software framework, CAFAna [64]. From there, an analysis with MCMC, like this one, can be done with built-in analysis, physics-specific NOvASoft code, while eliminating the need for the user to directly interface with the source code of Stan.

With an overhead view of Bayesian statistics and Markov Chain Monte Carlo we are ready to progress onto applying MCMC to the NOvA Near Detector simulation to estimate the values of the uncertainties within NOvA’s uncertainty model. An important note regarding terminology as we progress next to MCMC sampling of the NOvA ND MC and data; an “uncertainty” within NOvA’s uncertainty model will henceforth be used synonymously with “parameter” or “systematic”. The next chapter will focus on the ND-only component to constrain the uncertainties, and then move onto constraining the oscillation parameters in the following chapter using

both the NOvA ND and FD information.

Chapter 6

Near Detector Fit

The first step to constraining the neutrino oscillation parameters requires fitting NOvA’s ND predictions to data in order to constrain NOvA’s systematic uncertainties. This chapter will discuss this process and describe the details of a fit to the NOvA Near Detector data. Section 6.1 describes and discusses the motivations for selecting the prior function. The following section, Section 6.2, walks through a NOvA Near Detector “fake data” fit of a small subset of NOvA’s uncertainties to demonstrate the proof of principle of fitting physics model parameters with MCMC to NOvA fake data. Section 6.3 shows the data and MC comparisons in the ND topological samples to provide context as to how we anticipate MCMC will improve the agreement. Section 6.4 describes initial attempts and challenges at fitting the NOvA ND data. Section 6.5 describes the work to introduce new model uncertainties for the RES and DIS systematics to achieve better agreement with the NOvA data. Next, Section 6.6 introduces the “Residual Difference” present following the ND fit to data. The last Section 6.7 summarizes the important conclusions made from fitting the ND data with MCMC.

6.1 Prior Choice for Uncertainties

In quantification of uncertainties we often discuss uncertainties in units of σ , or “pull”; each model parameter we are constraining has a default value with some

uncertainty. In a fit, a given parameter may produce better agreement to the data (or fake data) at a value different from its default, thus the parameter is “pulled”. In this analysis, we will discuss fit results of parameters in terms of “pulls” in units of σ that the fit takes the parameter value away from the default value.

In the Near Detector fit, a decision must be made to determine the functional form of the Bayesian prior for NOvA’s systematic uncertainties. A simple Gaussian ($\mu = 0, \sigma = 1$) is a practical choice of prior for the pull; the most probable outcome is a pull of 0.0σ and its probability quickly drops off for large negative and positive values of σ , indicating that NOvA has high confidence of its systematic uncertainties at the default value 0σ and more extreme predictions at larger pulls are less likely.

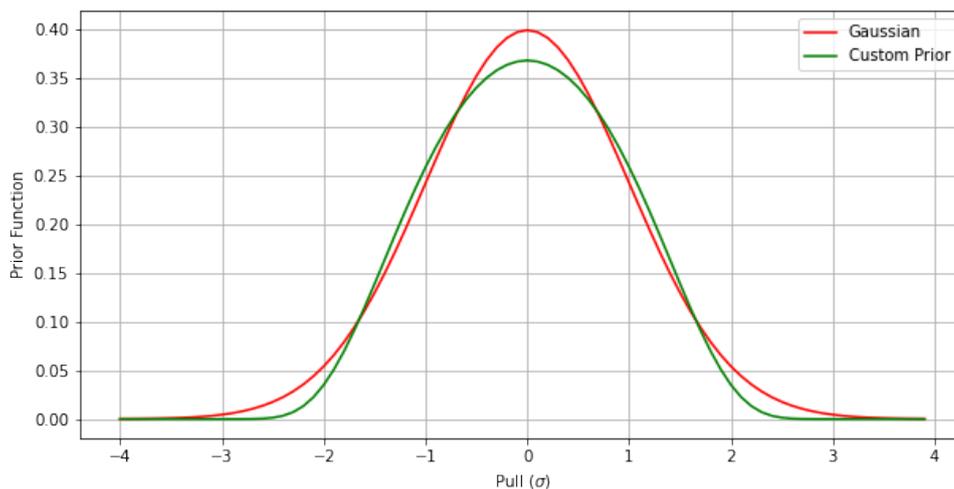


Figure 6.1: A Gaussian distribution in red. Prior function for NOvA’s uncertainties in green: $p(\sigma) = e^{-0.3025\sigma^2}$. This custom prior is differentiable for all σ values. Note the similarity in shape between the custom prior and the Gaussian except at the $\pm 2\sigma$ region, where the custom prior approaches zero at a smaller σ value than a standard Gaussian to restrict large pull values.

Many of the underlying models begin to exhibit nonphysical behavior if parameters are pushed beyond $\pm 3\sigma$, so limiting MCMC’s prior to this range ensures that our results are physically reasonable. This will allow us to better understand which model – and their uncertainties – may have fundamental problems if a parameter is pulled to an extreme value. For example, this could indicate the parameter’s uncertainty range may be too small, or the underlying physics model associated with the

given parameter may need attention.

With these two points in mind, we select a prior function that is Gaussian-like,

$$p(\sigma) = e^{e^{-0.3025\sigma^2}}, \quad (6.1)$$

where σ is the pull value of a given systematic uncertainty. The function maintains the characteristic of having the largest probability at the default value, 0σ , and restricts extreme pull values. This “custom prior” seen in Figure 6.1 (green) falls to zero probability for an even smaller pull value than a standard Gaussian – approximately $\pm 2.5\sigma$ rather than $\pm 3\sigma$ of the Gaussian. This further prevents a systematic from being pulled to an extreme (beyond $\pm 3\sigma$) value. Furthermore, our “custom” prior is also continuous and differentiable satisfying a necessary criterion for Hamiltonian Monte Carlo.

6.2 ND Fake Data Fits

Now we can perform a “fake data” fit to a subset of NOvA’s uncertainties. There are several motivations for performing a fake data fit. First, it allows us to test our fitting machinery by demonstrating MCMC can reproduce the exact, or known, pull values from the parameters in the fake data. Secondly, it will help us understand which parameters our ND data has the most sensitivity too, and which are poorly constrained in NOvA’s simulation. Lastly, we can investigate correlations between parameters by marginalizing over parameters and analyze the MCMC samples over pairs of parameters from the fit. This section will walk through a fake data fit with a subset of NOvA cross section uncertainties.

6.2.1 Fake data fit to QE, RES, and FSI Parameters

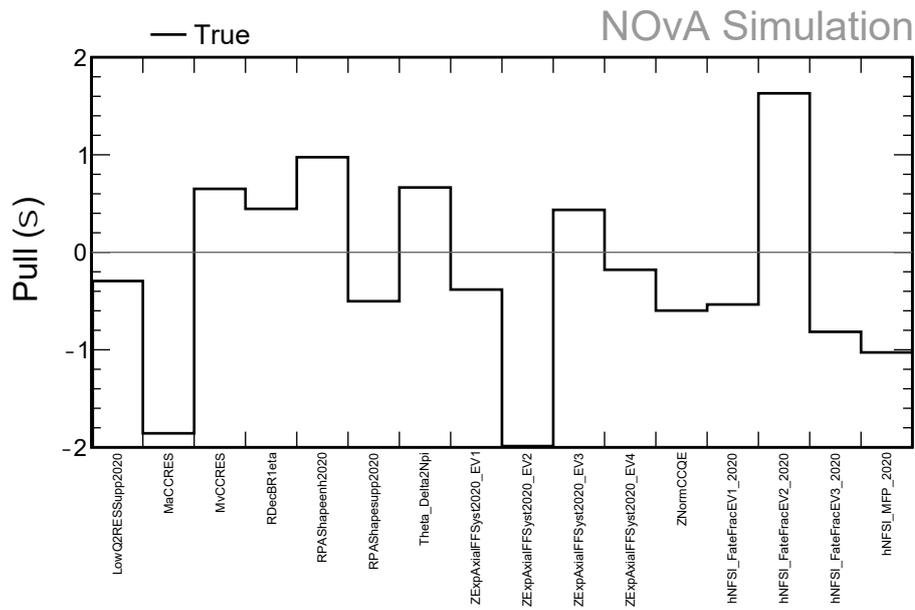
We produce a fake data MCMC fit to a subset of NOvA uncertainty parameters to validate the MCMC fitting machinery. To this end, the CCQE, CCRES, and FSI uncertainties are selected; random values are selected by drawing from a uniform

distribution ranging from $[-2\sigma, +2\sigma]$. From there, these shifts are applied to the nominal NOvA MC simulation in the $(Reco|\vec{q}|, E_{had}^{vis})$ 2D ND topological distributions to produce a fake data set, also known as an Asimov data set. No statistical fluctuations are applied; the random pulls, or true pull values can be seen in Fig. 6.2a. With these random pull values applied to the nominal NOvA MC simulation, we obtain a fake data set, or Asimov data set. The fake data, nominal simulation and uncertainties we seek to fit are input into MCMC to find the combination of pull values that successfully reproduce the fake data.

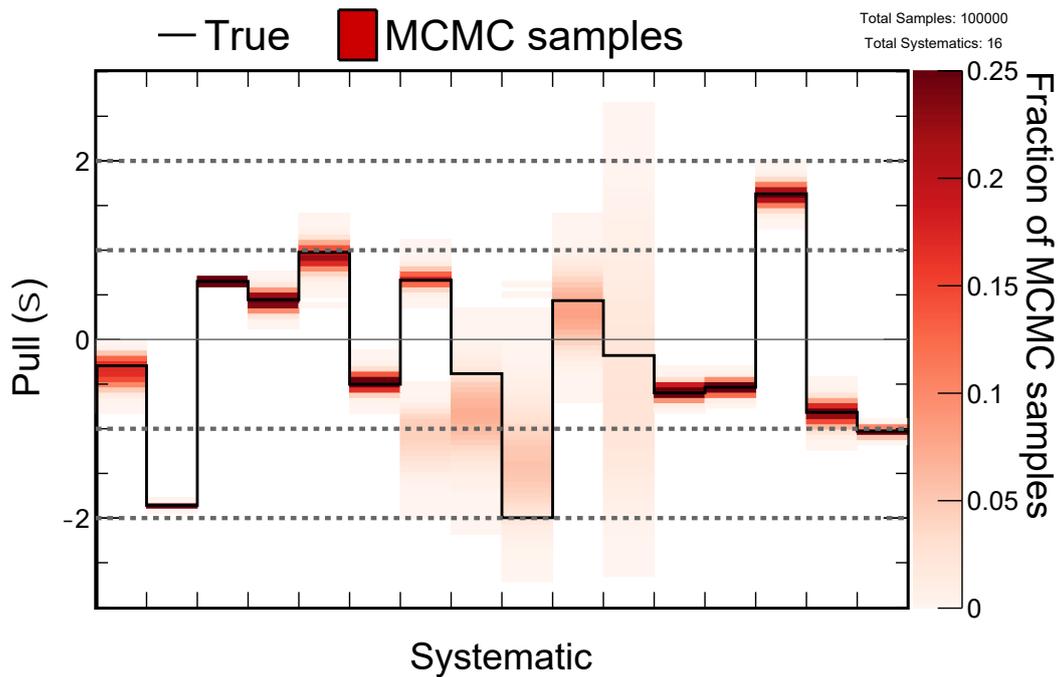
This fake data fit is performed using Fermilab’s computing center, where MCMC optimizes over approximately 250k degrees of freedom (1600 bins \times 10 topologies \times 16 uncertainties) to produce 100k samples in approximately 48 hours.

Pulls Summary

In an ideal fit, the pull values from the MCMC samples should match with the true values in Fig. 6.2a. One means of analyzing the agreement to the fake data is to plot the pull values for each parameter from all MCMC samples. We see these results in Fig. 6.2b where each true value uncertainty is a bin in the histogram and the MCMC samples are overlaid on top; in the lower figure plot, the darker red region indicates a larger density of samples and thus a more probable pull value. The CCRES and FSI parameters have MCMC samples that are very densely aligned at or near the true pull value, demonstrating these uncertainties are well constrained by the NOvA fake data. However, the CCQE parameters have a very broad distribution of MCMC samples indicated by the broad pink region that MCMC sampled many pull values from the posterior that are probable for these uncertainties, revealing these particular CCQE uncertainties (center bins of the plot) are not as well constrained by the fake data in the ND topological samples. Moreover, we see the CCQE systematic, `ZExpAxialFFSyst2020_EV1`, is best constrained by the fake data, while the last eigenvector, 4, is the least because of its broadest pink region indicating many pull values are probable.



(a) Distribution of the true pull for each uncertainty included in the fake data fit. Pulls are obtained from a uniform distribution ranging from $[-2\sigma, +2\sigma]$.



(b) Plot of the true pulls (Fig. 6.2a) with the MCMC samples overlaid in red.

Figure 6.2: All samples from the fit are overlaid onto the true pull distribution. The CC RES and FSI parameters are well constrained – the MCMC samples are densely populated around the true pull. The MCMC samples around the QE parameters are broadly distributed – these parameters are less constrained by the fake data.

Rep. Sample: 1D Projection onto E_{had}^{vis}

Next we can examine the 1D projections of E_{had}^{vis} following the fake data fit. The plots in Figures 6.3 and 6.4 show the nominal prediction in grey with the a priori $\pm 1\sigma$ error band from the 16 uncertainties used in the MCMC sampling. The black points are the fake data, and the blue distribution is the prediction using the pull values that correspond to the maximum \mathcal{LL} , the Rep. Sample. Recall the Representative Sample is simply a coordinate in the posterior parameter space, where each dimension in the space corresponds to the pull value for a parameter. Therefore, the Rep. Sample prediction is a prediction with the pulls on the systematics applied corresponding to the maximum \mathcal{LL} coordinate. Again, recall in Section 5.5 we acknowledge the Rep. Sample is not the ideal metric when using MCMC, as it does not account for where the *densest*, or highest probability, regions might be located within the posterior space. Regardless, it is useful here to quantitatively demonstrate the goodness-of-fit from MCMC sampling to the fake data set.

We see from these plots of E_{had}^{vis} that the Representative Sample successfully reproduces the fake data in every topology; another very convincing example that MCMC can constrain NOvA systematic uncertainties. One strong example to demonstrate the success of this fake data fit is the $\mu + \pi^\pm + X$ topologies for FHC and RHC. We see the fake data has distorted these two predictions significantly from the nominal prediction. In both plots the fake data is significantly less than the nominal prediction at the peak of $E_{had}^{vis} = 0.15$ GeV. Despite this significant difference between the nominal prediction and the fake data, MCMC is still able to perfectly reproduce the fake data when using the Rep. Sample prediction. Plots of the 1D projections onto the *Reco* $|\vec{q}|$ axis can be found in the Appendix (Figs. A.6 & A.7).

χ^2 metric

We can examine the accuracy of the fit quantitatively, again using the Representative Sample. Figure 6.5 shows the fake data-nominal MC χ^2 metric (grey)

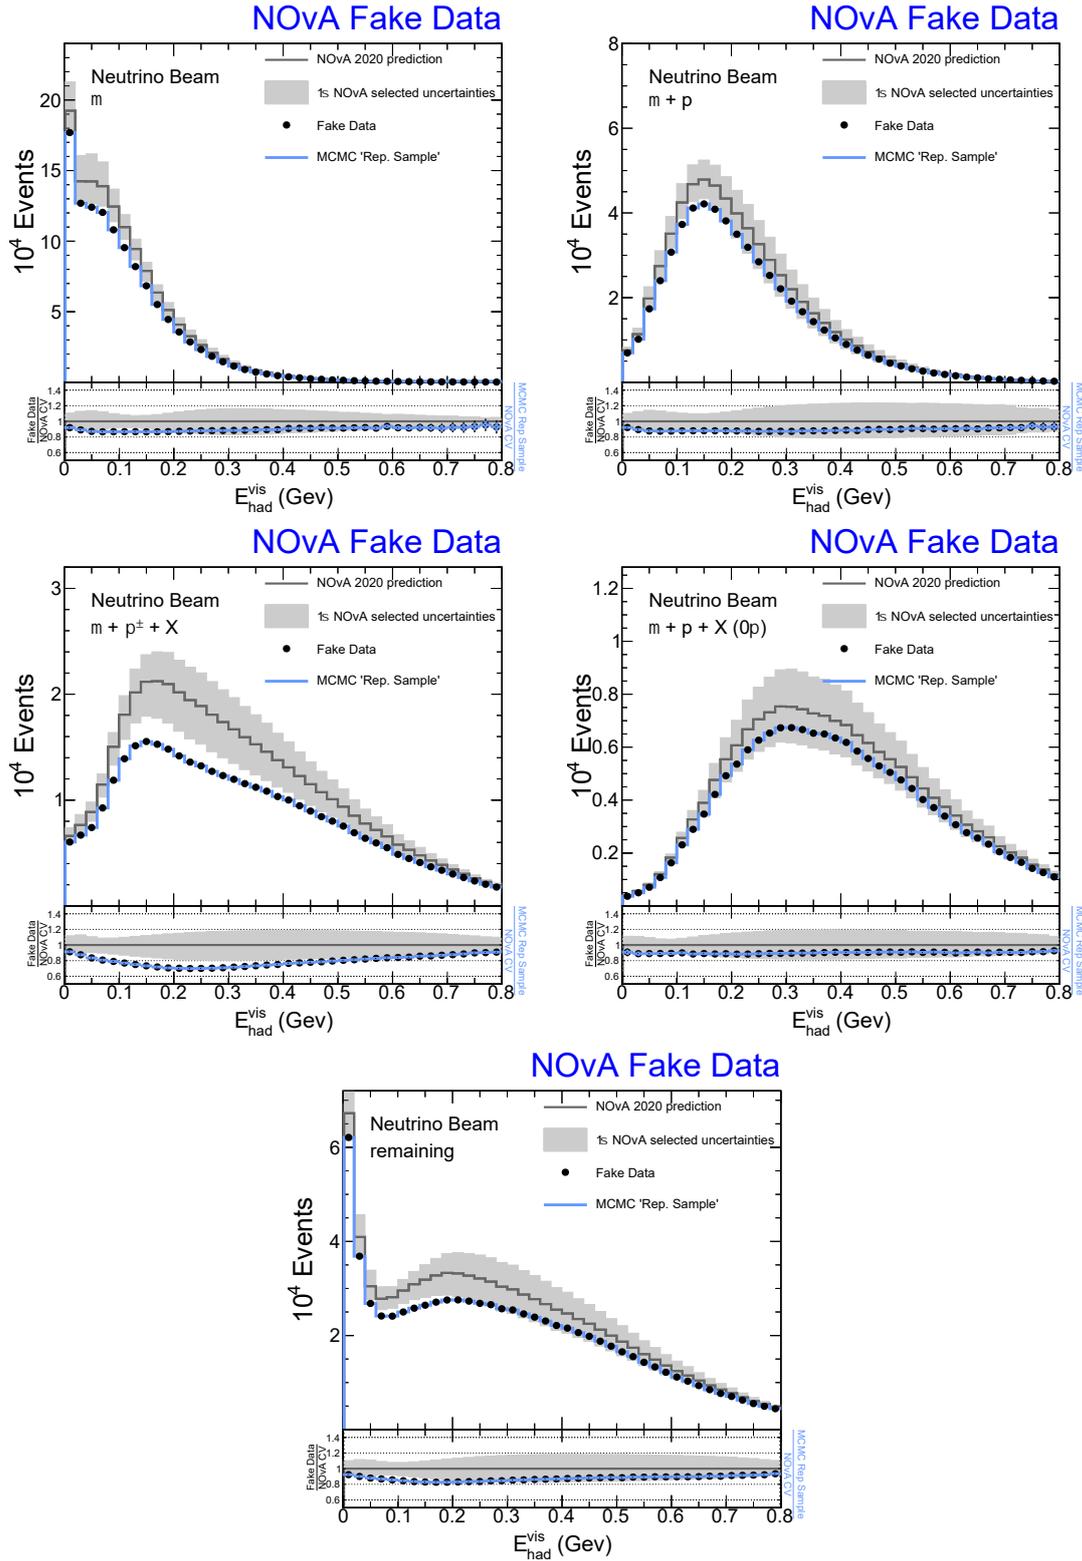


Figure 6.3: Representative Sample prediction, in blue, on E_{had}^{vis} projection of FHC topologies. The grey histogram represents the nominal prediction and the grey error band represents the a priori $\pm 1\sigma$ range for the 16 selected uncertainties. Note the Rep. Sample prediction (largest $\mathcal{L}\mathcal{L}$ MCMC sample) produces perfect agreement with the fake data, confirming MCMC can constrain NOvA uncertainties effectively. Plots of the 1D projections in the $Reco |\vec{q}|$ variable can be seen in the Appendix A.6.

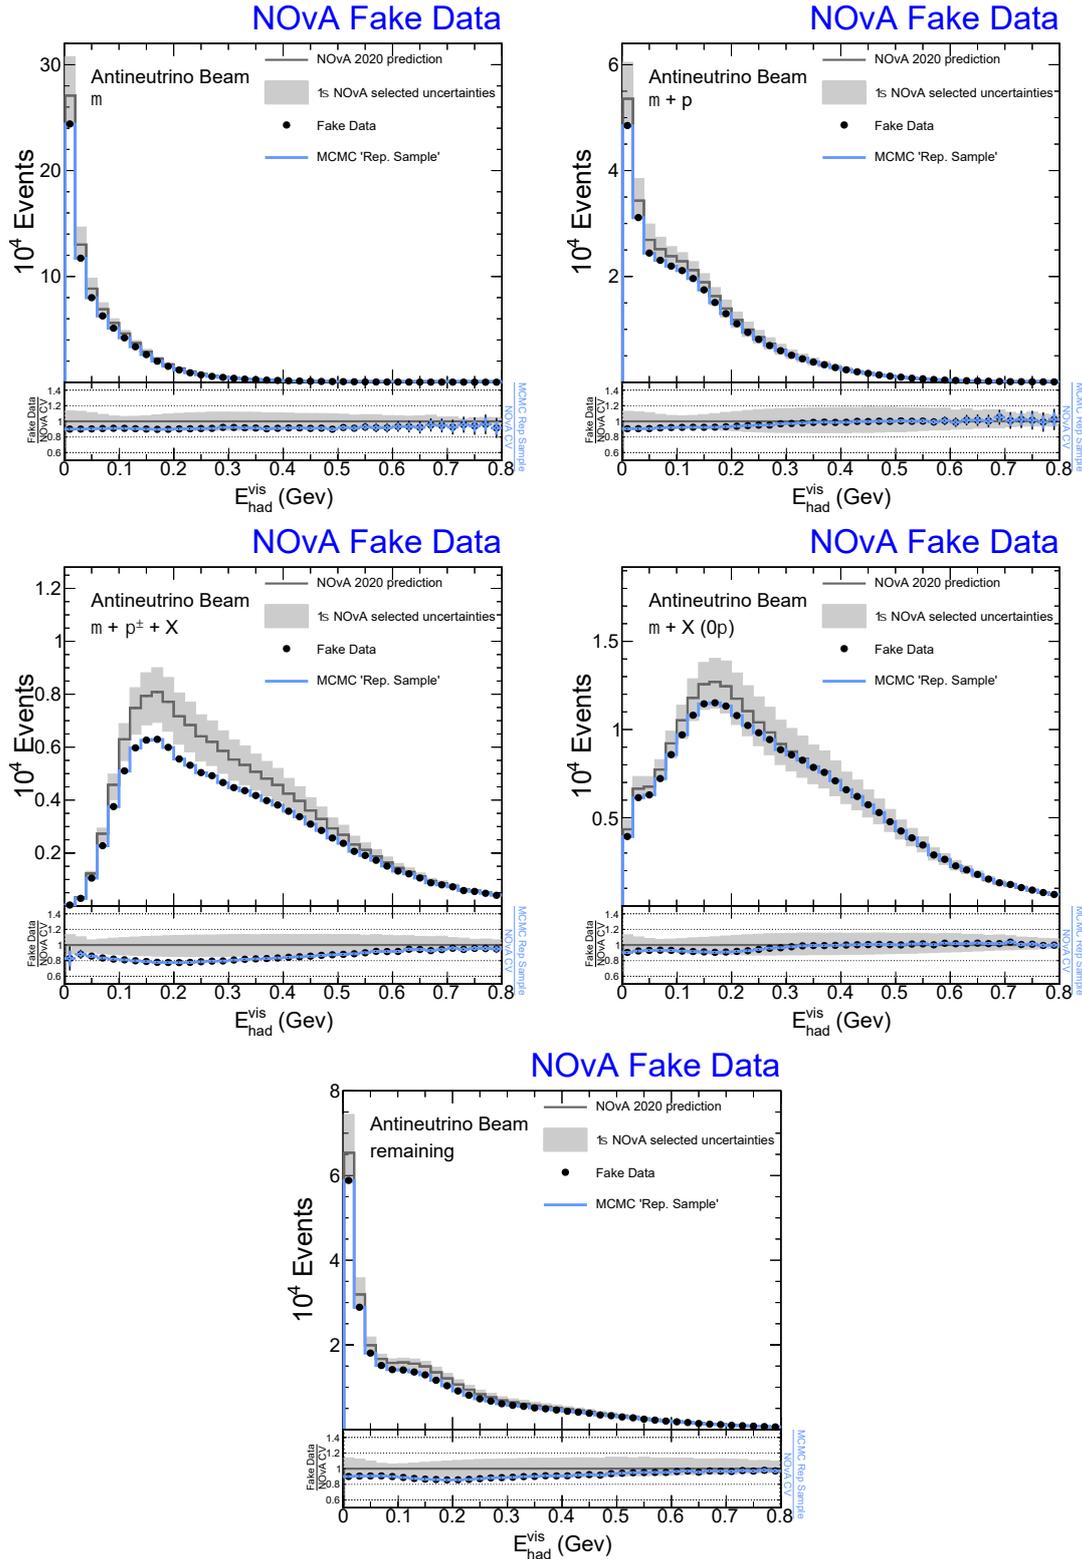


Figure 6.4: Representative Sample prediction, in blue, on E_{had}^{vis} projection of RHC topologies. The grey histogram represents the nominal prediction and the grey error band represents the a priori $\pm 1\sigma$ range for the 16 selected uncertainties. Note the Rep. Sample prediction (largest \mathcal{LL} MCMC sample) produces perfect agreement with the fake data, confirming MCMC can constrain NOvA uncertainties effectively. Plots of the 1D projections in the $Reco |\vec{q}|$ variable can be seen in the Appendix A.7.

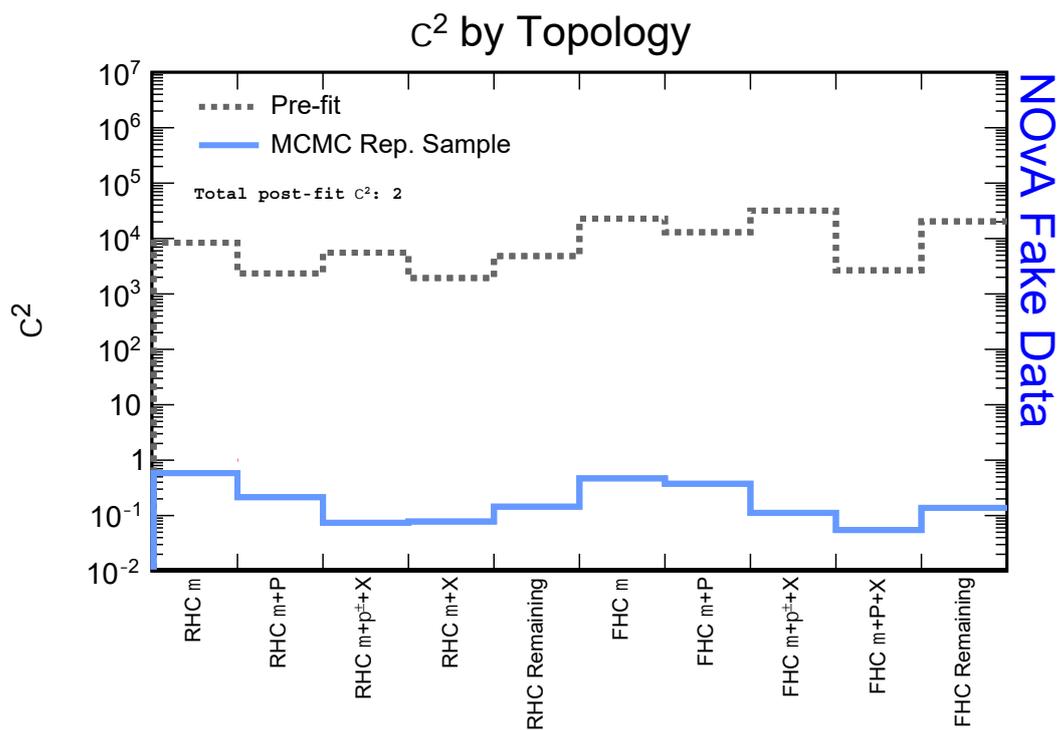


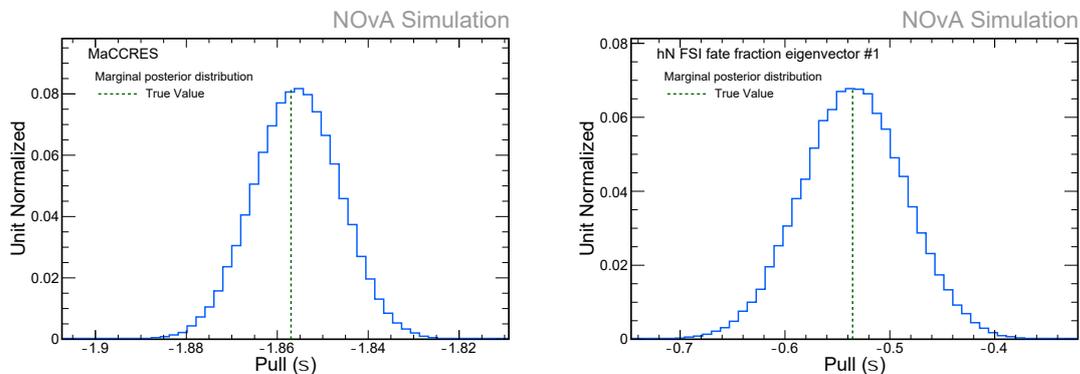
Figure 6.5: Plot of the fake data-MC χ^2 agreement as a function of the ND topology. The χ^2 is calculated using the Rep. Sample (maximum $\mathcal{L}\mathcal{L}$ MCMC sample) prediction. We see the χ^2 is poor before fitting (grey), however after MCMC sampling the agreement is near perfect with a cumulative $\chi^2 \approx 2$, quantitatively demonstrating MCMC can constrain NOvA uncertainties and fit the fake data.

and fake data-Rep. Sample MC χ^2 metric (blue) in each topological sample. Note the grey, “Pre-fit” χ^2 is large – intentionally so, this is before MCMC has sampled the posterior space. However, when applying the maximum \mathcal{LL} pull values to the nominal prediction, the blue distribution χ^2 drops dramatically. Moreover, when we consider the large number of events in the ND topological samples, a cumulative $\chi^2 \approx 2$ quantitatively demonstrates strong agreement between the fake data and Rep. Sample predictions.

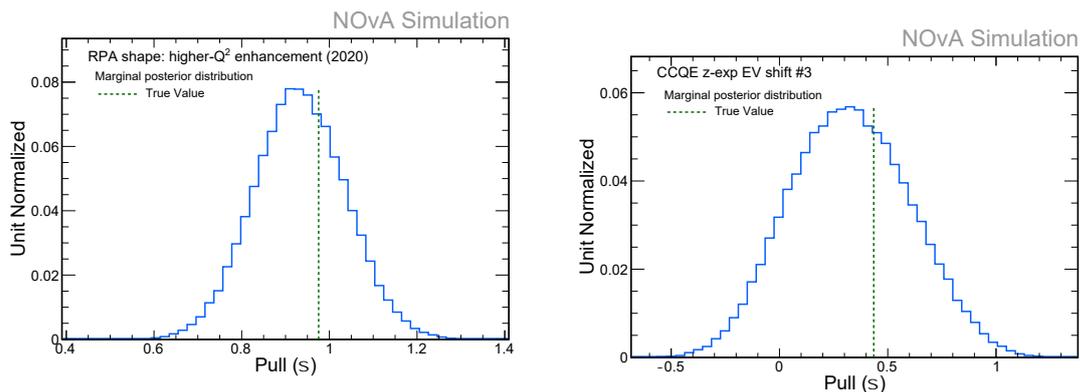
1D marginals

Another informative means of understanding the quality of the fit is to marginalize, or integrate out the probability of undesired, or nuisance, parameters as described in Subsec. 5.1.2. We implement this procedure here to produce a 1-dimensional distribution for each of the NOvA uncertainties provided in the fit. In other words, we project out the red MCMC samples distribution in Fig. 6.2b onto each uncertainty; Fig. 6.6 shows marginal distributions for four of the 16 parameters used in the fit. Focusing on Subfigs. 6.6a, 6.6c, and 6.6b, we see these are strongly constrained – the width of these marginal distributions is narrow indicating only a small range of values are the most probable pulls for these systematic uncertainties. We also observe the MCMC samples are accurate to the true pull; the most probable pulls, or modes, of these three distributions are centered along the green vertical line, which represents the true pull value that was applied to create the fake data (i.e. the black histogram in Fig. 6.2). These plots confirm the most probable pull value from MCMC matches exactly the true pull value. This is strong evidence the MCMC machinery is functioning properly.

However, the CCQE Z-expansion systematic is not as well constrained, Subfig. 6.6d. Although the green vertical line representing the true pull value is near the most probable MCMC pull value, the width of the distribution is fairly wide ranging from negative to positive pull values. This wide range of probable pull values reveal this fake data set provided into the fit cannot constrain this uncertainty as effectively



(a) 1D marginal distribution of the NOvA uncertainty on the axial mass for CC RES neutrino interaction, M_A . (b) 1D marginal distribution of the NOvA uncertainty of the first eigenvector of the FSI hN model.



(c) 1D marginal distribution of the NOvA uncertainty on the Random Phase Approximation (RPA) shape enhancement for QE interactions. (d) 1D marginal distribution of the NOvA uncertainty of the third eigenvector for the QE model Z-expansion.

Figure 6.6: Four 1D marginal distributions of NOvA uncertainties in this fake data fit. These distributions are probability densities for a given pull value. Note the vertical green line represents the true pull value from the random uniform distribution and applied to the fake data. This green line is centered on the peak, or mode, of each marginal distribution confirming MCMC has successfully reproduced the random pull values by sampling the posterior. These marginals are normalized to unit area and can be interpreted as the posterior probability distribution.

as the other three parameters in Fig. 6.6. This provides useful information: not every NOvA uncertainty can be constrained from the fake data set we provide. We will return to this concept and build on this further when fitting NOvA data with the complete set of NOvA uncertainties.

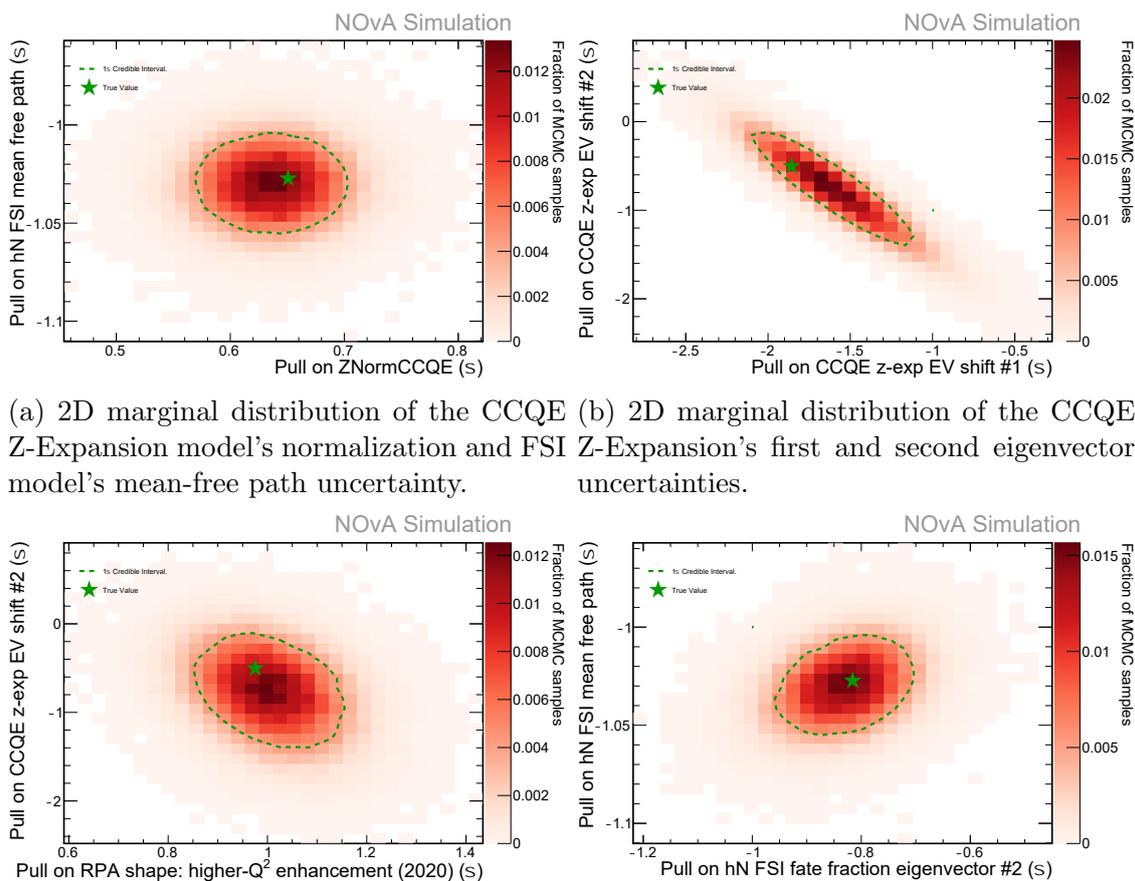
2D marginals

Another means of dissecting MCMC samples is to learn of the correlations of any pair of systematic uncertainties by again marginalizing over the remaining uncertainties. Looking at 2D marginal distributions is most helpful for exactly this use case – to learn of correlations (or anti-correlations) between pairs of systematic uncertainties. From the same fake data fit, we show this in Fig. 6.7.

In the first Subfig. 6.7a, the correlation between the two parameters, the CCQE Z-Expansion model’s normalization and FSI model’s mean free path is fairly weak; the distribution of samples is equally dense in all directions from the true pull value represented by the green star. Given these two uncertainties, this relationship is to be expected, changes to the CCQE model normalization have little bearing on the effect of the mean free path of pions as they traverse the nucleus.

In the next Subfig. 6.7b of the CCQE model’s Z-Expansion first and second eigenvectors, we find the two uncertainties are highly anti-correlated; as one these two uncertainty values changes, the other will change in the opposite direction. This is expected, as the these two uncertainties eigenvectors are attempting to adjust the MC in the same region of phase space of $(Reco |\vec{q}|, E_{had}^{vis})$.

The last two Subfigs. 6.7c and 6.7d show only moderate anti-correlation and correlation, respectively. The general shape appears sloped downwards and upwards, respectively, indicating there is some correlation between these parameters. Considering the physics corresponding to these two pairs of uncertainties, we should expect some level of correlation; Subfig. 6.7c both address the CCQE model and therefore should influence one another, while Subfig. 6.7d both address the hN FSI model.



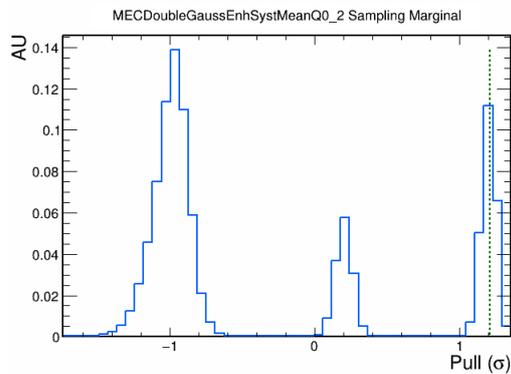
(a) 2D marginal distribution of the CCQE Z-Expansion model's normalization and FSI Z-Expansion's first and second eigenvector model's mean-free path uncertainty. (b) 2D marginal distribution of the CCQE Z-Expansion's first and second eigenvector uncertainties.

(c) 2D marginal distribution of the CCQE RPA shape enhancement and the second Z-Expansion eigenvector uncertainty. (d) 2D marginal distribution of the hN FSI model's second eigenvector and the mean-free path uncertainty.

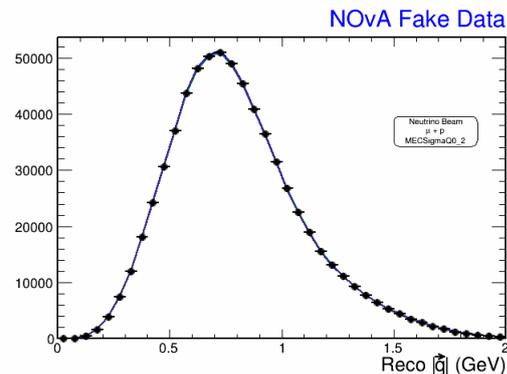
Figure 6.7: 2D marginalizations of a subset of the parameters utilized in the fake data fit. These distributions indicate correlations between uncertainties. From these plots, we see a breadth of the strength of correlations between uncertainties. The green star is the true pull value for the given pair of uncertainties. The green line represents the Bayesian 1σ credible interval. Marginalizations like these are one of the powerful advantages of MCMC sampling.

6.2.2 Degeneracies in MEC Double Gaussian Parameters

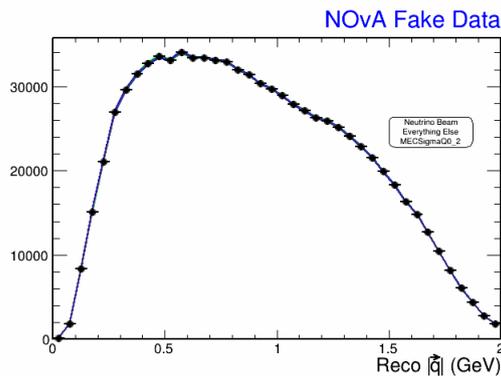
Now that an example of a ND fake data fit has been analyzed, there is a particular fake data fit result of interest. The fake data fit we will focus on is one that includes the MEC Double Gaussian parameters. When looking into the 1D marginalizations of the parameters used in this fake data fit, we observe several degeneracies – or multiple, high-probability modes – in several of the marginals of the MEC Double Gaussian parameters. In Figure 6.8, we highlight the degeneracies present within one of NOvA's MEC Double Gaussian parameters. Figure 6.8a shows the 1D marginal of



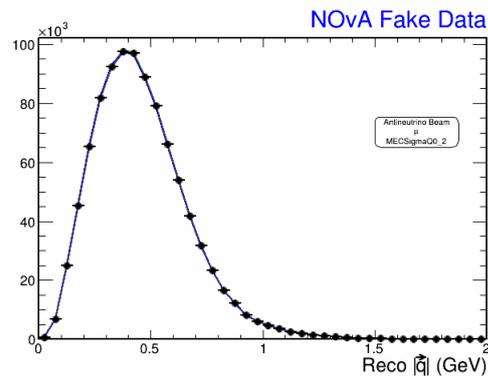
(a) The multi-modal 1D marginal distribution of the MEC Double Gaussian parameter `MECDoubleGaussEnhSystMeanQ0_2` certainty. The green line at $+1.2\sigma$ denotes the true pull value.



(b) 1D projection of $\mu + P$ FHC topology onto $Reco |\vec{q}|$ with all three modes applied.



(c) 1D projection of ν Else FHC topology onto $Reco |\vec{q}|$ with all three modes applied.



(d) 1D projection of μ RHC topology onto $Reco |\vec{q}|$ with all three modes applied.

Figure 6.8: **Top Left:** Plot of the three probable pull values for the MEC Double Gaussian parameter `MECDoubleGaussEnhSystMeanQ0_2`. There are three unique, probable modes, at -1σ (blue), $+0.2\sigma$ (red), and $+1.2\sigma$ (green, the true pull value) indicating a degeneracy. **Remaining Figures:** Plots of the 1D projections onto $Reco |\vec{q}|$ with the predictions from all three pull values applied. All three pull values of this parameter produce perfect agreement to the fake data.

the lower MEC Gaussian’s Mean q_0 systematic, `MECDoubleGaussEnhSystMeanQ0_2`. The marginal of this systematic shows three distinct, clear high-probability modes centered at: -1σ , $+0.2\sigma$, and $+1.2\sigma$; the next question raised is: do the each of these three unique pull values produce the same quality of agreement? To this end, the remaining three plots, Figs. 6.8b, 6.8c, and 6.8d answer this question. These three plots contain predictions, in $Reco |\vec{q}|$, of all three of these modes -1σ (blue), $+0.2\sigma$ (red), and $+1.2\sigma$ (green). Regardless of any of the three pull values applied to the MC, we obtain the *same prediction, indicating a degeneracy in this systematic*, because all three unique pull values produce a strong agreement with the fake data (black dots). We note we observe this behavior in additional MEC Double Gaussian systematics: `MECDoubleGaussEnhSystMeanQ3_2`, `MECDoubleGaussEnhSystCorr_2`, `MECDoubleGaussEnhSystNorm_2`, and `MECDoubleGaussEnhSystSigmaQ0_2`. This behavior of degeneracies from the MEC Double Gaussian parameters is an important discovery in the context of the ND fitting, and will be important as we move onto NOvA data. What’s more, it presents an opportunity to improve the MEC model by eliminating these degeneracies.

6.3 Data & MC Comparisons

Before we transition from MCMC sampling with fake data to sampling with real NOvA data, it is important to first understand where the NOvA data is positioned with respect to the NOvA Prod5.1 MC. This question is crucial as it helps to more meaningfully understand where MCMC is attempting to correct the prediction when sampling the posterior space of NOvA uncertainties.

Figures 6.9 and 6.10 show the 1D E_{had}^{vis} projections of the NOvA ND topologies in FHC and RHC for both Prod5 data and MC *and* Prod5.1 data and MC ($Reco |\vec{q}|$ projections in Appendix A.8 & A.9). First, we must highlight the improvement in the agreement between data-MC from NOvA’s recent production, Prod5, to its current one, Prod5.1. This is most noticeable in the lowest bin of the μ topologies. This improvement in agreement is even more significant as the number of events

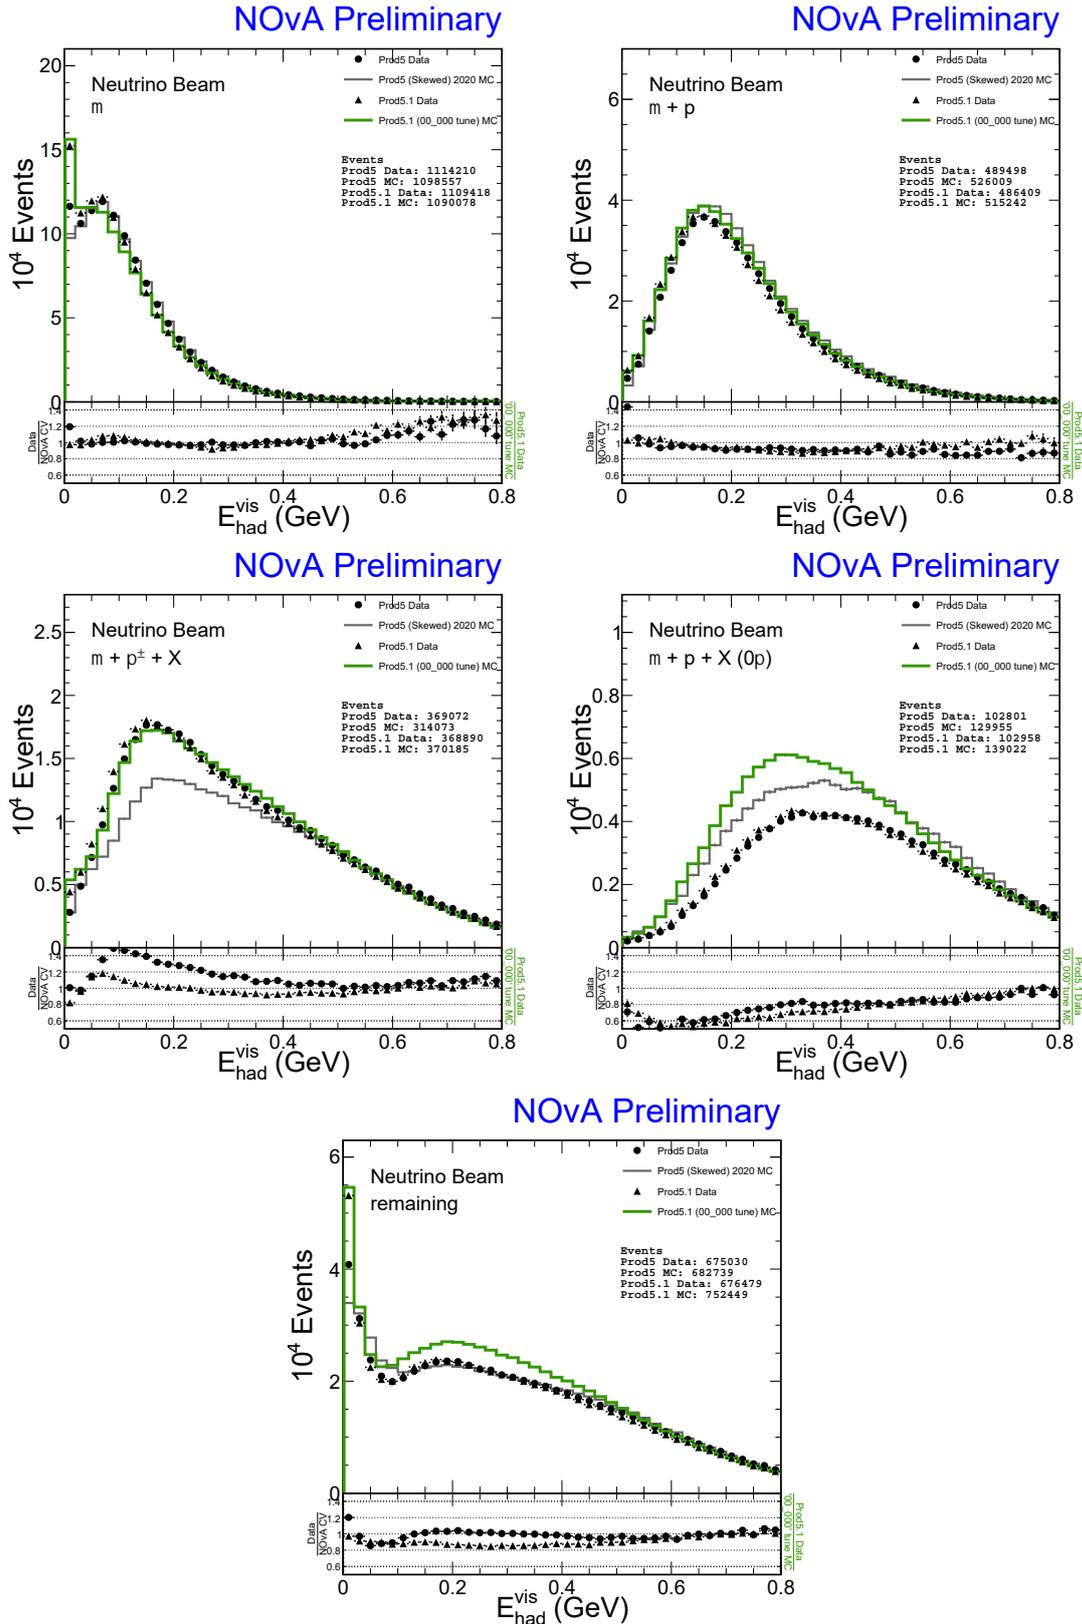


Figure 6.9: Comparisons of Prod5 data (black dots) and Prod5 MC (grey) and Prod5.1 data (black triangles) and Prod5.1 MC (green) predictions projected onto E_{had}^{vis} in the FHC topological samples. This plot highlights the changes in the underlying physics models from Prod5 to Prod5.1. The focus here is to show the disagreement between the Prod5.1 data and Prod5.1 MC, as this is what is used in the analysis. Plots of the 1D projections in the $Reco |\vec{q}|$ variable can be seen in the Appendix A.8.

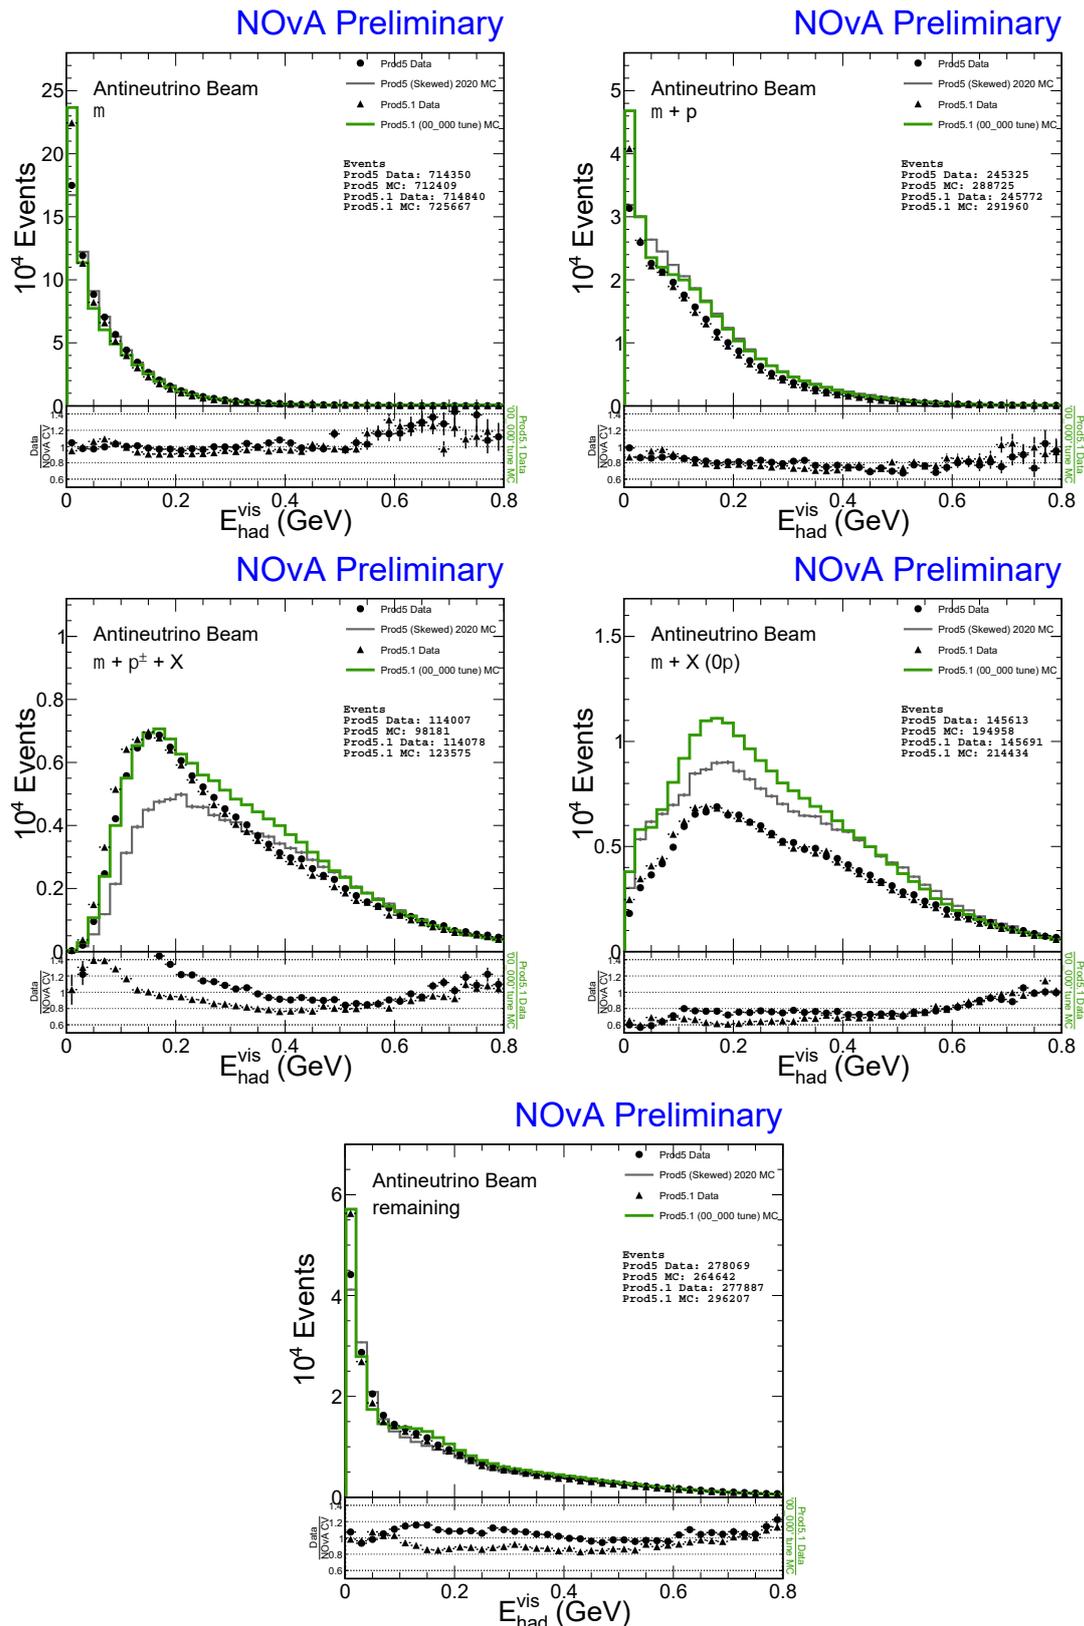


Figure 6.10: Comparisons of Prod5 data (black dots) and Prod5 MC (grey) and Prod5.1 data (black triangles) and Prod5.1 MC (green) predictions projected onto E_{had}^{vis} in the RHC topological samples. This plot highlights the changes in the underlying physics models from Prod5 to Prod5.1. The focus here is to show the disagreement between the Prod5.1 data and Prod5.1 MC, as this is what is used in the analysis. Plots of the 1D projections in the $Reco |\vec{q}|$ variable can be seen in the Appendix A.9.

in this topology is larger than all other topologies, therefore this will drive the χ^2 metric significantly. Moreover, despite an improvement of the data-MC agreement, because these low E_{had}^{vis} bins in the μ topologies have so many events, the statistical uncertainty is sub-percent level. This means even if the overall agreement is at the few-percent level bin-by-bin, the statistics will drive the χ^2 contribution significantly. This region of phase space is dominated by QE and MEC, an important reminder when we consider the pulls on the systematics from these models when fitting the ND data.

We observe a noticeable improvement to the $\mu + \pi^\pm + X$ topology from Prod5. The nominal Prod5.1 MC predicts this $\mu + \pi^\pm + X$ topology well. This means there is little for the ND fit to correct directly in this topology.

We also plot the data-MC ratios of the ten topologies in Figure 6.11 & 6.12. This provides us more insight into where the disagreement lies, especially because it is in the 2D distributions that the MCMC sampling is performed in. We observe an over-prediction, in blue, in the $\mu + P + X$, $\mu + X$, and **EvElse** in Prod5.1 MC. The disagreement we see in the $\mu + P + X$ topology is one that we believe we have sufficient systematic freedom to bring the prediction into agreement with the Prod5.1 data. Regardless, these are large regions of phase space where MCMC will need to improve the agreement.

Lastly, we see the over-prediction also lives along the low- Q^2 diagonal kinematic boundary. This is very informative and helpful to understand what type interactions may be causing this over-prediction; we must be observant of systematics that are associated low- Q^2 events when fitting the ND data.

To summarize, there are several qualitative features we want to focus on:

- Moderate agreement, but very high statistics in the low ($Reco |\vec{q}|, E_{had}^{vis}$) region of phase space, especially in the two μ topologies (see Figure 6.9 & 6.10) that will drive the χ^2 .
- A general over-prediction of events in several topologies: FHC $\mu + P + X$, RHC $\mu + X$, RHC $\mu + P$, and **EvElse** (see Figure 6.9 & 6.10).

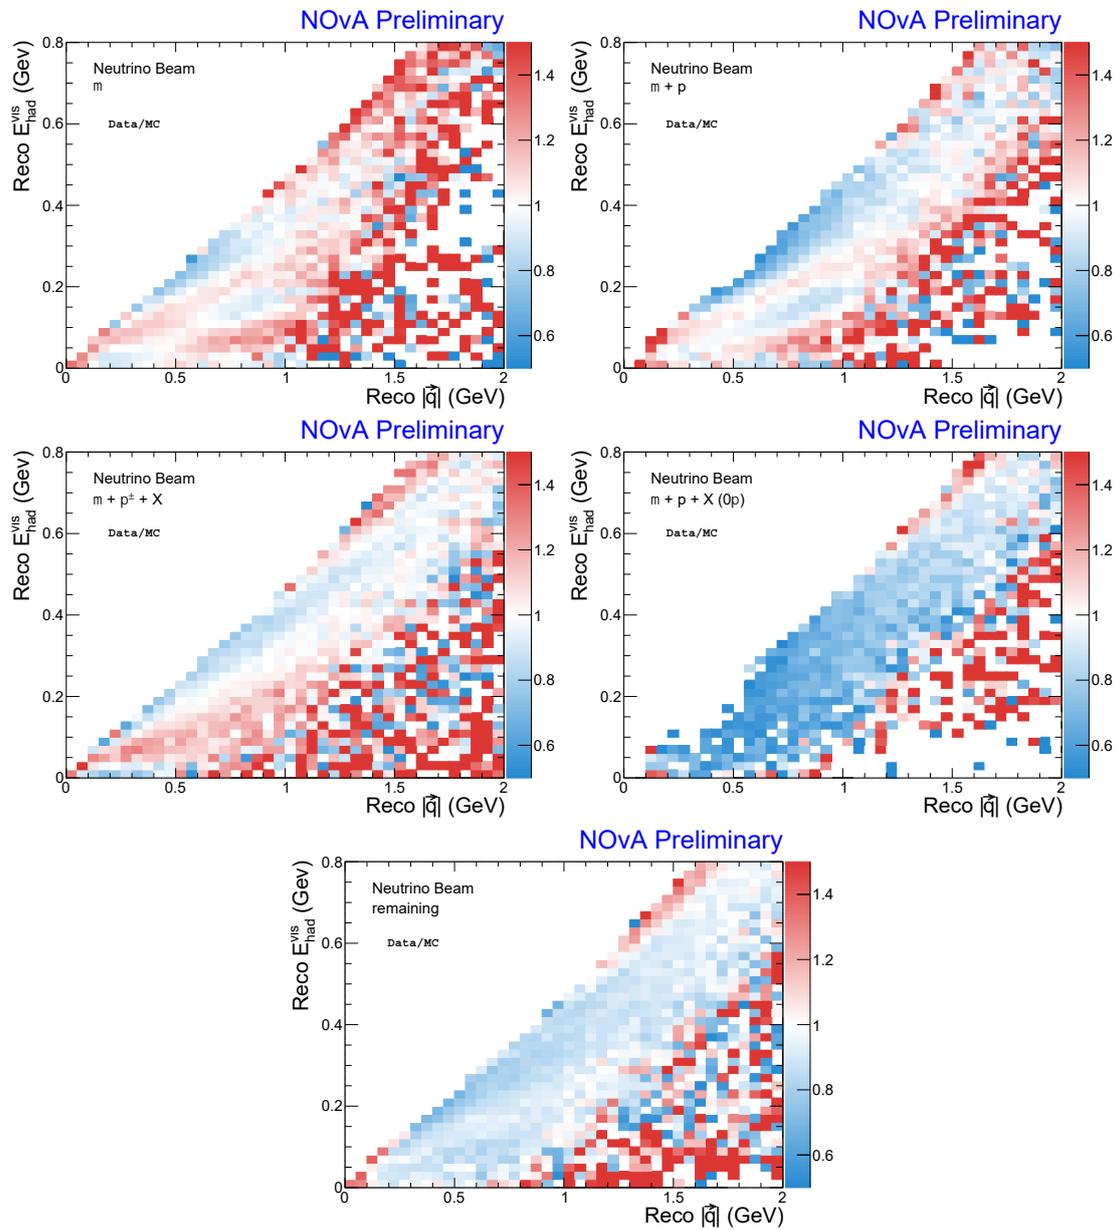


Figure 6.11: FHC 2D topological samples plotted as the data-MC ratio. Note the over-predicted (blue) region along the kinematic boundary. This boundary is along the $Q^2 \approx 0 \text{ GeV}^2$ diagonal.

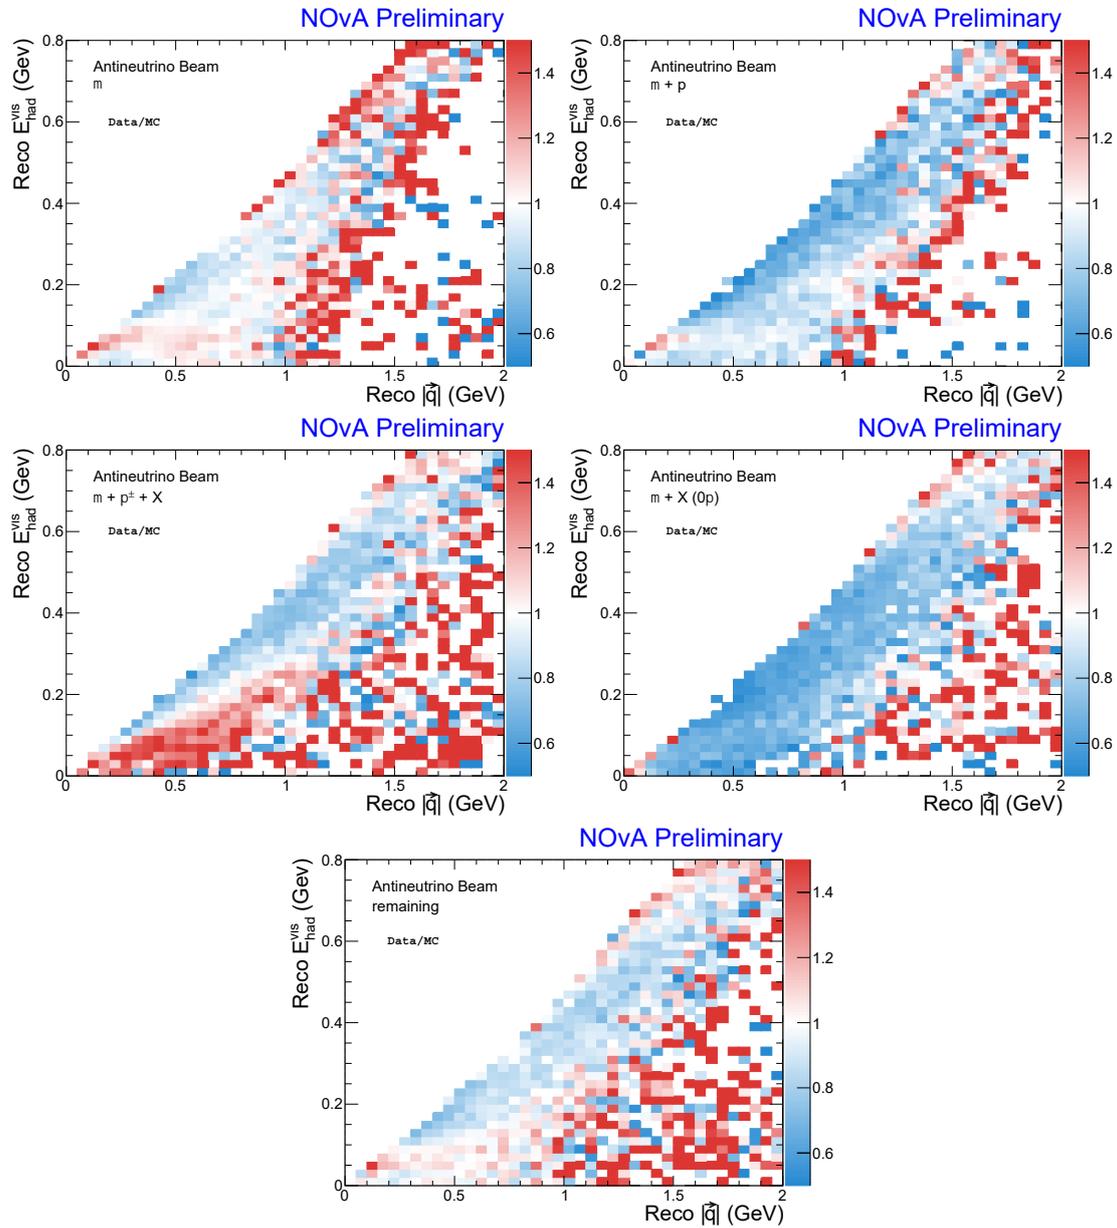


Figure 6.12: RHC 2D topological samples plotted as the data-MC ratio. Note the over-predicted (blue) region along the kinematic boundary. This boundary is along the $Q^2 \approx 0 \text{ GeV}^2$ diagonal.

- An over-prediction of events in the low- Q^2 region of phase space (see Figures 6.11 & 6.12).

Understanding where the disagreement between data and MC is within the topologies *and* the physical reasoning behind it is extremely critical when drawing meaningful conclusions from MCMC sampling results, as we must have a sense of what and where in the phase space needs improvement in these ND topologies. With this understanding, we can advance to fitting the NOvA data.

6.4 Initial Fitting Attempts to ND Data

In a preliminary fit to NOvA data, the results were unsuccessful. In fact, the Representative Sample produced worse agreement with the χ^2 metric in certain topologies, seen in Figure 6.13. Figure 6.14 shows that many systematics have very broad 1D marginal distributions, indicating they were not constrained by the data. This motivated the need to adjust the information we provide into MCMC.

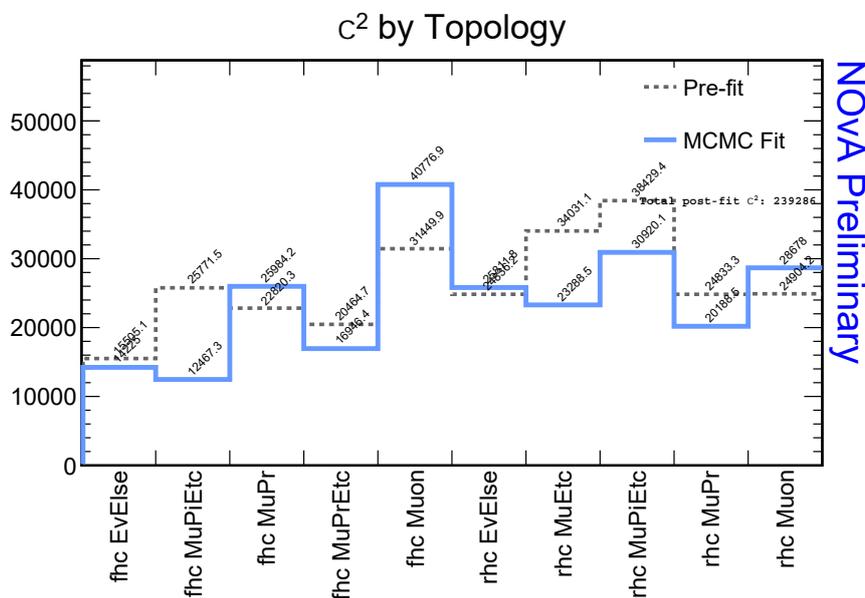


Figure 6.13: The χ^2 , Eq. 5.18, per topology before fit (grey) and after fitting, using the “Representative Sample” prediction (blue). In most topologies the agreement between the data and NOvA simulation improves, however in the Muon topologies the agreement is markedly worse.

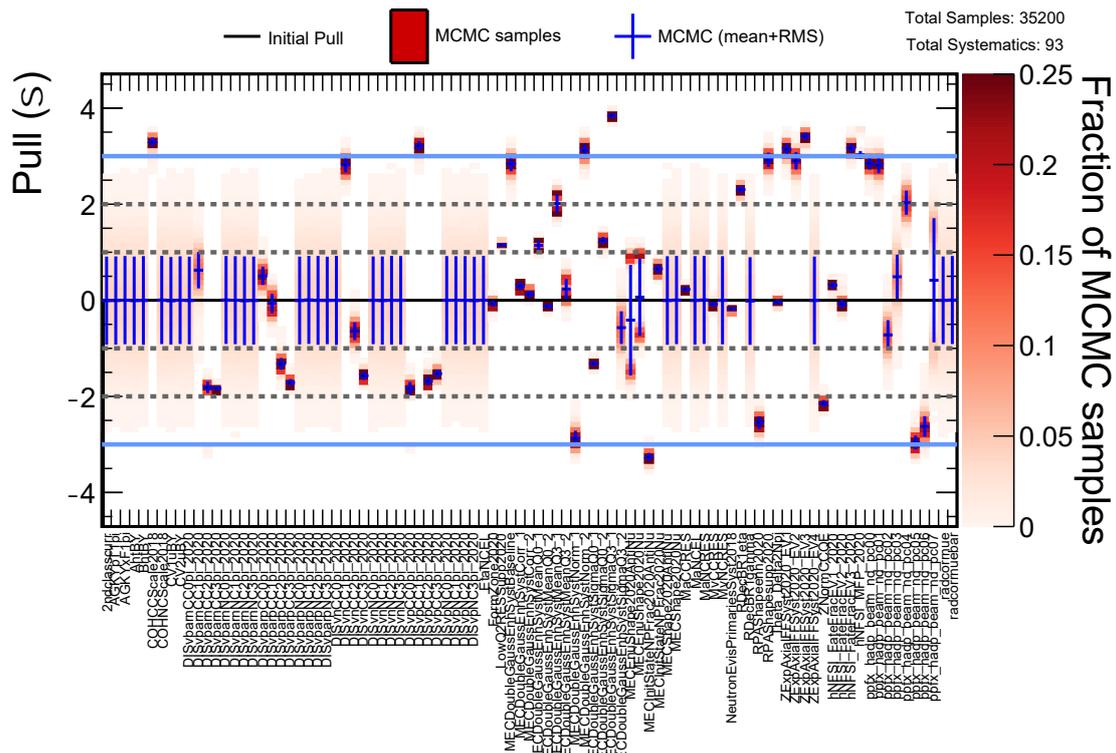


Figure 6.14: There are many uncertainties (columns) that have many probable pull values, or broad pink distributions. Effectively, MCMC is recreating our prior function. This means the NOvA ND data does not constrain these uncertainties. This motivates the need to change how we perform MCMC sampling.

This poor level of agreement prompts the need to further improve the fit. The approach to achieve the goal of improving the MCMC fit is divided into two categories: **quality** and **quantity**. We seek to improve the **quality** of the fit by improving the χ^2 between the data and the prediction. To improve the **quantity** we aim to increase the number of MCMC samples that are produced, or, in other words, reduce the run-time of the MCMC sampling.

To achieve these goals, we must first understand the primary limitations preventing and an improvement of quality and quantity of an MCMC fit:

- The computation demand of Hamiltonian Monte Carlo algorithm is large – approximately 8 GB of CPU memory is required to run on the FermiGrid. There were too few MCMC samples obtained from the exploration to obtain meaningful conclusions.
- Not every NOvA systematic is constrained by the ND ν_μ data.

The computational demand of MCMC to explore a 93 dimensional parameter posterior space is immense. Given the constraints of the FermiGrid, this task is impractical and we must reduce the demand to increase the **quantity** of MCMC samples.

The ND ν_μ topologies are complex and each one is sensitive to different models in different ways. Therefore further scrutiny of the underlying models is necessary: are there additional degrees of freedom in the underlying models that can be exposed and utilized in the fit? Are there important physics elements that have been overlooked? Have some model uncertainties been underestimated (which would lead to too-tight priors)? These are important considerations to improve the **quality** of the fit.

6.4.1 Adjustments to MCMC Input

This section outlines the changes made to the predictions and parameters that are input into Markov Chain Monte Carlo. The objective in performing these adjustments is to improve the **quality** and **quantity** of the fit.

Ranking uncertainties

From Fig. 6.14, it is clear there are numerous uncertainties that are not constrained by the NOvA ND ν_μ data, evidenced by broad, pink regions in columns of the Pulls Summary plot. This motivates the question if it is necessary to include every NOvA systematic uncertainty into the MCMC fit. To address this question, we shift each systematic uncertainty by $\pm 1\sigma$ with respect to the nominal prediction in each topology, one systematic at a time. To quantify the change in the shifted distribution from the nominal distribution, we calculate a χ^2 using Eq. 5.18. With these χ^2 values for each systematic, we can rank which value is the largest, and therefore which systematic shifts the NOvA ND topological distributions the most. The systematics that have the largest values will be uncertainties that *are* constrained by the NOvA data, while those that have small values are not well constrained by the data. The plots of these χ^2 values for each systematic in each ND topological

sample can be seen in Figures 6.15-6.19 (RHC figures can be found in the Appendix A.2.3). We plot the χ^2 values on a log scale and note there are approximately 15-20 uncertainties that alter the nominal the prediction substantially. This means these systematics are likely the most important during MCMC sampling. Conversely, there are about 10-15 systematics that alter the ND topologies very little, by their small χ^2 values on the right side of the x -axis on these plots. To summarize which systematics *are* significant in changing the prediction, a summary is provided below:

- Flux systematics.
- MEC Double Gaussian systematics.
- QE systematics.
- Neutron Response systematic.
- Resonant systematics.
- DIS $1,2\pi$ systematics.
- FSI systematics.
- Detector systematics.

The NC DIS and COH scattering systematics are examples that have low χ^2 values on these plots and are not as relevant in changing the NOvA prediction, which is expected as our selection captures ν_μ CC interactions.

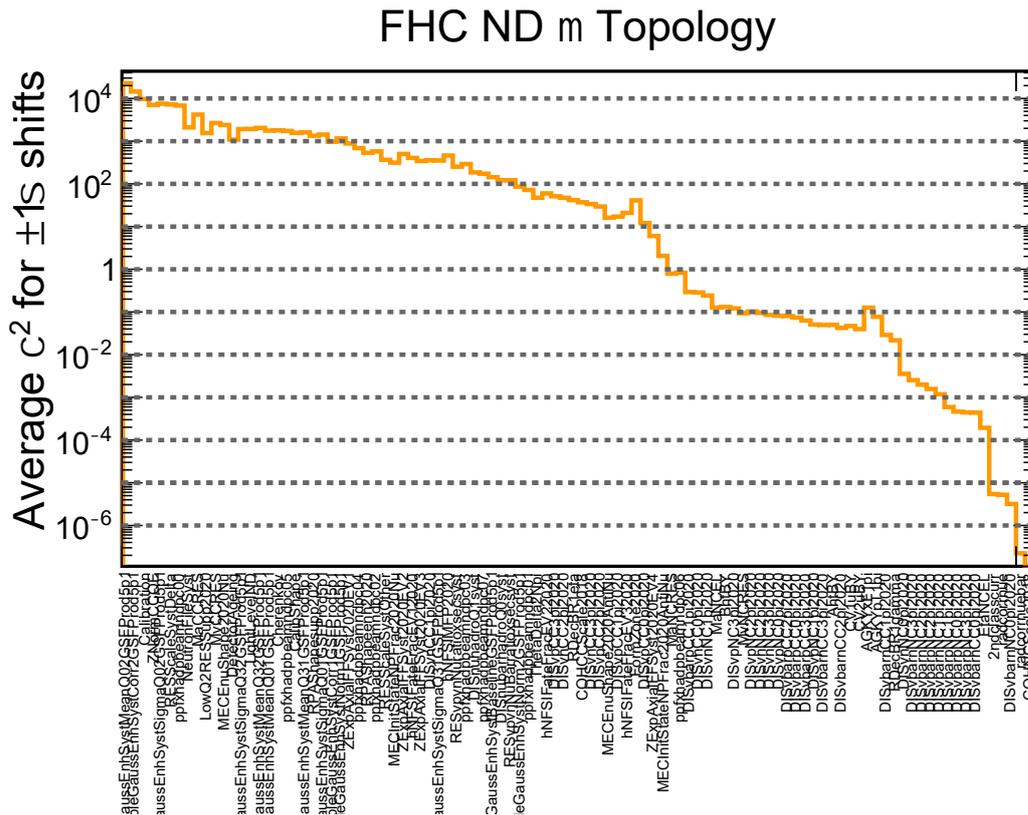


Figure 6.15: ND FHC μ topology χ^2 systematic ranking between the nominal prediction and the $+1\sigma$ shifted prediction. Note the log-scale y -axis, which illustrates the range of χ^2 values for NOvA’s systematic uncertainties. There are 10-15 systematics that alter the ND prediction substantially, and 15-20 systematics that alter the ND predictions very little.

These five plots (and the five RHC plots) contain much information. To distill this information down, we average the χ^2 together from each systematic over all ten (FHC and RHC) topologies. The result is a single plot of the χ^2 for each systematic, Fig. 6.20. Our objective is to reduce the dimensionality of the posterior space by eliminating parameters that do not change the prediction significantly. These parameters with smaller χ^2 values are likely to also change the posterior probability very little from one pull value to another. Thus, due to the fact the predictions don’t change meaningfully, from a topological perspective, the posterior space in these dimensions are likely very flat. This means the HMC NUTS criteria is likely not met and a MCMC chain can run on almost indefinitely, drastically slowing the down the speed of the MCMC sampling, and producing fewer samples. If we remove these parameters from MCMC sampling, the posterior space is less likely to

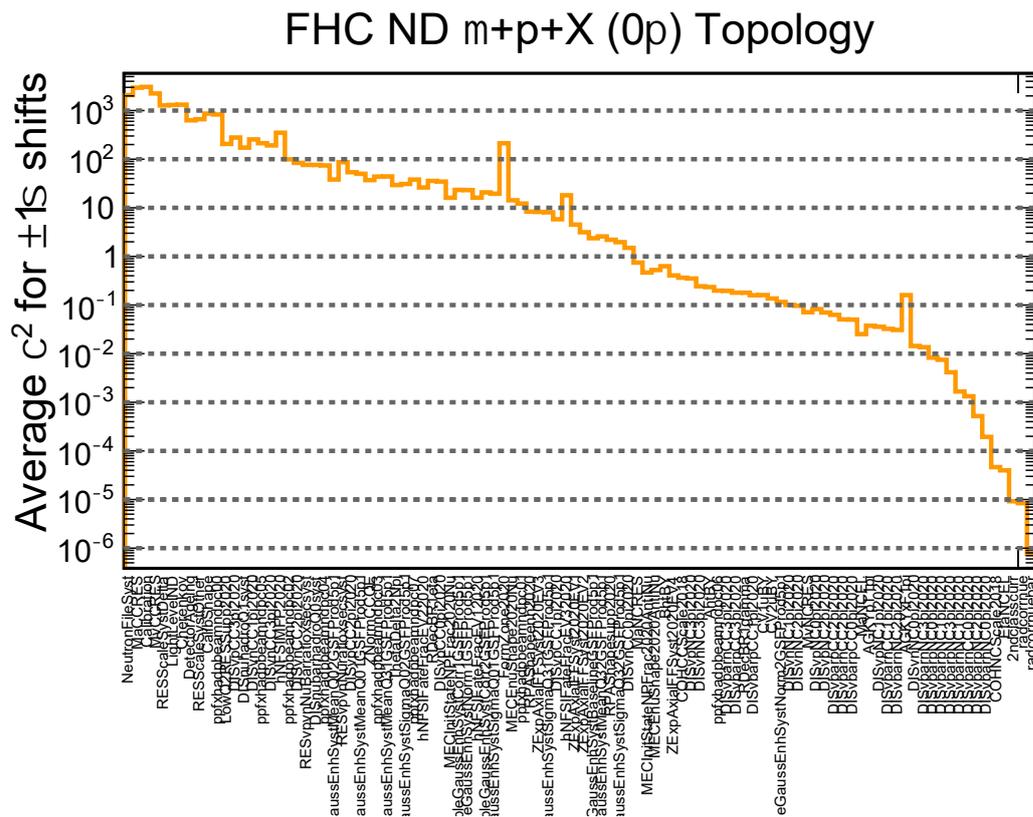


Figure 6.18: ND FHC $\mu + P + X$ topology χ^2 systematic ranking between the nominal prediction and the $+1\sigma$ shifted prediction. Note the log-scale y -axis, which illustrates the range of χ^2 values for NOvA's systematic uncertainties. There are 10-15 systematics that alter the ND prediction substantially, and 15-20 systematics that alter the ND predictions very little.

be flat and should allow MCMC chains to explore the posterior more effectively and trigger the NUTS criteria. With this χ^2 information, we estimate the systematics that meaningfully alter the ND predictions possess a $\chi^2 > 100.0$; we select a cut of $\chi^2 = 100.0$ to use for ND fits; any systematic uncertainty with a $\chi^2 < 100.0$ in this plot is removed from the ND MCMC sampling. The result of this cut is notable: *a reduction of approximately 30 systematic parameters that are not well constrained by the NOvA ND ν_μ data*, reducing the dimensionality of the posterior space, and thereby reducing the CPU demand, thus increasing the **quantity** of MCMC samples.

Variable binning & phase space cut

Introducing a cut on the $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space and altering the binning scheme are two developments that significantly reduce the computational demand of these MCMC fits, and thereby increase the number of statistics returned from a fit, with an approximately 30% reduction in the CPU wall time.

The 2D distributions are binned in 40×40 uniformly sized bins. Observing that events are kinematically prohibited beyond the $Q^2 = 0$ diagonal motivates the change to variable binning, where we can use larger bins widths for bins with few or no events. As has been shown, the majority of the neutrino interactions occur at low E_{had}^{vis} within the NOvA detectors, for both FHC and RHC beam. Therefore, when reducing the number of bins, we must be careful to maintain this resolution at low E_{had}^{vis} as this is where we would like to achieve optimal data-MC agreement. This means we also we aggregate events into a few larger bins in the regions of high $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space, which are statistically limited.

To further reduce the sampling time, we introduce a cut on the $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space. This cut is applied to all ten topologies and intended to remove events in the statistically limited region of phase space at low- E_{had}^{vis} and high- $Reco |\vec{q}|$. These cuts are diagonal lines, parallel to the $Q^2 = 0$ diagonal. Figure 6.21 and Figure 6.22 show the result of these two developments on the ten topological samples.

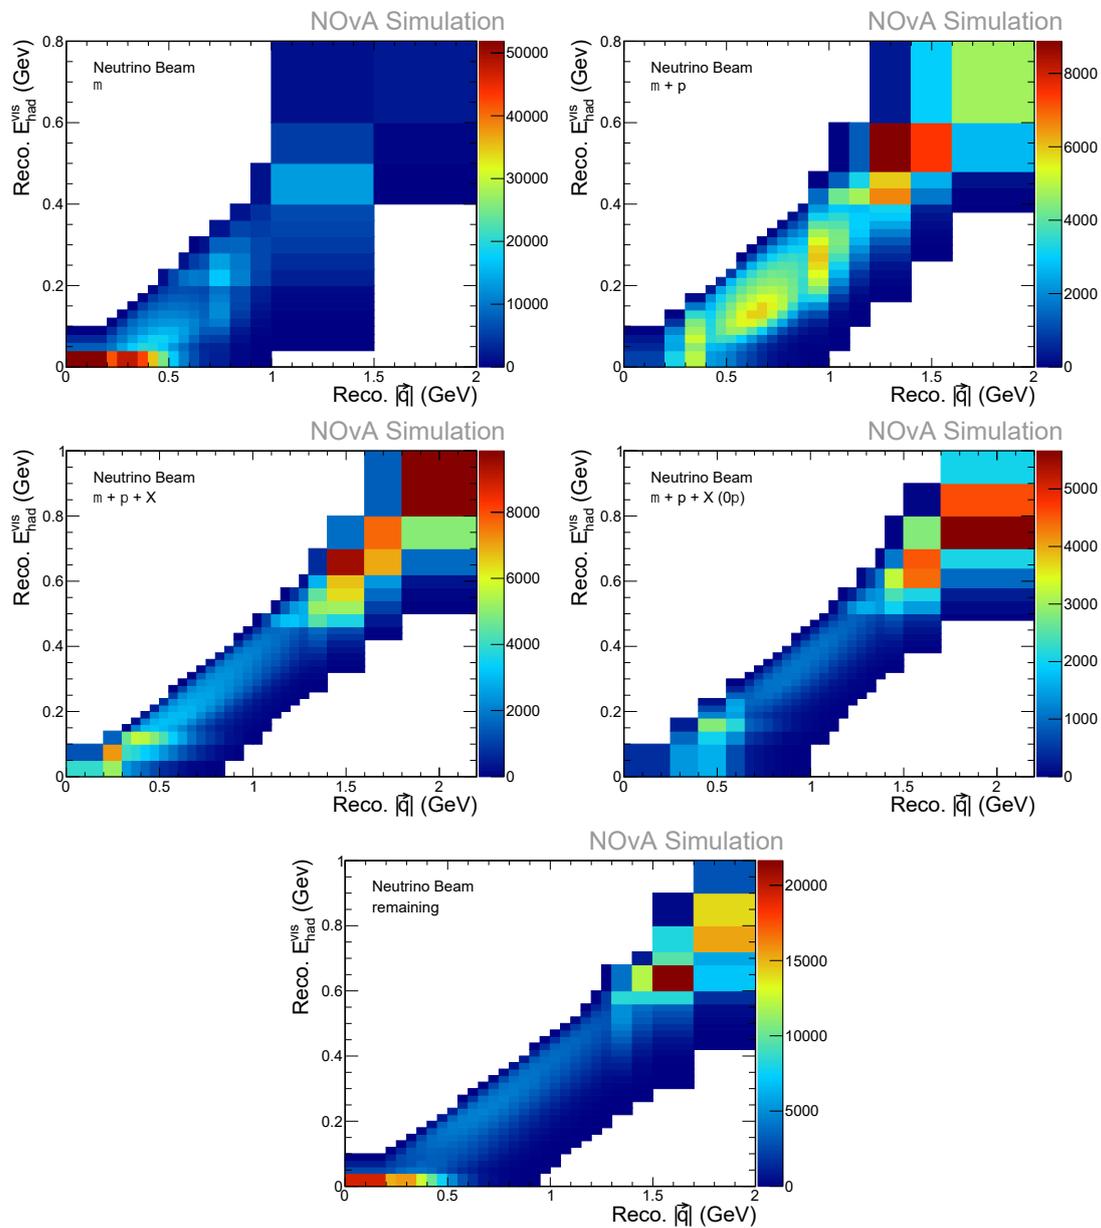


Figure 6.21: The five FHC ND topologies after applying the variable binning scheme and the phase space cut to remove events at low- E_{had}^{vis} , high- $Reco\ |\vec{q}|$. The fine binning at low- E_{had}^{vis} is preserved and allows MCMC to continue to fit this complex region of phase space.

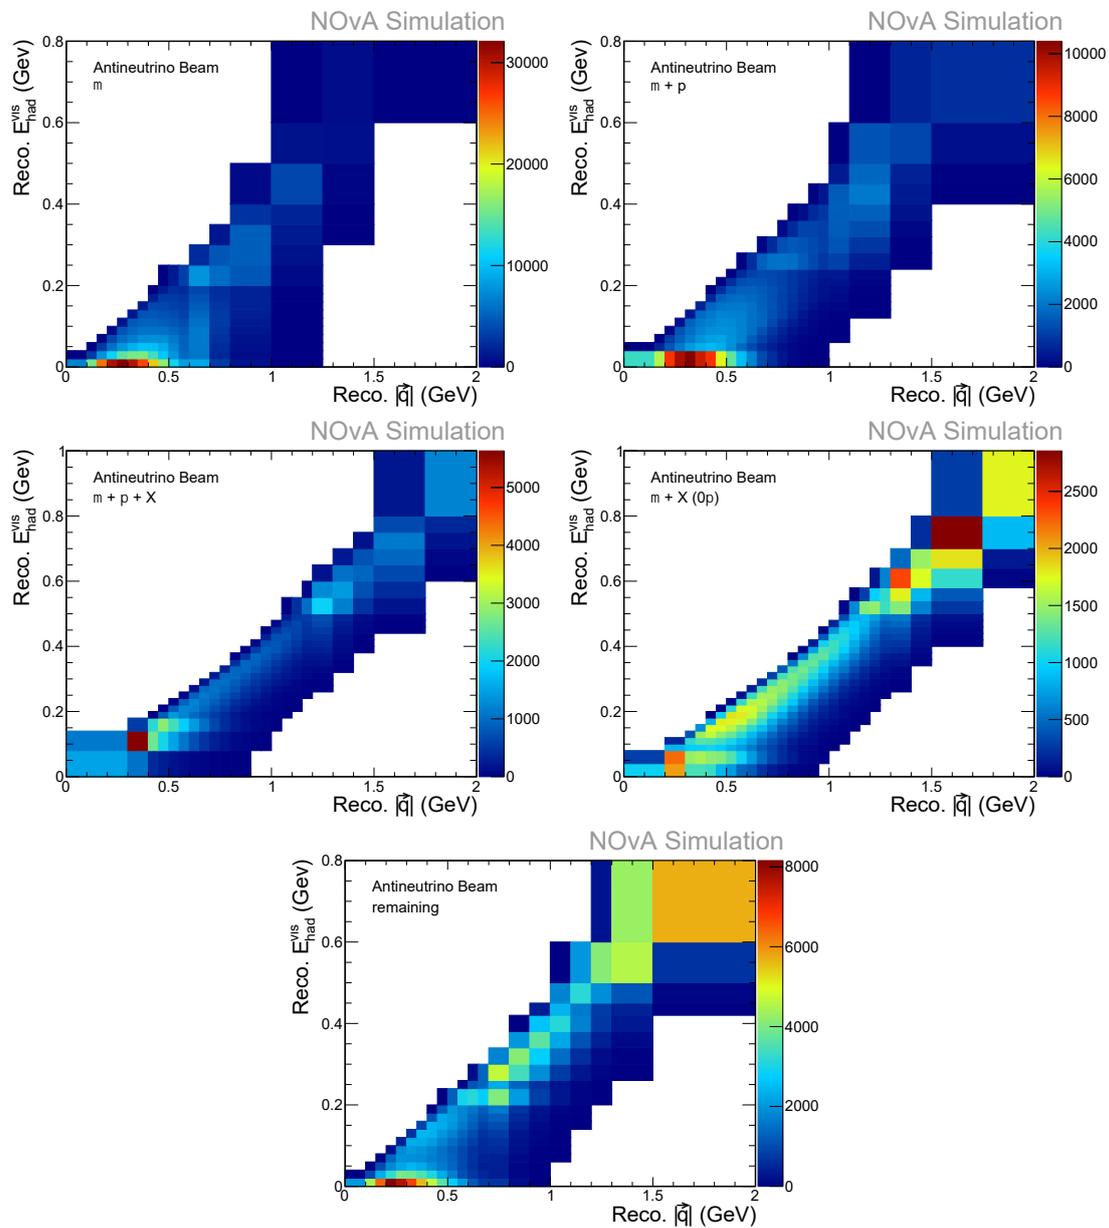


Figure 6.22: The five RHC ND topologies after applying the variable binning scheme and the phase space cut to remove events at low- E_{had}^{vis} , high- $Reco |\vec{q}|$. The fine binning at low- E_{had}^{vis} is preserved and allows MCMC to continue to fit this complex region of phase space.

It is particularly useful to remove the events in these low- E_{had}^{vis} and high- $Reco |\vec{q}|$ bins because systematically shifted predictions produce large statistical fluctuations when compared to the nominal prediction; this would cause MCMC to attempt to correct the prediction in these individual bins rather than where the regions of phase space where most events live (along the $Q^2 = 0$ diagonal). The result from these two developments will both improve the **quality** and the **quantity** of the MCMC sampling by reducing bins with statistical fluctuations and reducing the number of bins to calculate the target distribution in, respectively. These are the final adjustments to the ND predictions we will use for both the ND fit and the joint ND + FD fit.

Covariance matrix input

To further reduce the MCMC sampling time, a covariance matrix of the uncertainties can be input into a MCMC to replace the warmup phase of MCMC exploration of the posterior thereby saving a significant amount of CPU time. This covariance matrix would replace the mass term from Eq. 5.14 with a dense matrix that quantifies the correlations between each parameter. We calculate a covariance matrix of the parameters from a previous fit to data and use this matrix as input into a new MCMC sampling to the NOvA data. However, this option did not yield the desired results as the sampling phase of MCMC posterior exploration was slowed drastically. There are several reasons this could have occurred, for example when providing MCMC with a covariance (and replacing the warmup phase) the user must provide parameters to determine the “step size” the chain will take in the posterior space. This is done automatically in the warmup phase, and therefore further tuning may be needed. However, the source of the problem was not identified, and so was abandoned to maintain the warmup phase of MCMC.

6.5 Newly Created Cross-Section Uncertainties

In the work of the ND fit, we have identified the need for additional degrees of freedom – specifically in RES interactions – to improve the agreement between the data and the MC. For example, we already know GENIE predicts many RES events at low- Q^2 from Figure 4.6. If we look at the RES prediction for the $\mu + \pi^\pm + X$ topologies (Fig. 6.23, top two plots), we see much of the prediction in this topology are RES events, especially at low- Q^2 . In the middle plots of Fig. 6.23, we see the DIS contribution for the same topology; here the peak of the distribution is higher up, also associated with higher W , along the kinematic boundary. The RES and DIS interactions are the two most dominant interactions in this topology. Next, the bottom two plots in Fig. 6.23 display the motivation to create new systematic degrees of freedom. The bottom two plots show the data-MC χ^2 contribution per bin; bins that are dark red indicate where the disagreement between the data and MC is strongest. Note how the largest dark red region in the two bottom plots are along the kinematic boundary at $Q^2 \approx 0$ GeV², confirming this region is poorly modeled with NOvA’s model configuration of GENIE.

To address this region of disagreement, we have created three new sets of new parameters that will be incorporated into our fits:

- **RES relative cross section scaling** (RES_vpvn_Nu{Bar}XSecRatio).
- **RES Δ and Other Resonance scaling** (kRESScaleSyst{Delta,Other}).
- **SIS/DIS π production scaling** (kDISNuHadroQ1Syst & kDISNuBarHadroQ0Syst).

6.5.1 RES Relative Cross Section Scaling Systematics

In a RES neutrino interaction, a ν_μ can strike a proton or a neutron and create a Δ particle via a resonance. The significance is the resulting secondary particles decaying from the Δ will differ; in the case of $\nu_\mu + n \rightarrow \mu^- + \Delta^+$, the Δ^+ will decay into different secondary particles than a neutrino striking a proton: $\nu_\mu + p \rightarrow \mu^- + \Delta^{++}$. The difference between decay particles from the Δ is crucial

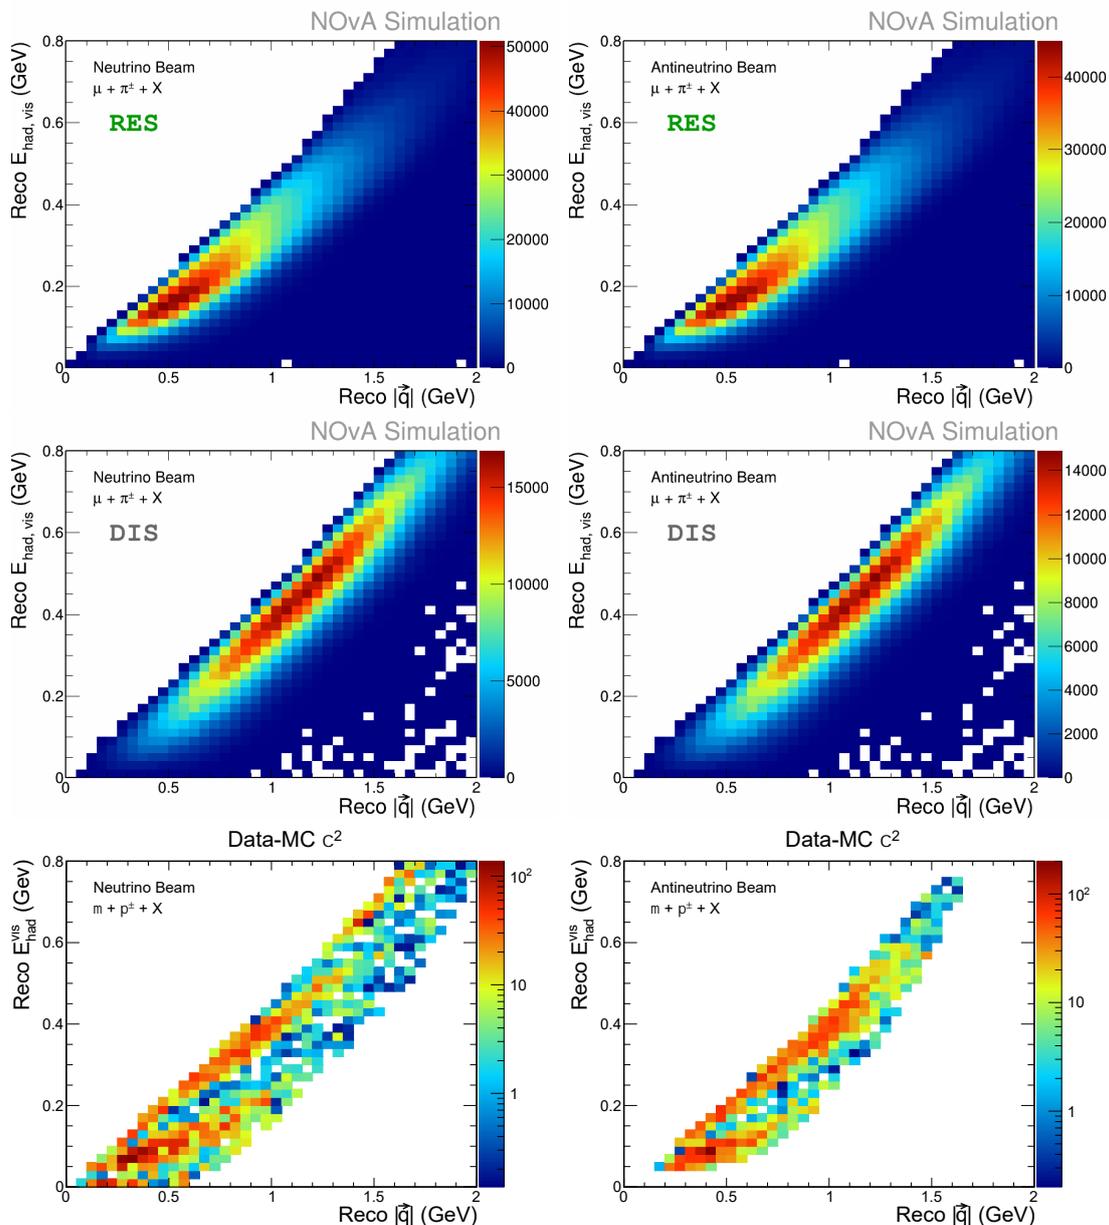


Figure 6.23: Plots of the $\mu + \pi + X$ topology in FHC and RHC. **Top:** 2D distributions of true RES interactions. **Middle:** 2D distributions of true DIS interactions. **Bottom:** 2D distributions of the data-MC χ^2 per bin. For the RES and DIS truth plots, the majority of events live in the middle of the distributions, at approximately 300 MeV of E_{had}^{vis} and 750 MeV of $Reco |\vec{q}|$, where the χ^2 contribution is largest. This region of phase space is notably along the $Q^2 \approx 0$ GeV² diagonal (see to Fig. 4.6). Systematic uncertainties are motivated to attempt to address this discrepancy.

because they are observed in the ND, and therefore determine which ND topology the interaction would appear. For example, if a neutrino strikes a *proton* and creates a Δ^{++} resonance, the final decay products will be $p + \pi^+$ and likely be observed in the $\mu + \pi^\pm + X$ topology. Meanwhile, if a neutrino strikes a *neutron*, the Δ^+ resonance will decay into $\pi^0 + p$ or $\pi^+ + n$ which would be observed in the $\mu + P / \mu + P + X$ and $\mu + \pi^\pm + X$ topologies, respectively, illustrating the struck nucleon is significant in RES interactions. The complete table of the Δ decays and the likely ND topologies the interaction would appear in can be found in Table 6.1.

ν Interaction	Δ Decay	Relevant ND Topologies
$\nu + p \rightarrow \mu^- + \Delta^{++}$	$\Delta^{++} \rightarrow p + \pi^+$	$\mu + \pi^\pm + X$
$\nu + n \rightarrow \mu^- + \Delta^+$	$\Delta^+ \rightarrow \pi^+ + n$ $\Delta^+ \rightarrow \pi^0 + p$	$\mu + \pi^\pm + X$ & EvElse $\mu + P$ & $\mu + P + X$
$\bar{\nu} + p \rightarrow \mu^+ + \Delta^0$	$\Delta^0 \rightarrow \pi^0 + n$ $\Delta^0 \rightarrow \pi^- + p$	$\mu + X$ $\mu + \pi^\pm + X$
$\bar{\nu} + n \rightarrow \mu^+ + \Delta^-$	$\Delta^- \rightarrow n + \pi^-$	$\mu + \pi^\pm + X$

Table 6.1: Breakdown of the common Δ RES neutrino interactions for FHC and RHC. The type of Δ resonance will determine which ND topology the neutrino interaction will be placed into. It is important to note that not every π , for example, will be successfully identified, thereby impacting the classification of events in the ND topologies.

We discuss how determining the resonance decay in a neutrino interaction can dictate which ND topology the neutrino interaction corresponds to. This is the motivation for our first pair of systematics: `RES_vpvn_Nu{Bar}XSecRatio`. With these, we can modulate how events migrate from one ND topology into another by controlling the ratio of the total RES $\frac{\nu+p}{\nu+n}$ cross section. The systematic predicts $+N\sigma$ shifts by increasing the numerator $N \times 5\%$, and decreasing the denominator $N \times 5\%$. Meanwhile, negative sigma shifts decrease the numerator and increase the denominator, making the weight smaller and reducing the total RES cross section. We see a summary of how extreme shifts of these systematics impact the topologies in Table 6.2.

For completeness, we show the $\pm 1\sigma$ shifts of these two systematics on the $\mu + \pi^\pm + X$ topology projected onto E_{had}^{vis} for FHC and RHC in Fig. 6.24. At the maximum

Systematic	Shift	Impact on ND Topologies
RES_vpvn_NuXSecRatioSyst	+3 σ	$\uparrow \mu + \pi^\pm + X, \downarrow \mu + P, \mu + P + X$
	-3 σ	$\downarrow \mu + \pi^\pm + X, \uparrow \mu + P, \mu + P + X$
RES_vpvn_NuBarXSecRatioSyst	+3 σ	$\downarrow \mu + X, \uparrow$ and $\downarrow \mu + \pi^\pm + X$
	-3 σ	$\uparrow \mu + X, \downarrow$ and $\uparrow \mu + \pi^\pm + X$

Table 6.2: Summary of the impact for each of the two systematics at the extreme ($\pm 3\sigma$) pulls and their impact on the ND topologies. It is crucial to note that applying these systematics will not directly map RES events from one topology into another due to imperfect reconstruction and particle identification.

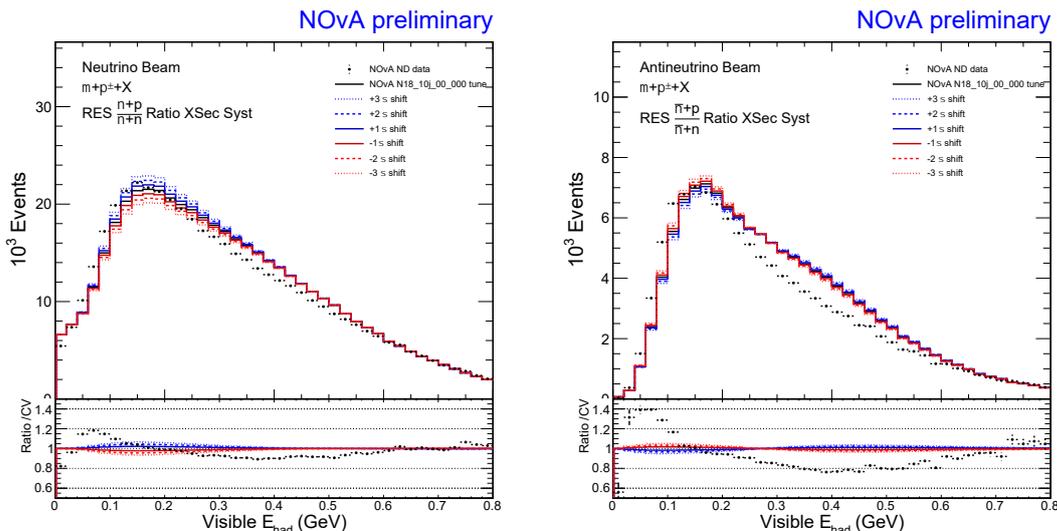


Figure 6.24: The RES_vpvn_Nu{Bar}XSecRatio pair of systematics. The impact of this pair of systematics shifts events around in the phase space, indicating it produces the desired effect.

+3 σ shift of RES_vpvn_NuXSecRatio in the FHC $\mu + \pi^\pm + X$ topology (left), we observe more events in this topology. Meanwhile a -3 σ shift predicts less events. For the RHC $\mu + \pi^\pm + X$ topology (right) the opposite effect is true with the $\bar{\nu}$ systematic. These added degrees of freedom should improve the **quality** of the ND fit, in particular for RES events.

6.5.2 Δ and Other Resonant Scaling Systematics

Two systematics are created to address the normalization of RES interactions: kRESScaleSystDelta and kRESScaleSystOther. As their names suggest, the first systematic scales the number of events from the Δ resonance, while the second systematic scales all higher order resonances, $W > 1.4$ GeV, referred to here as

“Other”. The scaling of each systematic is $\pm 20\%$ for a $\pm 1\sigma$ shift, scaled linearly ($\pm 60\%$ for a $\pm 3\sigma$ shift). Thus, these systematics only provide a normalization change and no shape change to the RES model, other than the explicit change to the W distribution of RES events. This change in W distribution can be seen in Fig. 6.25, where the effect of the systematics are shown. The physics motivation to create

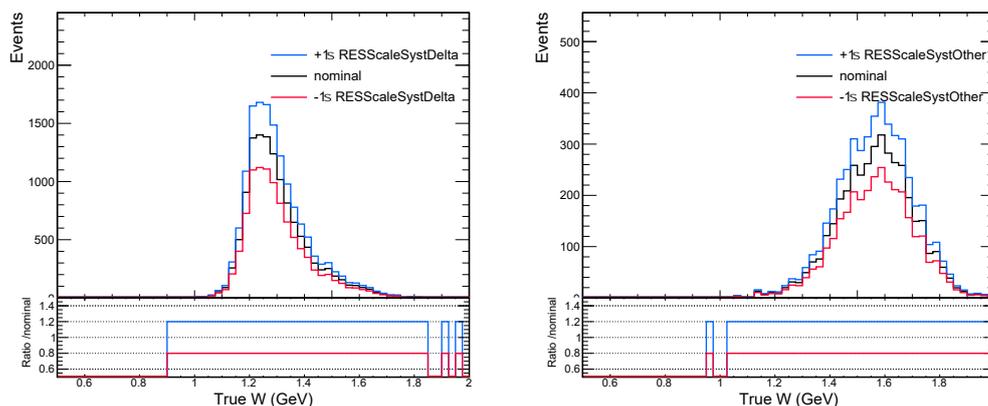


Figure 6.25: Validation of the `RESScaleSystDelta` (left) and `RESScaleSystOther` (right) on the NOvA simulation. The systematics are working as intended, scaling the W distribution from RES events by $\pm 20\%$ at $\pm 1\sigma$.

these systematics, originates from the over-prediction of RES interactions within GENIE, particularly at the low- Q^2 region of phase space. This over-prediction, as we have seen, is a leading contributor to the data-MC χ^2 disagreement in our ND topological samples. Therefore, we think it necessary to introduce this pair of systematics with sufficient freedom (20% at 1σ) to successfully “fit-out” this disagreement.

To further verify the intended effect, we show the effect of these systematics on the 2D $\mu + \pi + X$ topology. In Figure 6.26, we see the ratio of the $\frac{\pm 1\sigma}{Nominal}$ MC prediction for the `kRESScaleSystDelta` systematic in the 2D phase space. The top (bottom) row shows the FHC (RHC) $\mu + \pi^\pm + X$ prediction for the $\frac{-1\sigma}{Nominal}$ MC prediction on the left and the $\frac{+1\sigma}{Nominal}$ prediction on the right. We see that even $\pm 1\sigma$ shifts of these systematics alter the prediction by the expected 20% . This added degree of freedom will help to improve the agreement between the data and MC in the ND fit.

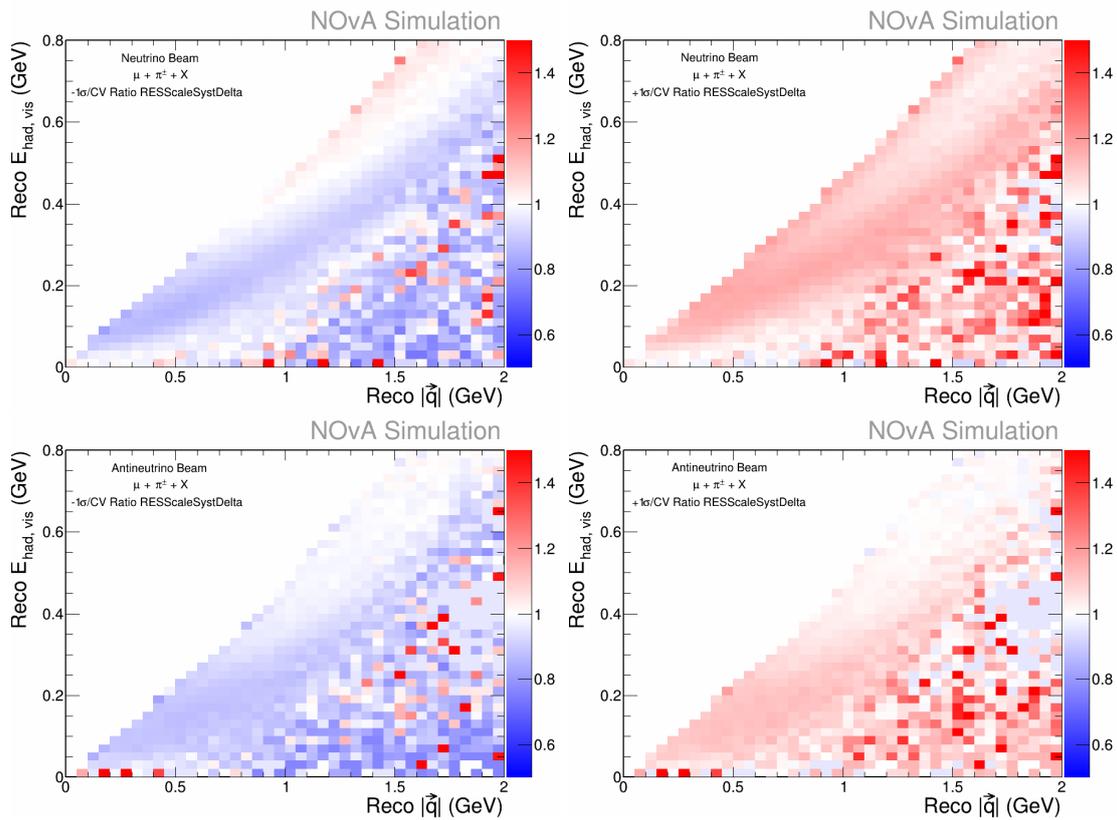


Figure 6.26: **Top:** Plots of the -1σ , left, ($+1\sigma$, right) divided by the nominal prediction ratio of the FHC $\mu + \pi^\pm + X$ topology for the `kRESScaleSystDelta` systematic. **Bottom:** Same as above, but for RHC $\mu + \pi^\pm + X$ topology. We see these systematics significantly alter the prediction from the nominal MC, given by the large ratio in many of the bins.

6.5.3 DIS Hadronization Systematics

Two new systematic uncertainties were created related to DIS hadronization: `kDISNuHadroQ1Syst` and `kDISNuBarHadroQ0Syst`. This pair of uncertainties aim to modulate final state hadron production in the Shallow Inelastic Scattering (SIS) and DIS regimes of neutrino interactions. Specifically, they address neutrino (antineutrino) DIS interactions that strike a neutron (proton) and contains final state multiplicity 2 – before FSI, with final state charge $Q=+1$ ($Q=0$ for antineutrinos). This means for neutrino interactions there are two final state hadron pair options with multiplicity 2 and charge $Q=+1$: 1 proton + $1\pi^0$ or 1 neutron + $1\pi^+$, while for antineutrinos the final state hadron pairs with $Q=0$ are: 1 neutron + $1\pi^0$ or 1 proton + $1\pi^-$. In relation to our ND fit topological samples, a ν_μ DIS interaction will change which topology the event is placed into depending on the final state particles produced. For example, a ν_μ CC DIS interaction that produces a hadron pair of 1 proton + $1\pi^0$ (rather than the 1 neutron + $1\pi^+$) will be placed into the $\mu + P$ or $\mu + P + X$ topology, rather than the $\mu + \pi^\pm + X$ topology if the final state hadron pair produced was 1 neutron + $1\pi^+$. In summary, identical CC DIS neutrino interactions will result in different ND topological categories depending on the final state hadron pairs produced.

Furthermore, the probability for a CC DIS interaction to produce one pair of final state hadrons over the other is equal, however, in GENE the probability is not equal; in GENIE, for a CC DIS neutrino interaction, the probability for producing a final state hadron pair – of multiplicity 2 and charge $Q=+1$ – is 1/3 for the proton pair and 2/3 for the neutron pair, while for antineutrinos the probability is reversed: probability 2/3 for the proton pair and 1/3 for the neutron pair. Seeing that the final state particles are predetermined and do not have sufficient flexibility illustrates a natural motivation to create uncertainties that can adjust the probabilities of the final state hadron pair. This is what the `kDISNuHadroQ1Syst` and `kDISNuBarHadroQ0Syst` systematics do for neutrinos and antineutrinos, respectively. We create weights that can modulate the production of one type of final state

Feature	kDISNuHadroQ1Syst	kDISNuBarHadroQ0Syst
incoming lepton	ν_μ	$\bar{\nu}_\mu$
interaction type	DIS	DIS
struck nucleon	n	p
final state hadron multiplicity	2	2
final state charge, Q	+1	0
hadron combinations	(1p + 1 π^0), (1n + 1 π^+)	(1n + 1 π^0), (1p + 1 π^-)
GENIE baryon prob.	p=1/3, n=2/3	p=2/3, n=1/3
uncertainty limits	2 σ , -1 σ	1 σ , -2 σ

Table 6.3: The characteristics of the two new DIS hadronization uncertainties created. One systematic (**kDISNuHadroQ1Syst**) only impacts neutrino interactions with final state multiplicity 2, charge $Q=+1$, and final state hadrons, while the other (**kDISNuBarHadroQ0Syst**) impacts antineutrino interactions with charge $Q=0$. Note the unique uncertainty limits for the two systematics; for example, a +2 σ (maximum) pull on the **kDISNuHadroQ1Syst** systematic will produce only 1 proton + 1 π^0 final state hadrons, while a -1 σ will produce only 1 neutron + 1 π^+ final state hadrons, maximally impacting which topological sample these CC DIS events fall into.

hadron pair over the other. This provides the desired flexibility to allow events to migrate from one topological sample to another.

The details of each systematic and the conditions relevant to these systematics can be seen further in Table 6.3. It is important to note the uncertainty limits for these two systematics. Setting an uncertainty to its limit only creates a single hadron final state pair and results in the other final state hadron pair having a probability of 0; this is the maximal effect.

Validations of this systematic can be seen in the plots of the ND topological samples, where we apply the **kDISNuHadroQ1Syst** systematic to the FHC topologies and **kDISNuBarHadroQ0Syst** to the RHC topologies. In Fig. 6.27 we show the systematics applied to the $\mu + \pi^\pm + X$ topology where this systematic is known to be the most effective. Note how the $\frac{-1\sigma}{Nominal}$ MC distribution of the **kDISNuHadroQ1Syst** systematic (top left plot) produces more events at -1σ than the nominal prediction, indicating that more 1n + 1 π final states are predicted and appearing in this topology.

At this stage we have introduced three new sets of uncertainties that aim to address the disagreement in the low- Q^2 region for predicted RES and DIS events.

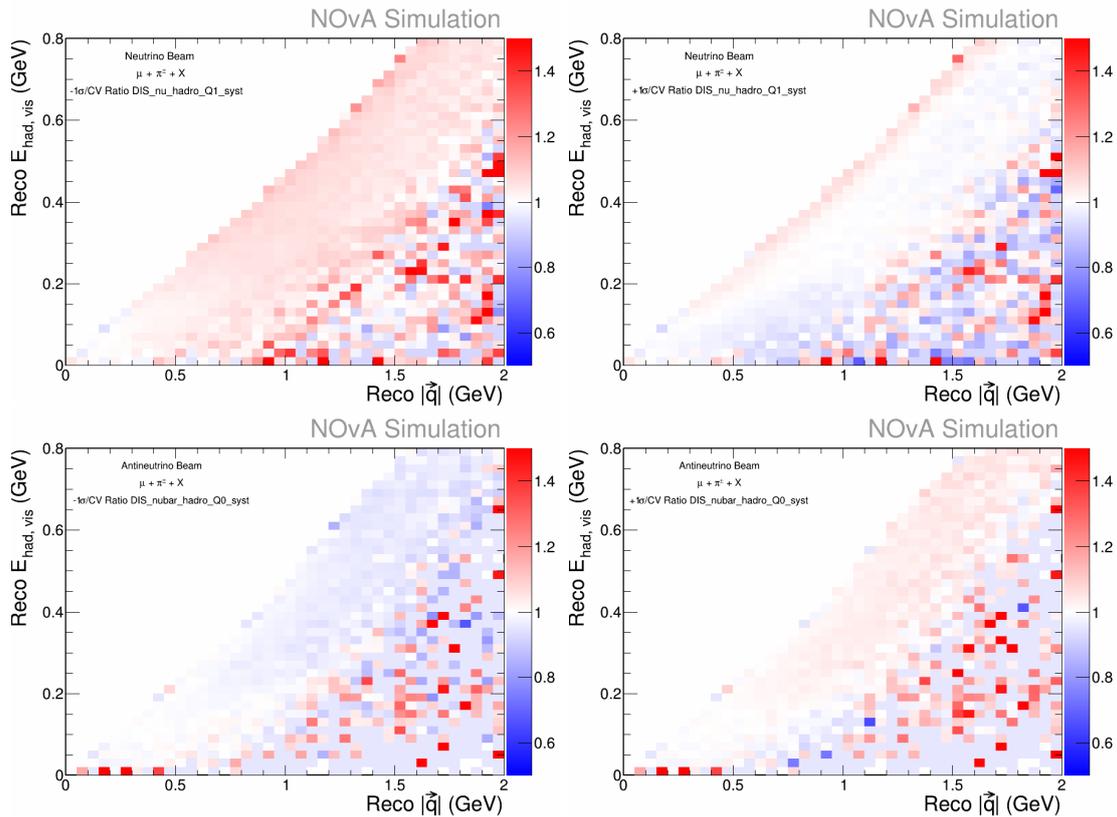


Figure 6.27: Plots of the new DIS hadronization systematics applied to the $\mu + \pi^\pm + X$ topologies of the ND fit. **Top:** the 2D $\frac{\pm 1\sigma \text{ shift}}{\text{nominal}}$ MC predictions of the `kDISNuHadroQ1Syst` systematic for the FHC topology. **Bottom:** the 2D $\frac{\pm 1\sigma \text{ shift}}{\text{nominal}}$ MC predictions of the `kDISNuBarHadroQ0Syst` for the RHC topology. For the FHC topology, increasing the `kDISNuHadroQ1Syst` knob to $+1\sigma$ reduces events in the $\mu + \pi^\pm + X$ topology to other topologies (at low E_{had}^{vis}). Conversely, pulling the knob to the other limit, -1σ , we see more events move into the $\mu + \pi^\pm + X$ topology. This same effect is reversed for the RHC $\mu + \pi^\pm + X$ topology.

With these six new added degrees of freedom, we believe we have introduced sufficient freedom into the uncertainty model for MC prediction to better agree with the NOvA ND data. Now we will attempt to fit the NOvA ND data again following these improvements made.

6.6 The ND Residual Difference

We are now ready to perform MCMC sampling again following the developments made to improve the **quality** and **quantity** of MCMC samples. We perform the *final* ND data fit and produce 500,000 MCMC samples. The final list of parameters used in the fit and their Representative Sample pulls from the final fit are listed in Table 6.4.

Table 6.4: MCMC Rep. Sample systematic Pull Summary from fit to ND data.

Systematic Parameter	Rep. Sample Pull (σ)
2ndclasscurr	0.513288
CalibShape	-0.0897525
Calibration	-0.290817
Cherenkov	-0.0581166
CorrMuEScaleSyst2020	3.01674
DISnuhadroQ1syst	1.02169
DISnubarhadroQ0syst	-1.77225
DISvbarnCC3pi2020	-1.77496
DISvbarpCC1pi2020	-2.35622
DISvbarpCC3pi2020	1.92755
DISvnCC1pi2020	-1.80565
DISvnCC2pi2020	1.36776
DISvnCC3pi2020	-2.24427
DISvpCC0pi2020	-2.07471

Continued on next page

Table 6.4 – *Continued from previous page* – ND Fit Pull Summary.

Systematic Parameter	Rep. Sample Pull (σ)
DISvpCC2pi2020	-0.924631
DISvpCC3pi2020	2.88617
FormZone2020	0.301691
LightLevelND	-0.330827
LowQ2RESSupp2020	0.95112
MECDoubleGaussEnhSystBaselineGSFProd5p1	2.85603
MECDoubleGaussEnhSystCorr1GSFProd5p1	1.89323
MECDoubleGaussEnhSystMeanQ01GSFProd5p1	-1.79964
MECDoubleGaussEnhSystMeanQ31GSFProd5p1	-3.28713
MECDoubleGaussEnhSystMeanQ32GSFProd5p1	3.33947
MECDoubleGaussEnhSystNorm1GSFProd5p1	0.0117823
MECDoubleGaussEnhSystSigmaQ01GSFProd5p1	-0.15476
MECDoubleGaussEnhSystSigmaQ02GSFProd5p1	0.81818
MECDoubleGaussEnhSystSigmaQ31GSFProd5p1	2.73205
MECDoubleGaussEnhSystSigmaQ32GSFProd5p1	2.5085
MECEnuShape2020AntiNu	1.61644
MECEnuShape2020Nu	0.104514
MECInitStateNPFrac2020AntiNu	1.10004
MECInitStateNPFrac2020Nu	0.0922145
MaCCRES	-0.643859
MvCCRES	0.836781
NeutronEvisPrimariesSyst2018	1.56971
RESDeltaScaleSyst	-1.07982
RESOtherScaleSyst	-2.68684
RESvpvnNuBarratioxsecsyst	-1.66419
RESvpvnNuratioxsecsyst	3.50826

Continued on next page

Table 6.4 – *Continued from previous page* – ND Fit Pull Summary.

Systematic Parameter	Rep. Sample Pull (σ)
RPAShapeenh2020	0.872877
RPAShapesupp2020	-2.4357
ThetaDelta2Npi	0.243223
ZExpAxialFFSyst2020EV1	2.08382
ZExpAxialFFSyst2020EV2	1.62305
ZExpAxialFFSyst2020EV3	-2.41613
ZNormCCQE	-0.766718
hNFSIFateFracEV12020	-1.8979
hNFSIMFP2020	3.37089
ppfxhadppc00	0.610238
ppfxhadppc01	-1.6092
ppfxhadppc02	0.332254
ppfxhadppc03	-2.72307
ppfxhadppc04	2.8194
radcornue	-0.75643
radcornuebar	-0.110261

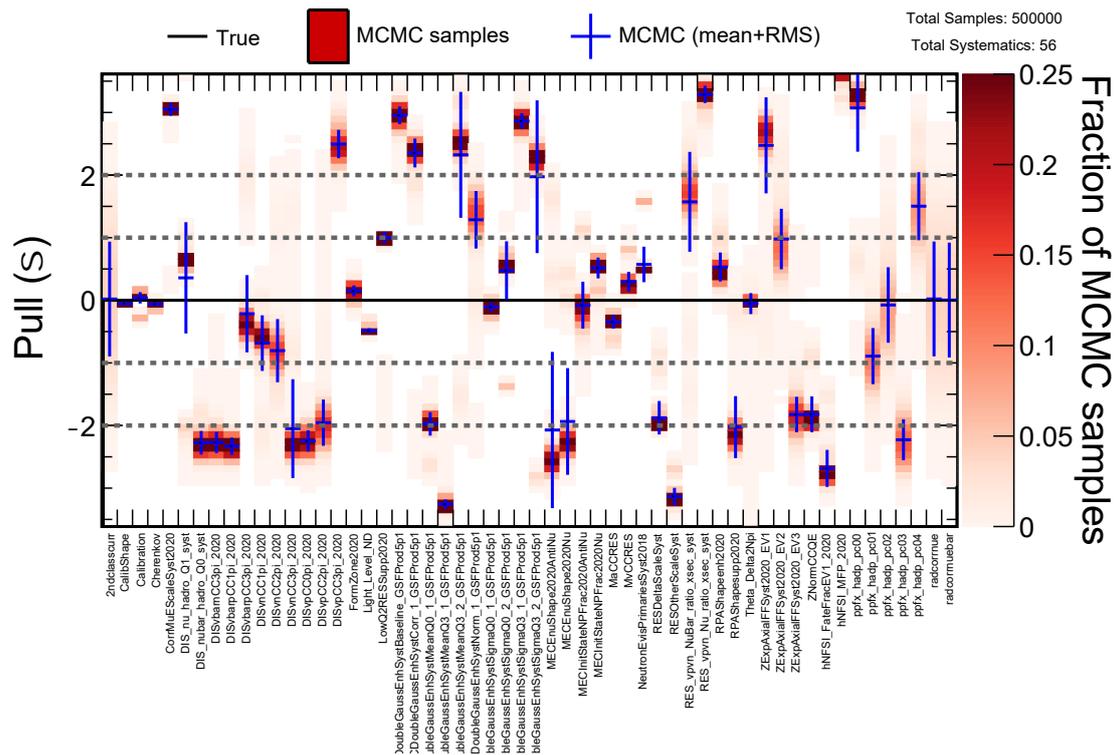


Figure 6.28: Pulls Summary plot from the ND fit to data. Note how nearly all systematics are well constrained by the data. There are also several parameters that are pulled to extreme values at or near $\pm 3\sigma$.

From examining the Pulls Summary plot (Fig. 6.28), we see now that nearly every parameter is well constrained from the ND data. Moreover, there are several systematic parameters that are pulled to extreme values. For example, the `LowQ2RESSupp` systematic is pulled to nearly exactly $+1\sigma$, indicating the most probable outcome for this parameter is to maximally suppress the low- Q^2 RES prediction. In addition the `ppfx_hadp_pc00` flux systematic is pulled to beyond $+3\sigma$. This suggests the systematic is providing an increase in the normalization to our ND predictions.

Looking deeper into the physics modeling, the `hNFSI_MFP_2020` parameter is pulled beyond $+3\sigma$, indicating the central value prediction from NOvA's hN FSI tune is insufficient to describe the data. In fact, the NOvA hN tune had reduced the mean free path by 40%, but the pull here reverses the tune and *increases* the mean free path further. This is strong evidence the hN FSI tune needs further attention. Moreover, we believe this increased MFP of the pion is related to the

neutron model and its uncertainty. For example, in a FSI charge-exchange scatter ($\pi^- p \rightarrow \pi^0 n$) a neutron can be produced, increasing the number of neutrons we expect; FSI processes are likely to produce neutrons and are intricately tied to the neutron response in NOvA. We must state how neutrons present themselves in the ND topologies. The neutron-dominant topologies are $\mu + P + X$ and $\mu + X$ (Figures 6.9 & 6.10) and generally appear as prongs (X). In these two topologies there is approximately a 40% over-prediction in the most populated bins of E_{had}^{vis} . Therefore, reducing the neutron activity is one means of improving the agreement. Returning to FSI, we now understand one possible explanation for such a large pull of the mean free path: reducing the number of neutrons. This is exactly what we see in the ND fit to data. It is worth noting that this is one of two possible ways to indirectly adjust the neutron model. The second is to adjust the detector response (via light level and detector calibration). This of overlapping FSI-neutron models is the strongest example of overlapping models we see in the ND fit.

We can examine the improvement in the data-MC agreement in the ND topologies by plotting the data-MC χ^2 in the 2D distributions before, Figure 6.29 & 6.31, and after, Fig. 6.30 & 6.32, the MCMC sampling. We must first note the different z -axis scale in the plots. Looking at the χ^2 distributions before the fitting – for both FHC and RHC topologies – we see that the largest contribution to the χ^2 is the low- Q^2 region along the kinematic boundary. However, when we examine the 2D χ^2 distributions after the fitting we see the this red region is improved significantly, by an order of magnitude in the μ and $\mu + P$ topologies. This is very strong evidence the developments made to the MCMC machinery improved the quality of the fit. Additional plots of the 2D data-Rep. Sample MC ratios can be found in the Appendix A.15 & A.16.

To better visualize the data-MC agreement after MCMC sampling, the χ^2 per topology is plotted again using the pulls associated from the Representative Sample (seen in Table 6.4) in Figure 6.33. We see the agreement improves relative to before the MCMC sampling, however, the $\chi^2 = 20,388$ for the Rep. Sample, and is still

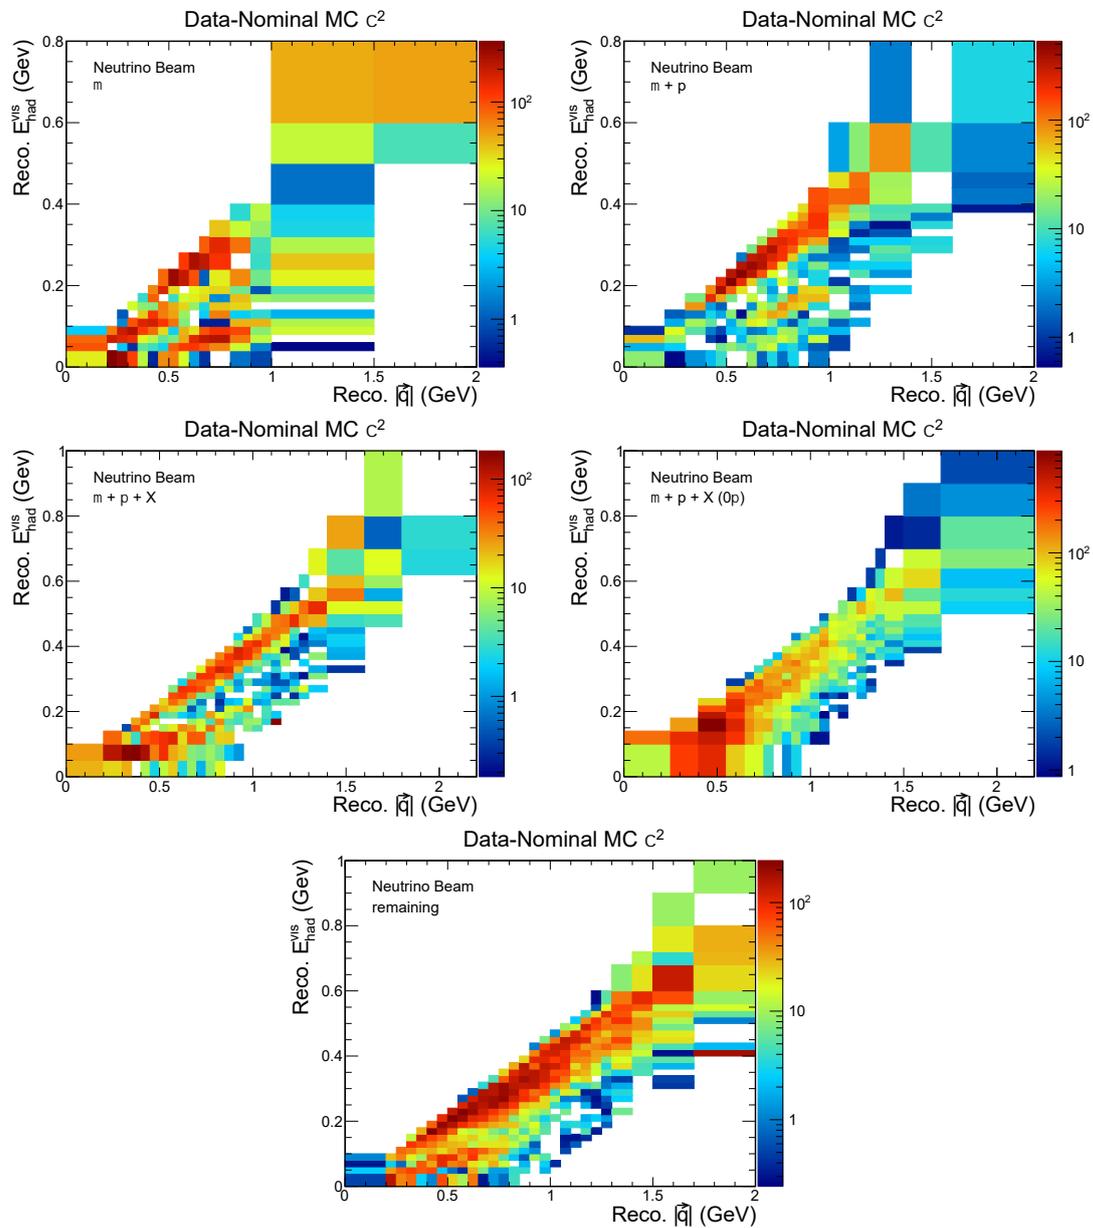


Figure 6.29: ND FHC 2D topological data-nominal MC χ^2 contribution per bin. This is showing the data-MC disagreement *before* the MCMC sampling. This demonstrates the regions where the disagreement is strongest (in red). Note the different z -axis scale in each plot.

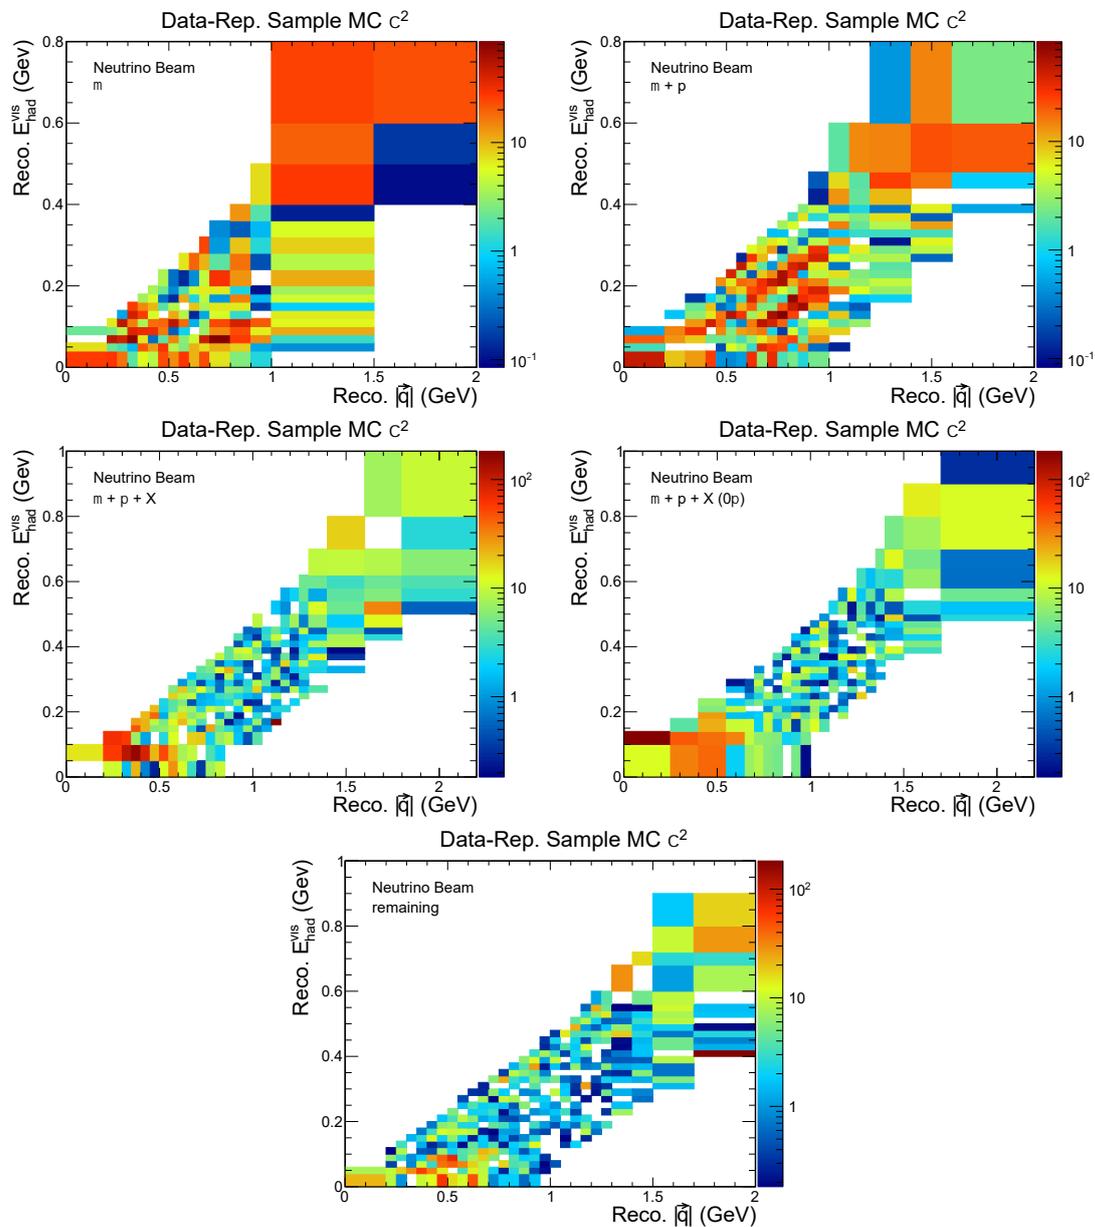


Figure 6.30: ND FHC 2D topological data-Rep. Sample MC χ^2 contribution per bin. This is showing the data-MC disagreement *after* the MCMC sampling when using the pull values associated with the Rep. Sample (Table 6.4). Many of the dark red regions in Figure 6.29 have been improved significantly, such as the region along the diagonal kinematic boundary. Note the different z -axis scale in each plot. The complimentary 2D data-Rep. Sample MC ratio plots can be found in the Appendix A.15.

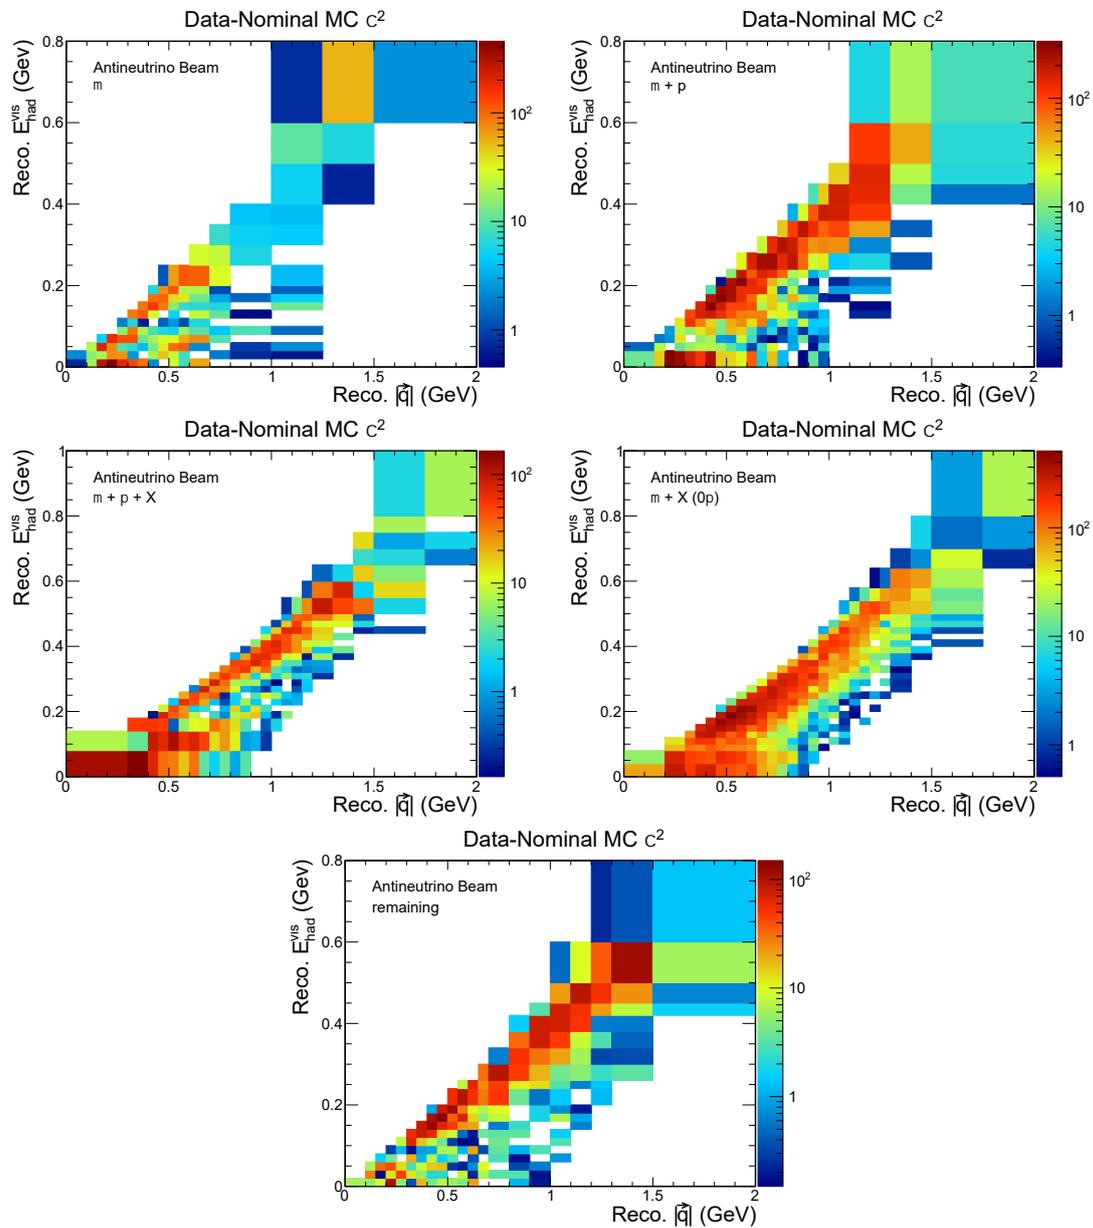


Figure 6.31: ND RHC 2D topological data-nominal MC χ^2 contribution per bin. This is showing the data-MC disagreement *before* the MCMC sampling. This demonstrates the regions where the disagreement is strongest (in red). Note the different z -axis scale in each plot.

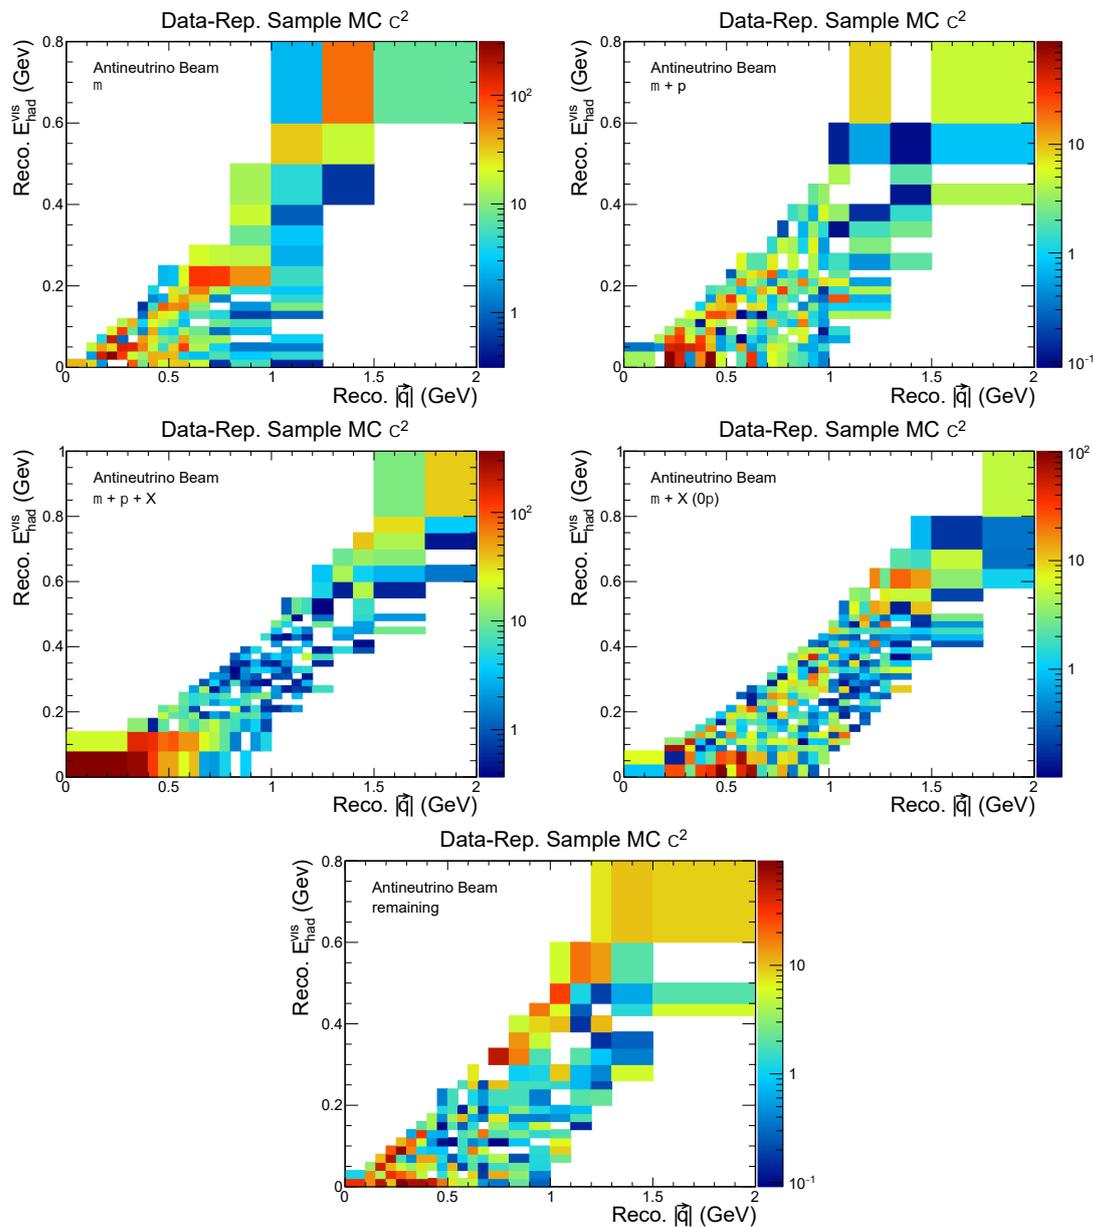


Figure 6.32: ND RHC 2D topological data-Rep. Sample MC χ^2 contribution per bin. This is showing the data-MC disagreement *after* the MCMC sampling when using the pull values associated with the Rep. Sample (Table 6.4). Many of the dark red regions in Figure 6.29 have been improved significantly, such as the region along the diagonal kinematic boundary. Note the different z -axis scale in each plot. The complimentary 2D data-Rep. Sample MC ratio plots can be found in the Appendix A.16.

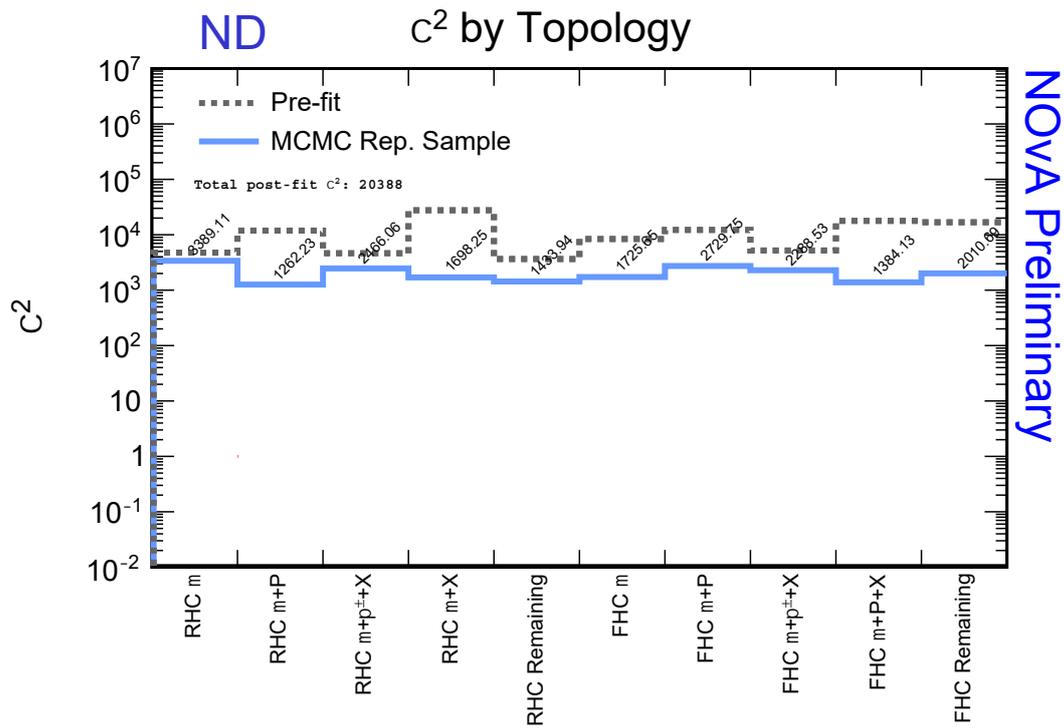


Figure 6.33: χ^2 per topology for the ND samples after improvements to the MCMC machinery. After these improvements, we still observe a non-zero χ^2 from the Representative Sample prediction. This indicates we have a **Residual Difference** in the data-MC agreement following the ND fit.

small enough to state we can fit the NOvA ND data exactly. Therefore, we have a **residual difference** in the ND fit. To effectively constrain the neutrino oscillation parameters, we must fit the data as best as possible, otherwise the Rep. Sample prediction could predict incorrect values of the oscillation parameters when we incorporate the FD predictions and data into an MCMC fit. This means we must address this residual difference. The next chapter will discuss this in more detail.

6.7 Conclusions from ND Fits

This chapter summarizes how NOvA can constrain systematic uncertainties with Markov Chain Monte Carlo. We describe the process of a NOvA ND fake data fit with a subset of parameters, illustrate how we analyze the MCMC samples, and demonstrate MCMC can successfully constrain the parameters. After, we move to fitting the NOvA ND data. Initially we find there are too many parameters to effectively sample the posterior and not every parameter in NOvA’s uncertainty model can be constrained from the ND data. To this end, numerous adjustments were made to the ND topological predictions to attempt to improve both the quantity of MCMC samples and the quality (i.e. produce better agreement to the data), such as creating new model-focused degrees of freedom. In summary, there are several important things learned from the NOvA ND fake data and real data fits, in particular relative to the physics modeling. The primary items learned can be highlighted:

Degeneracies in MEC model We have seen the MEC Double Gaussian parameters are highly degenerate *and* strong; we see multiple MCMC modes produce the same prediction of the ND topologies from a fake data fit. The adjustment of these predictions have the ability to significantly change the normalization in the topologies.

Over-prediction of RES events We have learned that NOvA’s simulation produces a non-trivial amount of RES events, especially at low- Q^2 . This motivated us to produce new systematic uncertainties to help combat this over-

prediction for RES events.

Inefficacy of hN FSI model Despite the tuning done to the hN model (summarized in Sec. 2.7), the consistently large pull values of the hN uncertainties – notable the mean free path systematic – in both fake and real data fits demonstrate the hN model is insufficient. Moreover, we believe this large pull of the MFP, which reduces the probability for FSI processes, is closely linked with the neutron response.

However, despite these important physics modeling conclusions, we still are unable to achieve perfect agreement with the NOvA ND data using the topological samples. In the next chapter we discuss how we plan to address the ND Residual Difference with respect to the NOvA Far Detector predictions and constraining the oscillation parameters.

Chapter 7

Near + Far Detector Joint Fit

Despite the analysis and improvements made to the MCMC sampling for the ND fit, the MCMC Representative Sample ND prediction still does not agree exactly with the ND data. This presents a unique question at this juncture: does the ND post-fit discrepancy impact the accuracy/efficacy of a constraint to the oscillation parameters Δm_{32}^2 , $\sin^2 \theta_{23}$, δ_{CP} ? This chapter attempts to answer this question by creating a FD fake data set from the FD nominal MC that contains the ND-only Rep. Sample pull values **and** the ND Residual Difference.

Section 7.1 first briefly discusses several possibilities for assessing the impact of the Residual Difference on the constraint of the oscillation parameters, and identifies which option this thesis performs. Section 7.2 describes the priors used for the oscillation parameters. Section 7.3 describes how the Residual Difference is applied to the FD ν_μ and ν_e predictions, and discusses an important conclusion learned about E_μ in this process. Section 7.4 analyzes the result of the ND+FD joint fit, with ND data and FD fake data, to constrain the oscillation parameters. Section 7.5 highlights several important improvements that can be made to improve the constraint on the oscillation parameters. Section 7.6 summarizes the work of the ND-only fit, and the ND+FD joint fitting.

7.1 Evaluating the Residual Difference

To assess the impact of the ND Residual Difference on a measurement of the neutrino oscillation parameters is a challenge. What's more, there are multiple means of asking this question. We list three possible options:

Fit the NOvA ND data perfectly, eliminating any Residual Difference. This

option is conceptually the simplest – improve the quality of the ND fit where the χ^2 per degree of freedom is unity in each ND topology. This could be done from, but not limited to, improving interaction modeling or incorporating additional model-driven uncertainties to MCMC. At which point a measurement of the oscillation parameters is maximally constrained from the ND data. The ideal case.

Introduce a systematic parameter to absorb the Residual Difference. This

parameter would likely adjust the corresponding truth quantities of $(Reco | \vec{q}|, E_{had}^{vis})$ to match the ND Rep. Sample prediction. Thus, the truth level information is required. Moreover, this would also require information of E_ν , the FD fit variable. Such a parameter would allow perfect agreement of the ND data, and any Residual Difference would take the form of a pull describing the difference between the $(Reco | \vec{q}|, E_{had}^{vis})$ and its truth quantities. This would likely require a systematic for each ND topological sample, further increasing the dimensionality of the posterior space.

Assess impact on oscillation parameters with fake data studies. This would

require applying the Residual Difference to a FD fake data set. From there, complete the ND+FD joint fit and determine if the oscillation parameters can still be constrained, despite the Residual Difference from the ND fit.

This thesis selects the third option – via a fake data study – to assess the impact of the ND Residual Difference on the oscillation parameters. There are multiple methodologies to perform such a fake data study. We mention two possibilities. One option is produce a fake data set with known inputs to the fake data, thereby

replicating the Residual Difference. The advantage here is the results of a study are easy to interpret as a clear comparison can be made between the inputs and the MCMC result. However, the challenge is to produce a fake data set that resembles the Residual Difference *with known inputs*. Crucially, if there was a known set of inputs that would produce the ND Residual Difference, we can correct the Residual Difference at the source; the expectation is to correct the models producing the differences rather than simply produce a fake data study. This is the crux to the using a fake data set with known inputs.

The second option for a fake data set is to apply the Residual Difference directly to create a FD fake data set. The advantage to this strategy is the Residual Difference from the ND is captured exactly to the FD fake data. This assumes the Residual Difference from the ND would be identical in the FD. The disadvantage here is the results can be difficult to interpret, as the inputs into the fake data are not clearly known. This is counter to the previous option. This final chapter elects to perform a fake data study via this option: **directly applying the Residual Difference to a FD fake data set**. The following sections outline this procedure in more detail.

7.2 Priors of Oscillation Parameters

The oscillation parameters Δm_{32}^2 , $\sin^2 \theta_{23}$, and δ_{CP} are constrained by external data, therefore we possess prior knowledge of what we anticipate the values of these parameters to be. This section discusses the prior choice for each parameter.

7.2.1 Prior $\sin^2 \theta_{23}$

We use a uniform prior for the measured variable $\sin^2 \theta_{23}$. NOvA is sensitive to $\sin^2 \theta_{23}$ (rather than θ_{23} of the PMNS matrix), so the prior used is in the measured variable. We use a uniform prior in this variable because the octant of θ_{23} is not known ($\theta_{23} < 45^\circ$ or $\theta_{23} > 45^\circ$). With this prior, MCMC will sample both octants

uniformly from the posterior.

7.2.2 Prior Δm_{32}^2

The parameter Δm_{32}^2 refers to the energy at which the oscillation “dip” region occurs. The value has been measured within NOvA to be $2.41 \times 10^{-3} \text{ eV}^2$ for Normal Ordering (NO) [13]. A “Top Hat” prior is selected for this parameter [65]. This function creates a uniform probability density within the “Physical Region” where we expect the true value and quickly falls away to zero probability in the “Smoothing Region”, all while still being differentiable. This allows MCMC to sample only this region and restricts MCMC from sampling beyond the edges of the “Top Hat.”

As discussed in Sec. 1.5, the value of Δm_{32}^2 can be positive for Normal Ordering (NO) or negative for Inverted Ordering (IH). This presents a degeneracy in the parameter value; because the mass eigen states m_2 and m_3 are squared, we do not know the individual values, only the squared difference. Therefore, Δm_{32}^2 can be either negative or positive, $\Delta m_{32}^2 \approx \pm 2.4 \times 10^{-3} \text{ eV}^2$, and thus, we must consider both mass orderings in our MCMC sampling. Moreover, the value of Δm_{32}^2 is unlikely to be between these two values. We implement a switching within our sampling that will force the MCMC chain to “jump” between the two mass orderings with a 50% probability; this ensure that we sample both mass orderings efficiently.

7.2.3 Prior δ_{CP}

The prior choice for δ_{CP} is unique because this parameter is periodic from $0 - 2\pi$. Therefore we require only this parameter range *and* must ensure that the end points of this parameter space (0 and 2π) are continuous and differentiable. We use a ‘Bounded Step’ Prior, Figure 7.1. This prior restricts MCMC samples to within the single phase, $0 \leq \delta_{CP}/\pi \leq 2$, while values beyond this phase are permitted, but quickly diminish to zero in the regions $[-1\pi, 0\pi]$ and $[2\pi, 3\pi]$ (grey regions). To account for the values beyond the single phase, the areas of the grey regions and the central region sum to unity, conserving probability, making the prior effectively

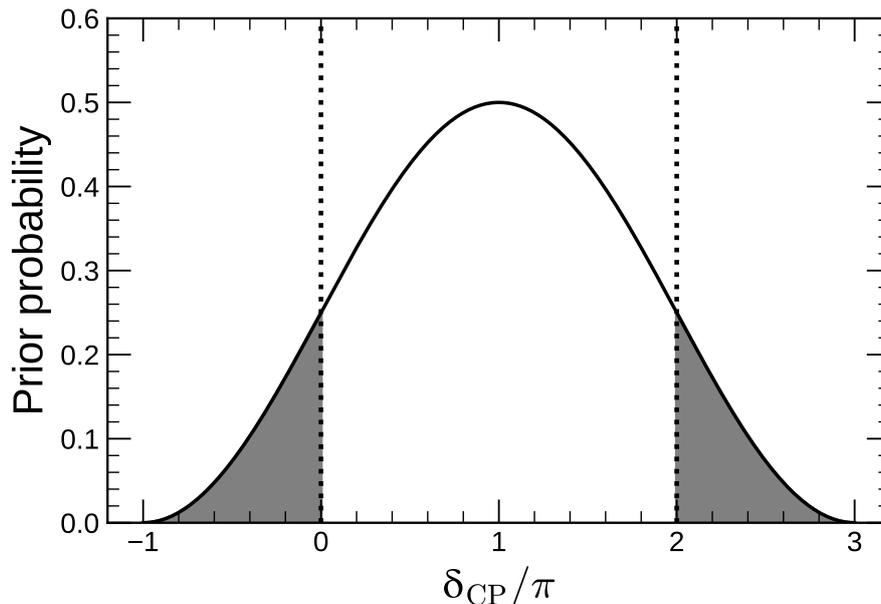


Figure 7.1: Prior function for δ_{CP} oscillation parameter. We use this specific function because this parameter is periodic. The integral of the shaded regions is equal to the difference between unity and the function within $0 \leq \delta_{CP}/\pi \leq 2$, therefore the probability within $[0\pi, 2\pi]$ is uniform.

uniform within the single phase. This prior is also differentiable making it accessible for HMC.

7.3 Creating the FD Fake Data with the ND Residual Difference

This section outlines the process for creating the FD fake data from the residual, post-fit discrepancy that remains from fitting the NOvA ND prediction to its data in the topological samples. It is important to note the “Residual Difference fit” does not ascribe any source to the residual difference in the FD, but assumes this difference appears equally in the FD as it does in the ND. Therefore we anticipate the Residual Difference to appear in the same way in reconstructed quantities in both detectors.

The FD fake data we are constructing contains three pieces:

- NOvA 2020 Best Fit Point for oscillation parameters (Table 7.1).

- Rep. Sample pull values for systematic parameters from ND-only fit to NOvA data (Table 6.4).
- ND Residual Difference from ND-only fit to NOvA data (Figure 6.33).

This is the FD fake data we will use to perform the ND+FD joint fit of the Residual Difference.

The following two Subsections 7.3.1 & 7.3.2 describe the process of applying the ND Residual Difference to the FD ν_μ and ν_e samples, respectively. Applying the Residual Difference is unique for the ν_μ and ν_e samples. This process is outlined below for the two neutrino flavors.

7.3.1 Application to ν_μ

The ND topological samples are all of ν_μ events; applying the Residual Difference to the FD ν_μ sample is straightforward. However, this is an imperfect process. We note this does not capture *exactly* the residual differences due to the different volumes and acceptances between the ND and FD, but this process should approximately capture the effect. We seek to extrapolate the ND Residual Difference to the FD ν_μ quantile distributions. First we must divide the ND ν_μ interactions into the same quantile cuts as the FD selection, where we divide the events into four quantiles of hadronic energy fraction. From there, the ND-fit Rep. Sample prediction (blue) is plotted in each of the four quantiles for FHC and RHC alongside the data (black points) in Figure 7.2 and Figure 7.3. By observing the Residual Difference in *Reco* E_ν , we learn more about the quality of the ND fit agreement. These figures show the agreement between the data and the Representative sample (blue) is most discrepant in Quantile 1, in particular at higher energies of $E_\nu > 2$ GeV. This can be seen in the event counts (top plot) and in the ratio (bottom plot) where the disagreement is as much as 20% in certain bins.

To observe a 20% disagreement in the ν_μ energy spectrum in the ND *after* MCMC sampling is an unexpected result. In other words, the Rep. Sample ND

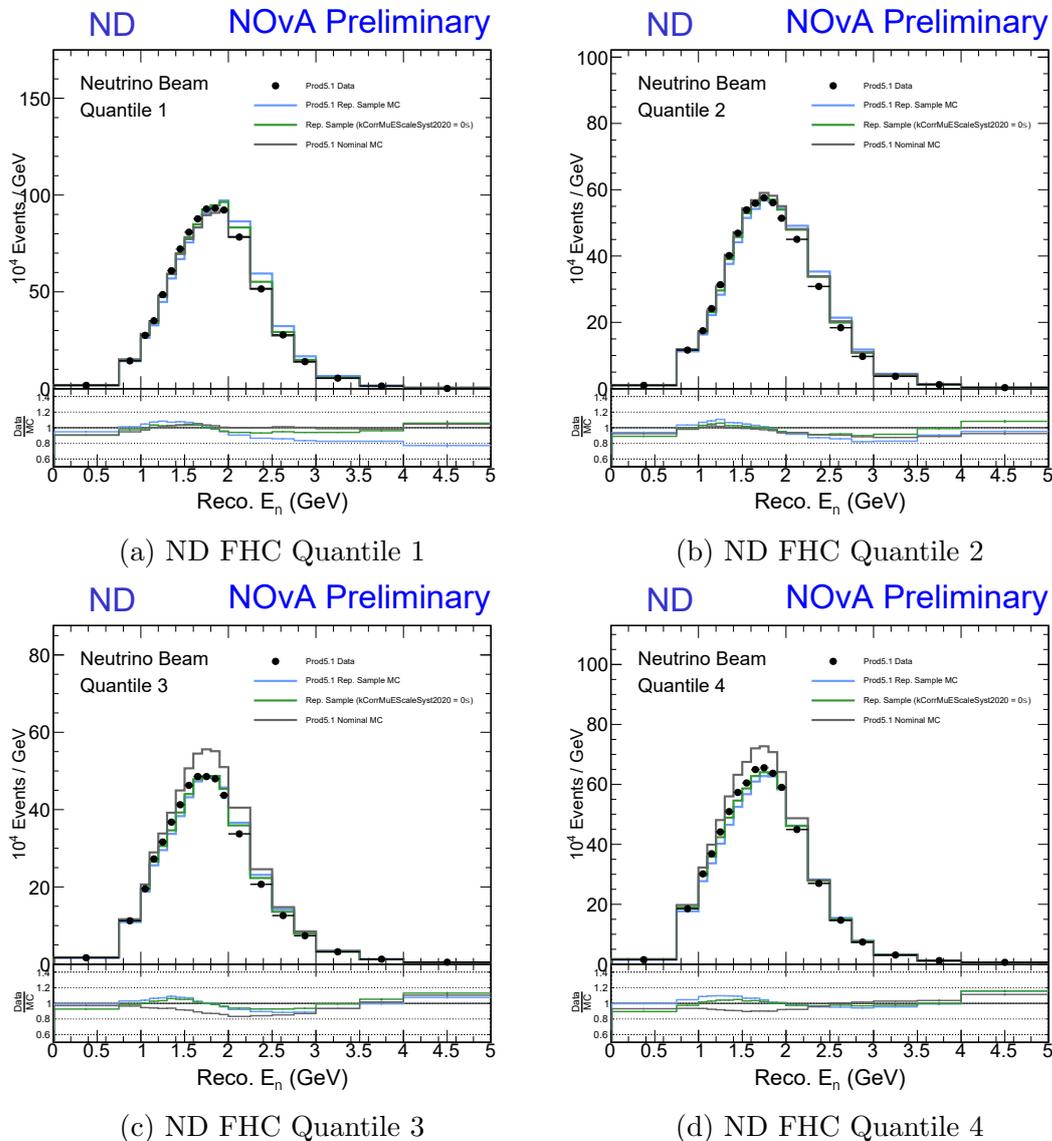


Figure 7.2: ND FHC Rep. Sample Prediction (blue) in the 2020 FD quantile cuts. Note the agreement improves with increasing quantile. This is because Quantile 1 contains the most E_μ , which is not directly constrained in the ND fitting. The green prediction is the Rep. Sample, but without the muon energy scale systematic, `kCorrMuEScaleSyst2020`. Note how the prediction improves when removing this systematic, especially in Quantile 1. We will use the green prediction to apply the Residual Difference to the FD MC.

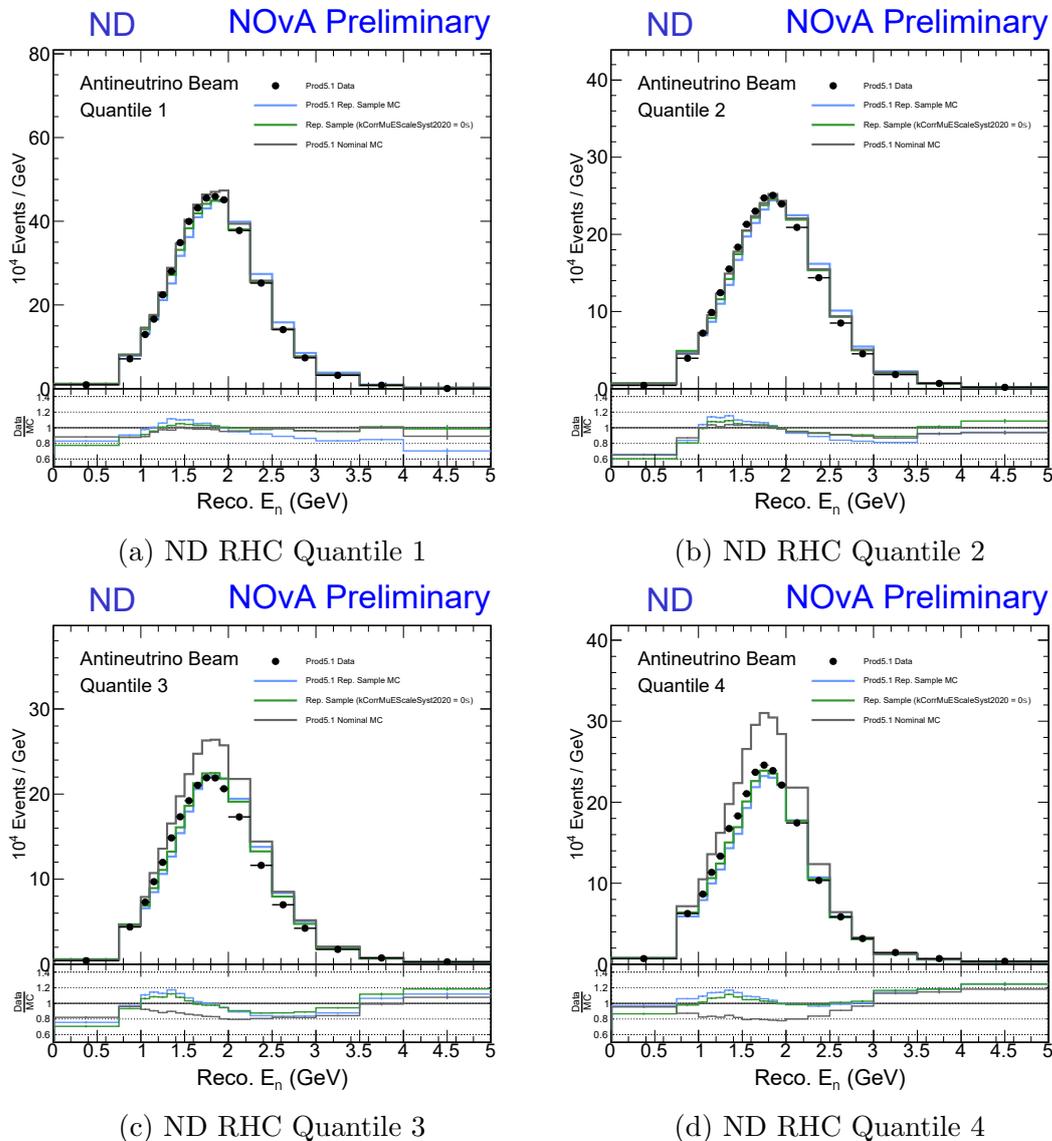


Figure 7.3: ND RHC Rep. Sample Prediction in the 2020 FD quantile cuts. Note the agreement improves with increasing quantile. This is because Quantile 1 contains the most E_μ , which is not directly constrained in the ND fitting. The green prediction is the Rep. Sample, but without the muon energy scale systematic, `kCorrMuEScaleSyst2020`. Note how the prediction improves when removing this systematic, especially in Quantile 1. We will use the green prediction to apply the Residual Difference to the FD MC.

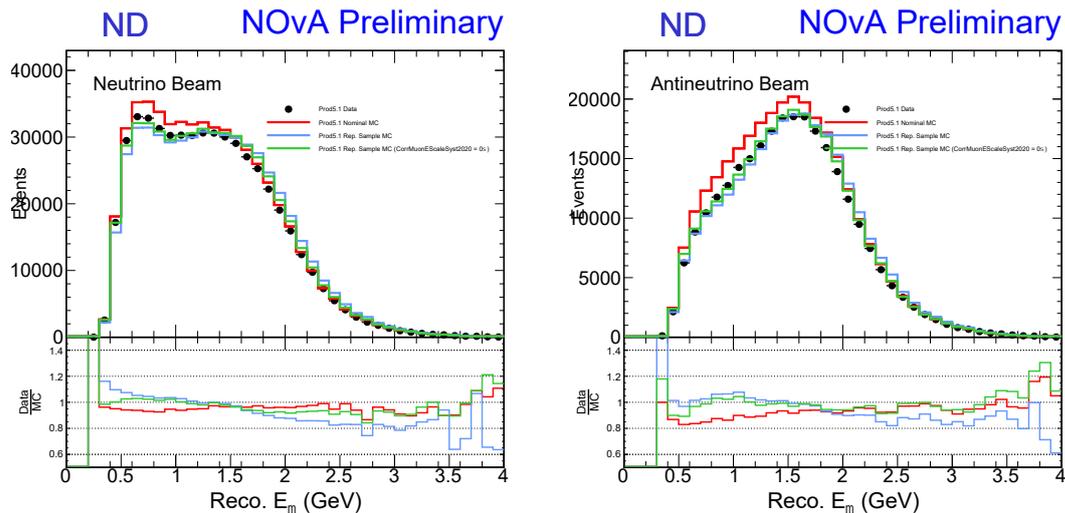


Figure 7.4: ND Reconstructed muon energy for FHC (left) and RHC (right). The disagreement – seen in the data-MC ratio – of the Rep. Sample (blue) distribution at higher E_μ shows the same disagreement as the ND quantiles Rep. Sample (blue) distributions in Figs. 7.2 & 7.3. The muon energy at $E_\mu < 2$ GeV is under-predicted, and at $E_\mu > 2$ GeV is over-predicted relative to the data. Meanwhile, when the `kCorrMuEScaleSyst2020` systematic is removed from the ND fit result (prediction in green), the data-MC ratio is improved by as much as 20%. The large pull value of `kCorrMuEScaleSyst2020` is the cause of the 20% discrepancy in the ND $Reco E_\nu$ Rep. Sample predictions.

prediction – when plotted in the FD quantile cuts – produces *worse* agreement to the data than the nominal MC (grey). To understand why this is, we consider the variables used in the ND fit: E_{had}^{vis} & $Reco |\vec{q}|$. Trivially, E_{had}^{vis} includes no information of the leptonic component of the interaction and therefore does not constrain the muon component in any meaningful way. However, $Reco |\vec{q}|$ *does* contain a component of the muon energy, from Eq. 4.1:

$$\begin{aligned}
 |\vec{q}| &= \sqrt{Q^2 + (E_\nu - E_\mu)^2} \\
 &= \sqrt{2E_\nu(E_\mu - p_\mu \cos \theta_\mu) - m_\mu^2 - (E_\nu - E_\mu)^2}.
 \end{aligned}
 \tag{7.1}$$

It was anticipated that this E_μ dependence in $Reco |\vec{q}|$ would sufficiently constrain the muon energy. However, E_μ is only a second-order dependence in $Reco |\vec{q}|$ – after E_ν – in this variable. Figure 7.4 of the reconstructed muon energy shows the same disagreement in E_μ as we see in the ND quantile distributions in Fig. 7.2 & Fig. 7.3. The Rep. Sample prediction is under-predicted at $E_\nu < 2$ GeV and over-predicted

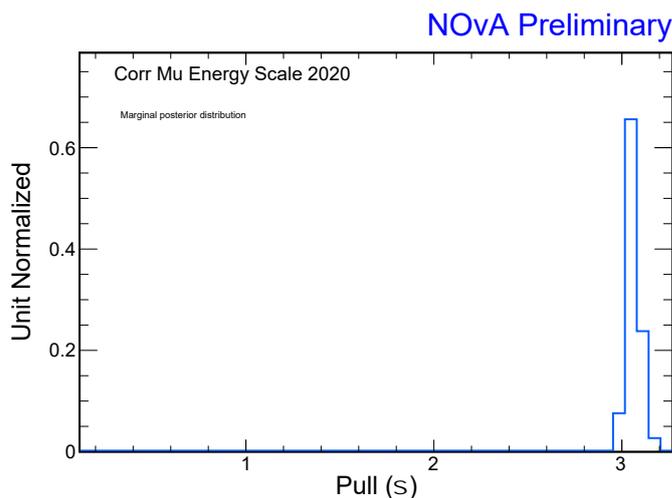


Figure 7.5: 1D marginal of the `kCorrMuEScaleSyst2020` systematic from the ND fit to data. The pull is strongly constrained to 3σ . The result is prediction of the muon energy that is shifted to higher E_μ respect to the data, seen in the blue prediction in Figure 7.2, Figure 7.3 of E_ν and Figure 7.4 of E_μ .

at $E_\nu > 2$ GeV. To understand which parameter from the ND fit would cause this disagreement, we note there is only one single systematic uncertainty that affects the muon energy and nothing else: `kCorrMuEScaleSyst2020`. This systematic was included in the ND fit and contains a mode in the 1D marginal distribution of $+3.1\sigma$, Figure 7.5. Therefore, the large, positive pull of this parameter predicts larger muon energies, shifting the predictions from Figs. 7.2 & 7.3 to the right (and down in the ratio plots), which is exactly the same effect we observe in E_μ (Fig. 7.4).

Again, the pull on the muon energy scale uncertainty is large from the ND fit because the muon energy is not constrained at all in E_{had}^{vis} and only has a second-order dependence in $Reco |\vec{q}|$. Therefore, a large pull value is highly probable to improve the quality of the agreement at the few %-level in the $Reco |\vec{q}|$ variable. In other words, MCMC slightly improves the χ^2 at the cost of a very large pull on the parameter. This is an important conclusion from the ND fit: a third variable, likely $Reco E_\mu$, is necessary to improve the accuracy of the prediction in $Reco E_\nu$. This is an important conclusion for future work to better constrain the oscillation parameters using a simultaneous fit.

Eliminating `kCorrMuEScaleSyst2020` from the ND Rep. Sample

We have learned the muon energy scale systematic adversely impacts the prediction of the neutrino energy compared to the ND data. It was expected that $Reco |\bar{q}|$ would be sufficiently sensitive to the muon energy, but we have learned it is not. Further, we see evidence that setting the systematic to a pull of 0σ is the correct choice as the data-MC agreement in $Reco E_\mu$ is improved significantly. We see the result of this choice in the Rep. Sample predictions in Fig. 7.2 and Fig. 7.3 in the green histograms. When we remove this uncertainty we see the improvement in the agreement in the ND quantiles is identical to the $Reco E_\mu$ distributions in Figure 7.4. Removing the uncertainty reduces the disagreement between the data and Rep. Sample prediction to approximately 10%, especially Quantile 1, which is dominated by the muon component and contains very little hadronic energy. As we apply the Residual Difference to the FD ν_μ prediction, we will now refer to the Representative Sample as the green prediction in the previous plots – the predictions without the muon energy scale systematic.

We have removed the uncertainty from the constraining ND predictions, however, we do not want to remove it from constraining the FD predictions; as we know from the quantile cuts, the muon energy is important to the ν_μ quantile distributions, especially Quantile 1. Therefore, when performing the joint ND+FD fit, we will allow `kCorrMuESscaleSyst2020` to be constrained by the FD fake data only. This is another advantage of being able to constrain the predictions from the ND and FD simultaneously.

We take the ratios from Figures 7.2 and 7.3 of the ND data and the *green* Rep. Sample prediction (without the muon scale systematic) and apply the ratio in each ND quantile bin to the corresponding FD prediction bin’s event count. This is how we apply the ND Residual Difference. Figures 7.6 and 7.7 show the ν_μ distributions after applying the ND Rep. Sample pull values and the ND Residual Difference in the orange dashed line. (The other colored histograms are the different predictions at the intermediate stages of creating the FD fake data, e.g. ND Rep. Sample

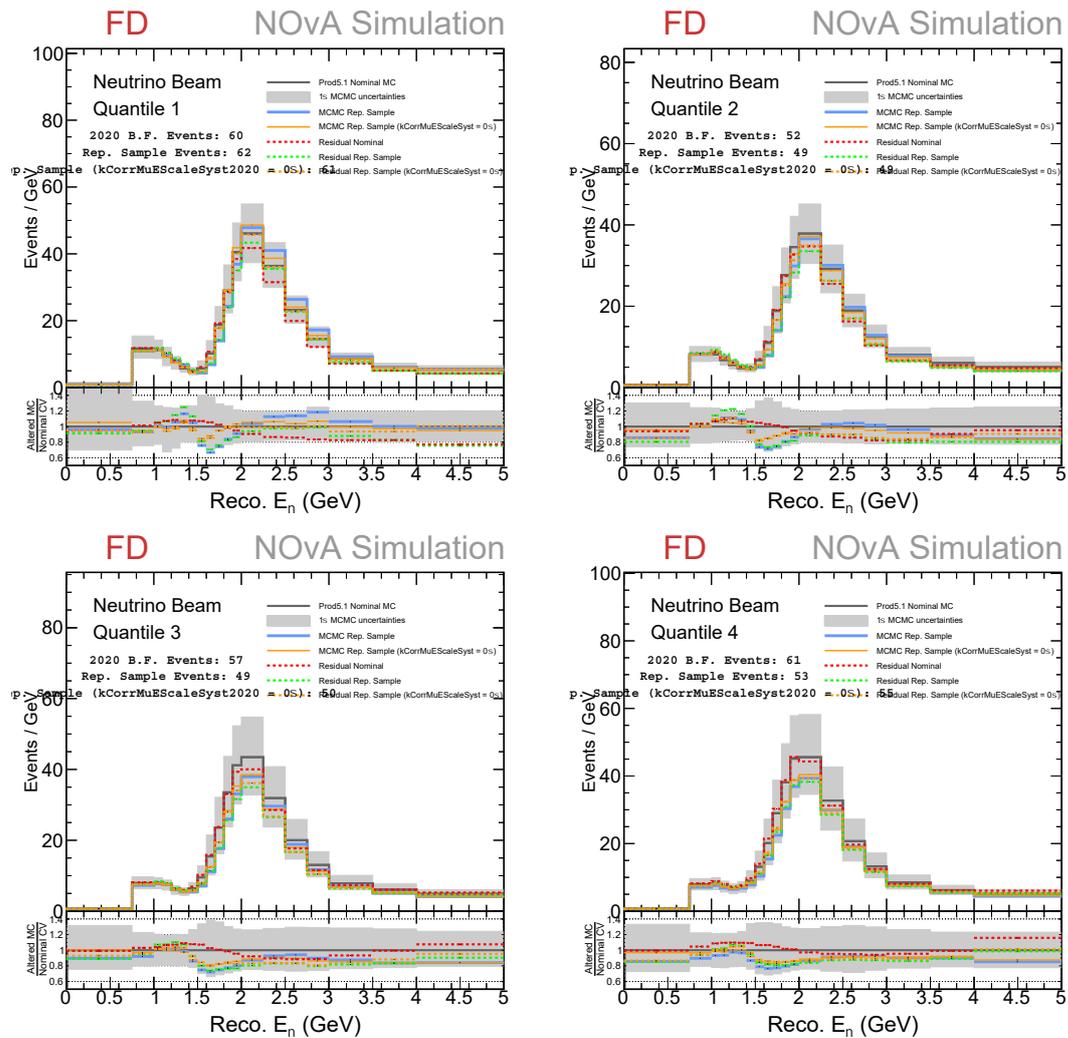


Figure 7.6: FHC FD MC predictions with different combinations of ND-only Rep. Sample pulls and/or ND Residual Difference applied in the four quantiles. The orange dotted histogram is the FD fake data for this fit.

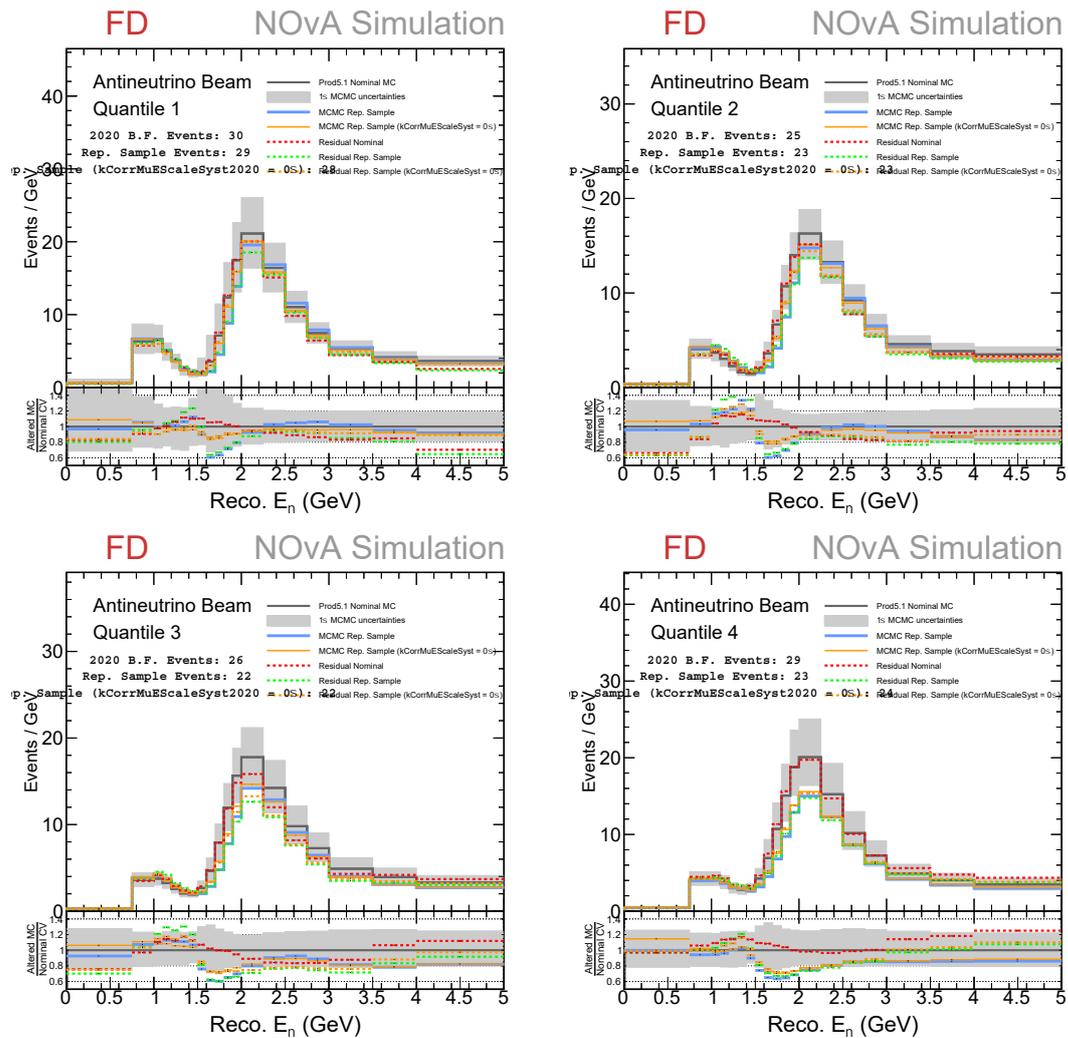


Figure 7.7: RHC FD MC predictions with different combinations of ND-only Rep. Sample pulls and/or ND Residual Difference applied in the four quantiles. The orange dotted histogram is the FD fake data for this fit.

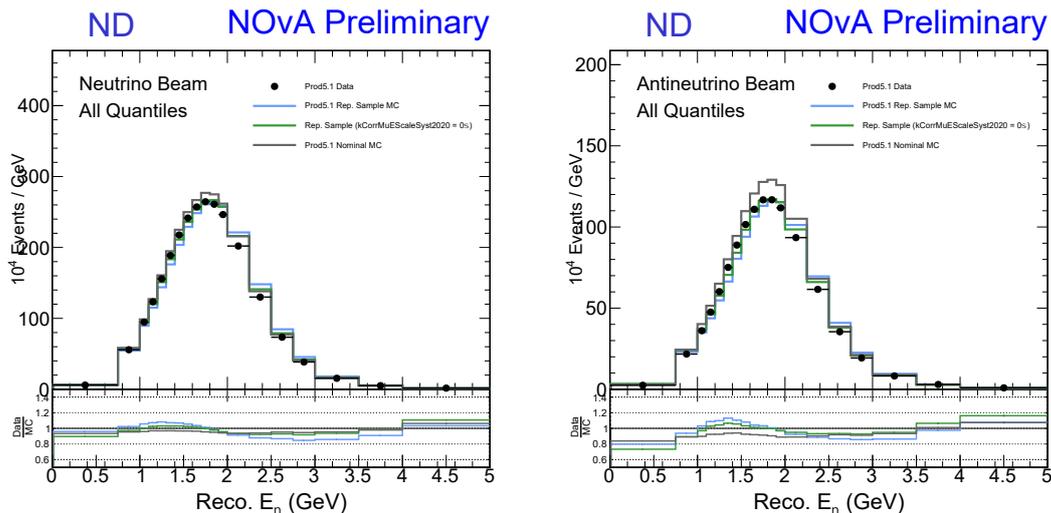


Figure 7.8: ND FHC (left) & RHC (right) Inclusive ν_μ Sample in E_ν . The blue prediction is the Representative Sample with the `kCorrMuEScaleSyst2020` included and the green is the prediction without the systematic, which will be used for the Residual Difference. To apply the Residual Difference to the Core ν_e sample, we interpolate the data-Rep. Sample MC ratio from these ν_μ inclusive samples in E_ν . For the Peripheral ν_e sample, we use the ratio of the data-Rep. Sample MC integrated events.

pulls only with and without the muon energy scale systematic, and the Residual difference only. These other histograms are for reference only and help to illustrate how the predictions change.) In these figures, it is notable the Rep. Sample pulls from the ND fit change the prediction the most (the grey to orange histogram). After applying these pulls, as you move from Quantile 1 to Quantile 4, the impact of the Residual Difference changes the prediction the least (solid orange to dotted). The orange dotted predictions of the ND Rep. Sample pulls and the ND Residual Difference will be the fake data for the FD ν_μ sample in our joint fit.

7.3.2 Application to ν_e

The ν_e sample is not explicitly constrained in the ND fits as all ten ND topological samples are ν_μ events. However, we can apply a constraint on the oscillated ν_e events in the NOvA Far Detector, which originated as ν_μ . Thus, we will apply the Residual Difference from the ND solely to the *appeared* ν_e in the NOvA FD simulation.

We apply this constraint on the FD ν_e samples by utilizing the inclusive ND

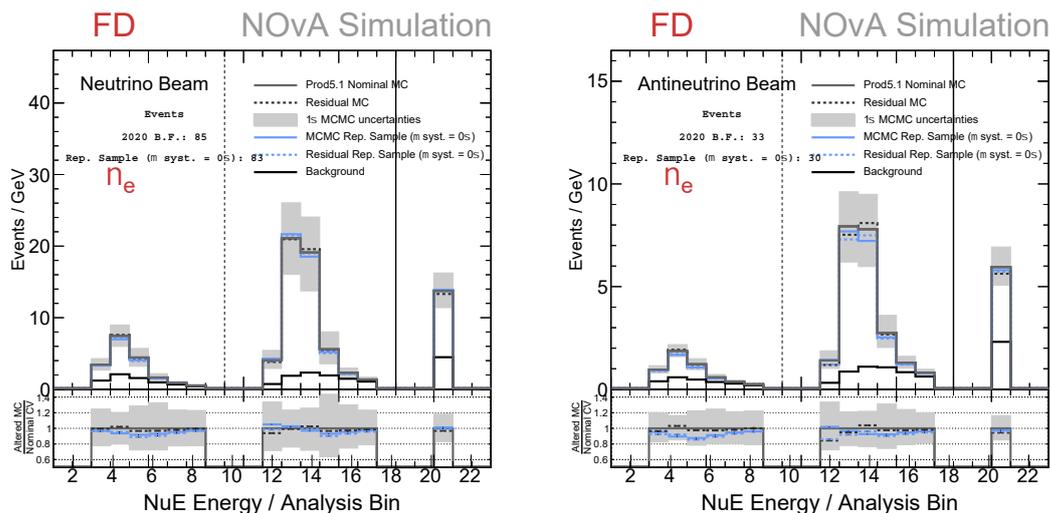


Figure 7.9: FD FHC (left) & RHC (right) ν_e distributions with the ND-only fit Rep. Sample pulls (blue solid) and distributions combined with the ND Residual Difference (blue dashed). The blue dashed prediction is the FD ν_e fake data.

ν_μ distribution in $Reco E_\nu$, rather than using a specific quantile prediction, since we do not discriminate ν_e events by hadronic energy fraction. For the two Core samples of the ν_e distribution, we interpolate the data-Rep. Sample MC ratio of the ND ν_μ inclusive sample to apply to the corresponding Core ν_e bins [66]. For the Peripheral sample, we use the integrated data-MC events ratio from the inclusive ν_μ sample to apply to the single Peripheral bin. The ν_e Residual Difference prediction is shown for FHC and RHC in Figure 7.9. We see the ND-fit's Rep. Sample pulls applied to the nominal simulation (blue solid) and the ND-fit Rep. Sample pulls with the ND Residual Difference applied (blue dashed) does not significantly change the prediction for ν_e events.

Similar to the muon energy scale treatment for the FD ν_μ sample, we allow the radiative corrections systematics to only constrain the FD ν_e sample when performing the joint fit; they have no constraining power in the ND event selection, but we anticipate they will impact the FD ν_e event sample.

7.4 ND + FD Residual Difference Fit Result

This section outlines the results of the ND+FD joint Residual Difference fit. This joint fit contains 60 parameters: 57 systematic uncertainties and 3 oscillation parameters. Again, recall this Residual Difference fit is a hybrid of real NOvA data and fake data; for the ND we use real, Prod5.1 NOvA data – the same data from the ND-only fit in the previous chapter. For the FD, we use fake data constructed in the following way (outlined in the previous Section):

- NOvA 2020 Best Fit Point for oscillation parameters (Table 7.1).
- Rep. Sample pull values for systematic parameters from ND-only fit to NOvA data (Table 6.4).
- ND Residual Difference from ND-only fit to NOvA data (Figure 6.33).

This hybrid of NOvA ND data and FD fake data is used for the joint fit. As we saw from the ND data fit, there are still non-trivial residual differences between the fitted MC (i.e. the Rep. Sample) and the NOvA ND data, despite the adjustments/improvements made to the ND machinery. This leads to the question of how significant is this Residual Difference in measuring the oscillation parameters. The analysis of this ND+FD joint fit is reported in this Section.

Parameter	Value
L	810 km
ρ	2.84 g/cm ³
θ_{12}	0.587
θ_{13}	0.184
Δm_{21}^2	7.53×10^{-5} eV ²
Δm_{32}^2	2.41×10^{-3} eV ²
$\sin^2 \theta_{23}$	0.57
δ_{CP}	0.82π

Table 7.1: The top section of the table represents the fixed parameters input into this joint ND+FD fit. The bottom half represents the true values of the parameters we are attempting to constrain.

This fake data fit is conducted using the NOvA 2020 Frequentist analysis best fit oscillation parameters as the Asimov – or true – point for our fake data. The input

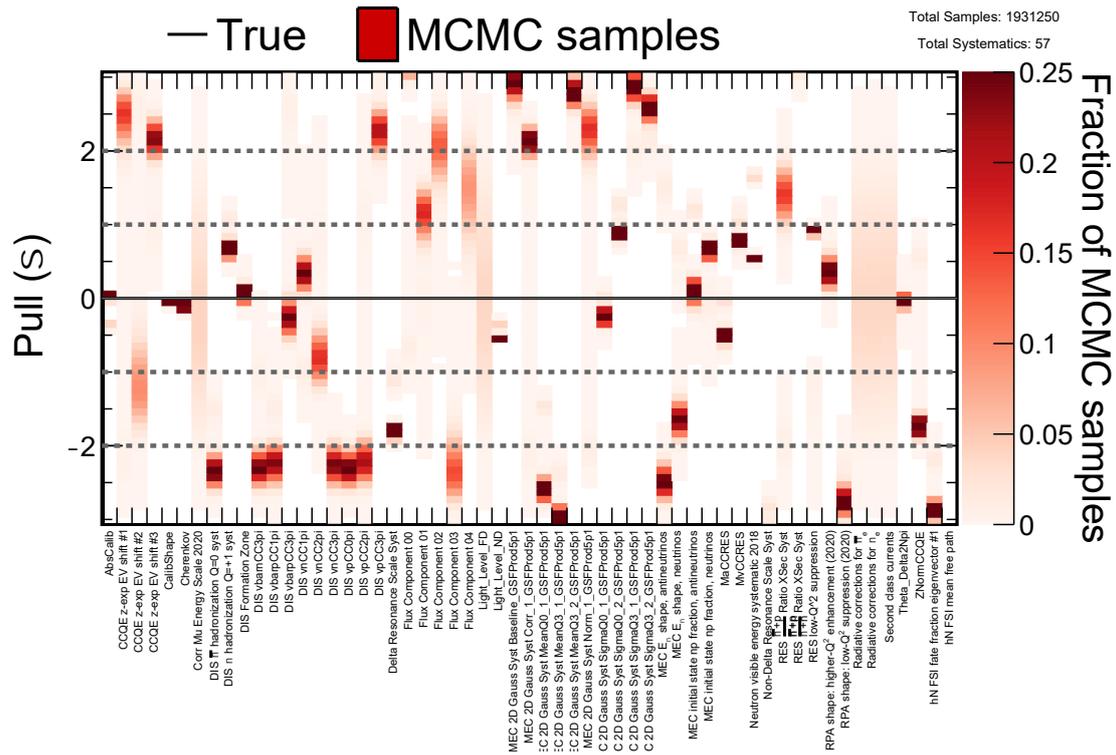


Figure 7.10: Pull Summary plot of the 57 systematic uncertainties from the ND+FD Residual Difference fake data joint fit. Many uncertainties are now well constrained, and we are able to produce many more samples as compared to our initial attempts to fit the ND data. This plot is useful to reference an uncertainty’s probable pull values.

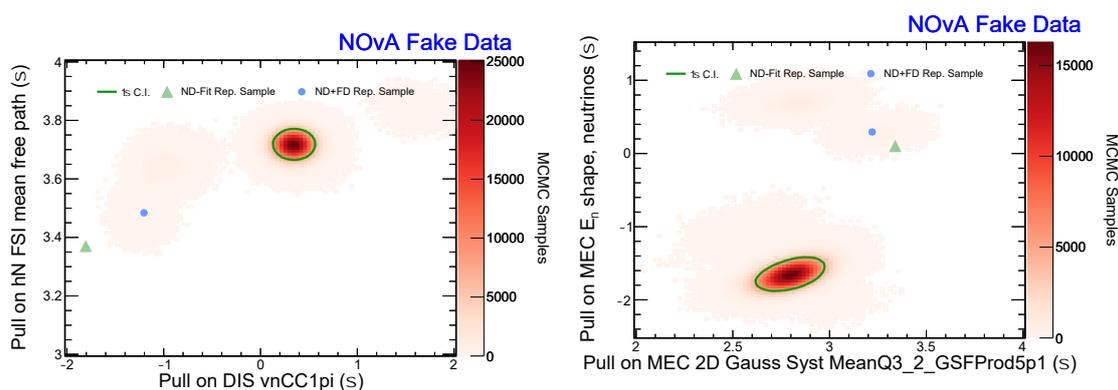
parameters into this fake data are listed in Table 7.1. From this study, a joint ND + FD MCMC fit should reproduce the Asimov point, indicating a robust analysis by successfully reproducing the 2020 NOvA Best Fit.

Pulls Summary

We analyze 1.93 million MCMC samples from 2,500 MCMC chains of the ND+FD joint fit. Figure 7.10 shows the pull summary of the 57 systematic uncertainty parameters constrained in this simultaneous fit. We see most uncertainties are well constrained, but there others that have second, smaller modes. We also see there are several parameters that are still not strongly constrained, such as the radiative corrections systematics, indicated by the lighter pink region. This is a useful reference to identify the modes of a specific uncertainty.

Correlations between uncertainties

As we have seen already from the ND-only fits, we can investigate correlations between pairs of systematics. We investigate a couple pairs briefly here in Figure 7.11. The figure shows the 2D marginal distribution for the hN FSI model's mean free path systematic and the DIS $\nu - n$ 1π systematic (left) and the MEC Double Gaussian parameter (Mean q_3) and the MEC neutrino energy dependence systematic (right).



(a) 2D marginal distribution of the hN FSI model's mean free path systematic and the DIS $\nu - n$ 1π systematic.

(b) 2D marginal distribution of one MEC Double Gaussian parameter (Mean q_3) and the MEC neutrino energy dependence systematic.

Figure 7.11: 2D marginals of systematic pairs from ND+FD Residual Difference MCMC samples. These pairs of systematics suggest they are correlated. The multiple regions of MCMC samples also illustrate the complexity of the posterior. The green triangle shows the ND-only Rep. Sample pull values. The blue circle is the Rep. Sample value from the joint ND+FD fit. The green line represents the Bayesian 1σ credible interval.

These marginals show correlations between the two pairs of systematics; both plots appear to be positively correlated with each other. The right plot is a marginal of two systematics that modulate the same model: MEC. We observe, perhaps unsurprisingly, a positive correlation between these two parameters given they control the same model. However, for the left plot, we see two systematics from two different models: FSI and DIS. This correlation is more revealing about the underlying physics. Large pull values of the MFP systematic translates to a longer mean free path, which means pions within the nucleus are less likely to re-interact (via an FSI process) and are less likely to be observed. As this systematic increases, the pull on

the DIS $\nu - n$ 1π systematic appears to increase too, predicting more CC DIS events from neutrino interactions that produce 1π . We learn MCMC has found high probability samples that seek to produce more neutrino events with 1π . Furthermore, if we recall the discussion from Sec. 6.6, we are already aware of the inter-connection between the FSI MFP parameter and the neutron model. This marginal illustrates how closely the physics models overlap in this analysis.

These marginals also reveal the complexity of the posterior space. Note the green triangle and blue circle in the plots. The green triangle shows the ND-only Rep. Sample pull values from the previous chapter. These pull values are applied to the FD fake data for the Residual Difference ND+FD fit. The blue circle is the Rep. Sample value from the joint ND+FD fit. The pull values from the Rep. Sample between the two fits are similar but they are not identical, further demonstrating how complex the posterior space is. What's more, the densest regions of these distributions (the dark red) is different than the ND and ND+FD Rep. Sample values – well beyond the 1σ credible interval, suggesting the effects of marginalizing over the 60-dimensional posterior space are at play. In addition note the 3-4 regions of MCMC samples in these two plots. For the DIS and MEC neutrino energy systematics negative and positive pull values are probable, which can produce meaningfully different predictions of the ND topologies. The features in these two marginal distributions demonstrate the complexity of the posterior space.

However, we will focus primarily on the oscillation parameters throughout the rest of the chapter, as the previous chapter analyzes the constraints on the uncertainties from the ND data. We will examine the 1D marginal distributions of the oscillation parameters next.

1D marginals of oscillation parameters

We can again marginalize over all other parameters, and samples from both Mass Orderings, to obtain a 1D probability distribution of an oscillation parameter in Figures 7.12, 7.13, and 7.14. The 1D marginal for δ_{CP} (Fig. 7.12) shows the true

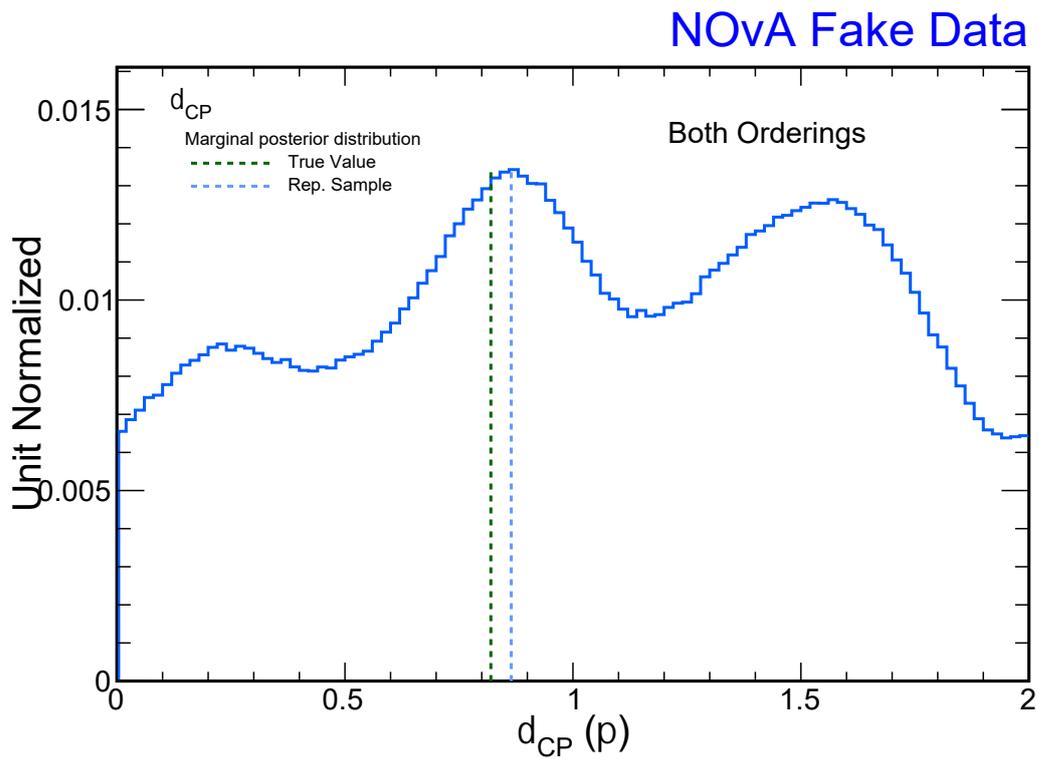


Figure 7.12: 1D marginal distribution for δ_{CP} in units of π , marginalized over both mass orderings. The true value is the green vertical line at $\delta_{CP} = 0.82\pi$. The blue line is the Rep. Sample at $\delta_{CP} = 0.864\pi$. This parameter's mode is also in line with the true and Rep. Sample values. The Normal Ordering and Inverted Ordering marginals are available in the Appendix A.20b & A.20d.

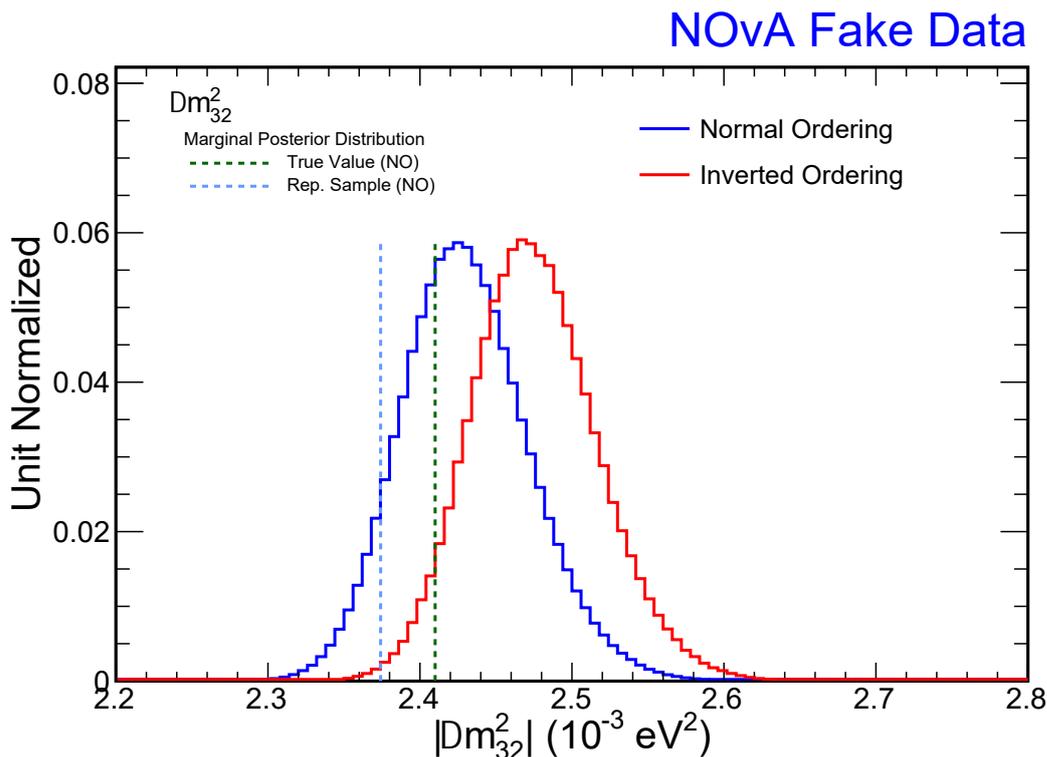


Figure 7.13: 1D marginal of Δm_{32}^2 for each mass ordering. The green line represents the true value at $\Delta m_{32}^2 = 2.41 \times 10^{-3} \text{ eV}^2$ (in NO). The blue line is the Rep. Sample value at $\Delta m_{32}^2 = 2.37 \times 10^{-3} \text{ eV}^2$ (in IO). The mode of the NO marginal distribution is in better agreement with the true value than the IO marginal, and the Rep. Sample pull is further from the mode. This suggests marginalizing over the other parameters can influence the shape of the 1D marginal.

value (green vertical line) is approximately the same as the Rep. Sample (max \mathcal{LL} , in blue). Moreover, the most probable value – indicated by the peak in the marginal distribution – also is in agreement with the true and Rep. Sample values. There are no indications the ND Residual Difference biases δ_{CP} . Moreover, this parameter is constrained by ν_e appearance, and NOvA is statistically limited, evident in the many probable values of δ_{CP} .

Figure 7.13 shows the 1D marginal for Δm_{32}^2 in each mass ordering. Again, we see the Rep. Sample and the true values are similar, while the true value is nearly centered on the mode of the Normal Ordering distribution at $\Delta m_{32}^2 = 2.42 \times 10^{-3} \text{ eV}^2$. Meanwhile, the Inverted Ordering MCMC samples show the mode is located at $\Delta m_{32}^2 = -2.48 \times 10^{-3} \text{ eV}^2$. This discrepancy between the most probable value and the Rep. Sample (max \mathcal{LL}) value suggests the influence of marginalization is

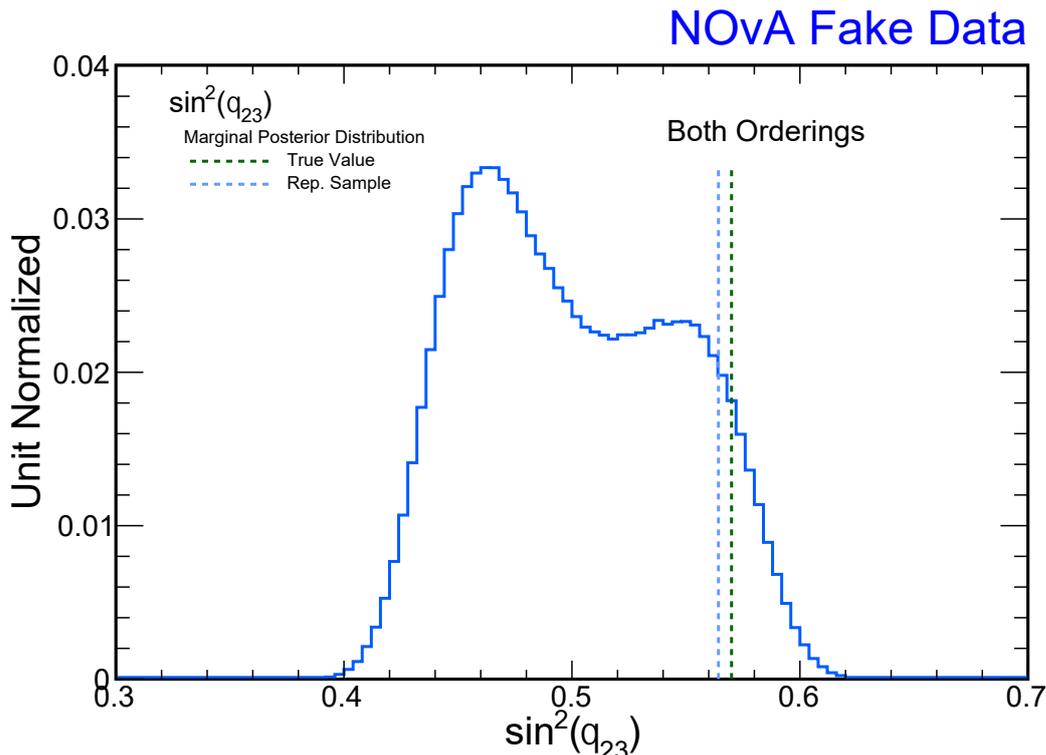


Figure 7.14: 1D marginal of $\sin^2 \theta_{23}$ parameter from joint fit over both mass orderings. The green vertical line is the true value ($\sin^2 \theta_{23} = 0.57$) and the blue vertical line is the Rep. Sample value ($\sin^2 \theta_{23} = 0.564$). The Rep. Sample value is very close to the true value, but the mode of the marginalized distribution is in the lower octant at $\sin^2 \theta_{23} = 0.46$. This represents a bias in this oscillation parameter when using the FD Residual Difference fake data. The Normal Ordering and Inverted Ordering marginals are available in the Appendix A.20a & A.20c.

in effect, as discussed from Figure 5.3.

The last 1D marginal is the parameter $\sin^2 \theta_{23}$ over both mass orderings. The true value (green) and the Rep. Sample (blue) value are very similar and prefer the upper octant. However, the mode of the marginal is located in the lower octant. From Figure 1.8 and NOvA’s 2020 result, we know NOvA has a slight preference for the upper octant. This presents tension as the true value and the max \mathcal{LL} sample both are in the upper octant. This may indicate the effects of marginalization from Section 5.5. Conversely, this may suggest the max likelihood point simply lies in a different region from the high density region of the posterior space.

For reference, the posterior probability distributions for each 1D marginal distribution split by Normal and Inverted Orderings can be found in the Appendix A.20.

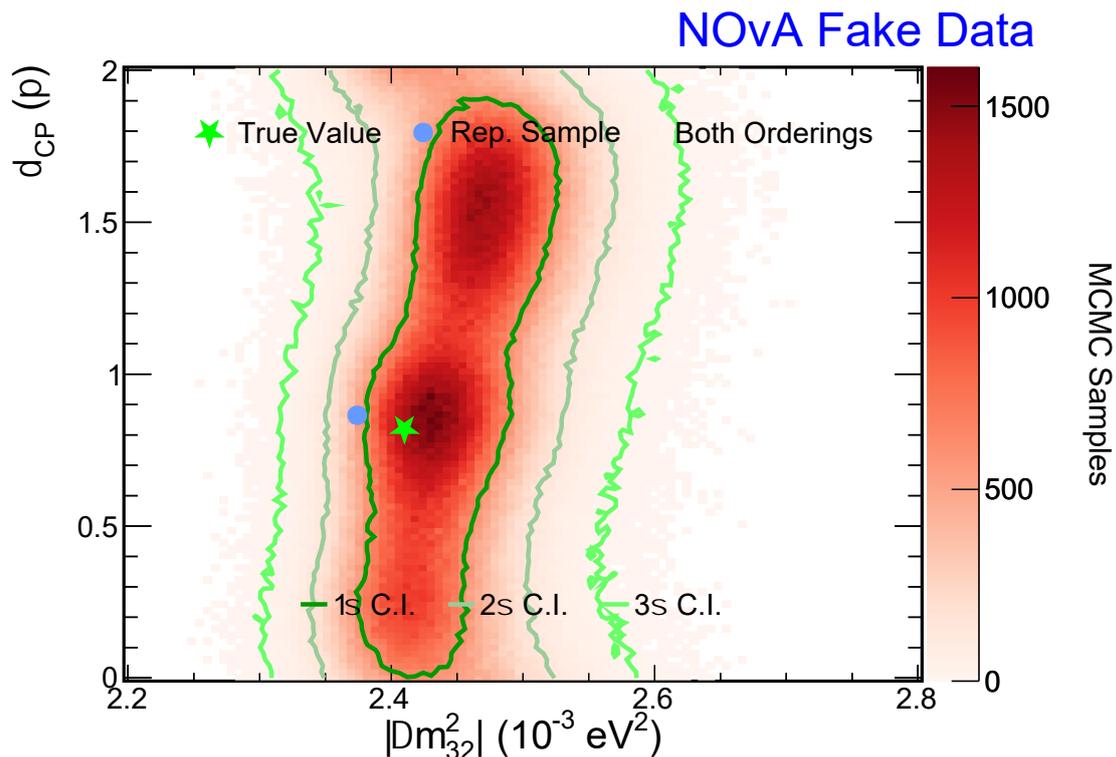


Figure 7.15: δ_{CP} and $|\Delta m_{32}^2|$ 2D marginal distribution over both mass orderings. The green star is the true value from the fake data. The blue is the Rep. Sample value. The *densest* region of MCMC samples (dark red) better agrees with the true value (also known as the mode in the 1D marginals). This Rep. Sample (max $\mathcal{L}\mathcal{L}$ sample) is further from this dense region, suggesting an effect of marginalization. The Normal Ordering and Inverted Ordering marginals are available in the Appendix A.21e & A.21f.

2D marginals of oscillation parameters

We can look at the marginalized oscillation parameters in pairs, as NOvA generally presents these results. Figure 7.15 shows the 2D marginalization of δ_{CP} (in π units) and $|\Delta m_{32}^2|$ over both mass orderings. This 2D distribution further highlights the observations made from the two 1D marginals; the Rep. Sample value is near the true value, however, the densest region of MCMC samples (dark red) is much closer to the true value. This 2D marginal further illustrates the notion that the *densest* region of MCMC samples (i.e. the high probability region) may lie in a different region of the parameter space. Again, this may also suggest the impacts of marginalizing over the $N - 2$ parameters in the posterior space. We can conclusively say, however, the 60-dimensional posterior space is highly complex.

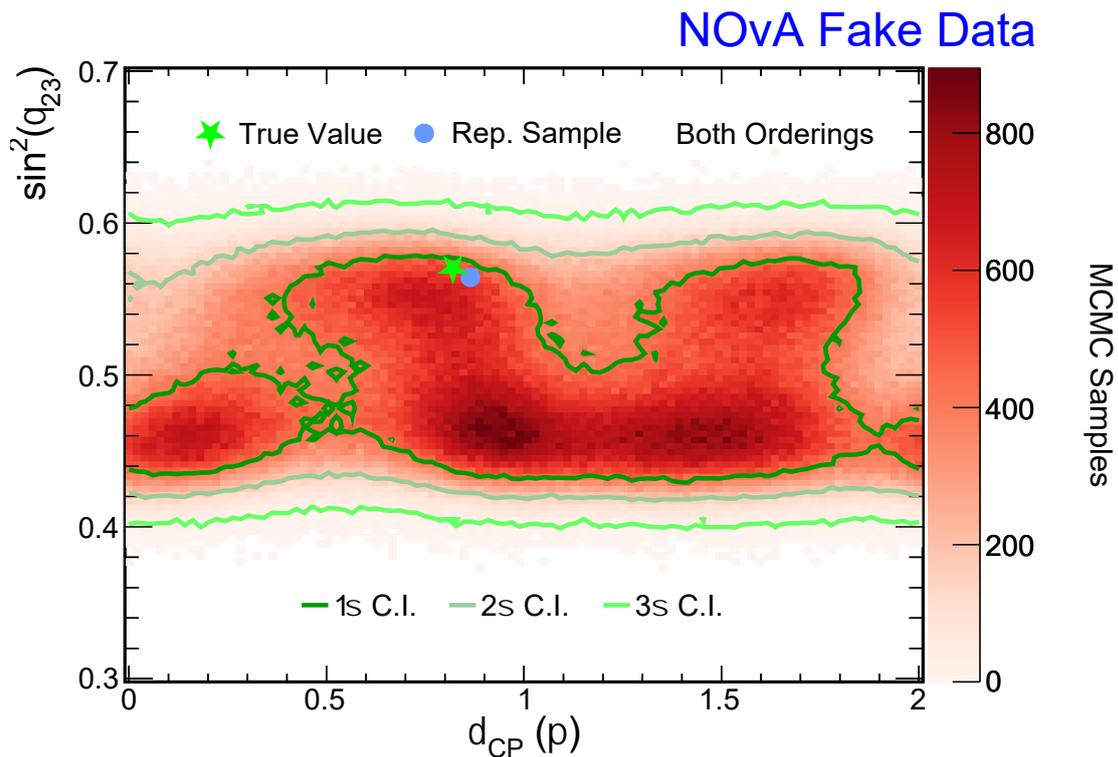


Figure 7.16: δ_{CP} and $\sin^2 \theta_{23}$ 2D marginal over both Mas Orderings. The green star is the true value from the fake data. The blue is the Rep. Sample. The true value and the Rep. Sample value are very near each other in this space, but the densest region is in the lower octant at approximately $\sin^2 \theta_{23} = 0.46$. This is evidence of a bias in the lower octant for this parameter. The Normal Ordering and Inverted Ordering marginals are available in the Appendix A.21b & A.21d.

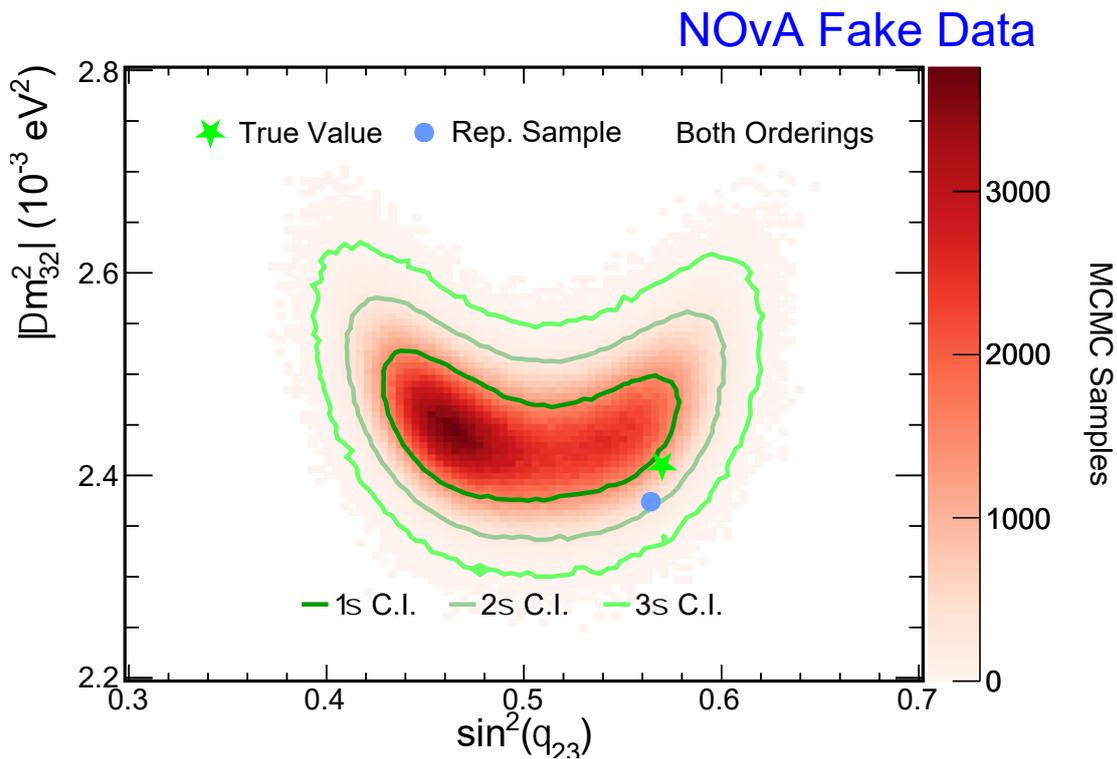


Figure 7.17: $|\Delta m_{32}^2|$ and $\sin^2 \theta_{23}$ 2D marginal over both the Mass Orderings. The bias of the lower octant is most evident here as the darker, denser region of MCMC samples is clustered around $\sin^2 \theta_{23} = 0.46$, while the true value is closer to the Rep. Sample in the higher octant. The 2D marginal of the two parameters in the Normal and Inverted Ordering is shown in Fig. A.21a & A.21c.

Figures 7.16 and 7.17 show the 2D marginals for δ_{CP} vs. $\sin^2 \theta_{23}$ and $|\Delta m_{32}^2|$ vs. $\sin^2 \theta_{23}$ over both mass orderings, respectively. These two 2D marginals show the Rep. Sample value is very similar to the true value, which from our ND-only metric, is classified as a robust fit. However, when we look at the distribution of the MCMC samples in these two plots we see the densest region is in the lower octant of $\sin^2 \theta_{23}$. In Fig. 7.17 we see the densest region of MCMC samples are located in the lower octant, while the Rep. Sample value is within the upper octant. These two plots further illustrate the complexity of the posterior suggesting: the effects of marginalization are present and the max likelihood sample does not neighbor the high density region in the posterior.

Rep. Sample Prediction

Now that we have seen increasing evidence of the discrepancy between the Rep. sample (max \mathcal{LL} sample) and the *densest* region of the parameter space, we will examine the quality of the agreement between the fake data and the Rep. Sample in the FD ν_μ and ν_e samples. Figure 7.18 and 7.19 show the FD FHC and RHC ν_μ distributions in the four quantiles plotted with the 1σ error band before the fit. We see the Rep. Sample prediction does not produce strong agreement, especially in Quantile 1, which has the most amount of muon energy – the parameter we discovered is poorly constrained from the ND data. The agreement, however, does improve with increasing hadronic energy fraction, which we would expect given the predictions are constrained primarily from the ND ($Reco |\vec{q}|, E_{had}^{vis}$) topological samples, where the hadronic energy is well constrained. Despite this, we conclude that the Rep. Sample ν_μ prediction does not agree well with the FD fake data.

Figure 7.20 shows the FD ν_e samples for FHC and RHC. We see, perhaps unsurprisingly, the Rep. Sample prediction is not significantly different from the fake data nor the nominal prediction. This coincides with the broad 1D marginal distribution for δ_{CP} , indicating NOvA does not constrain this parameter strongly due to reduced statistics.

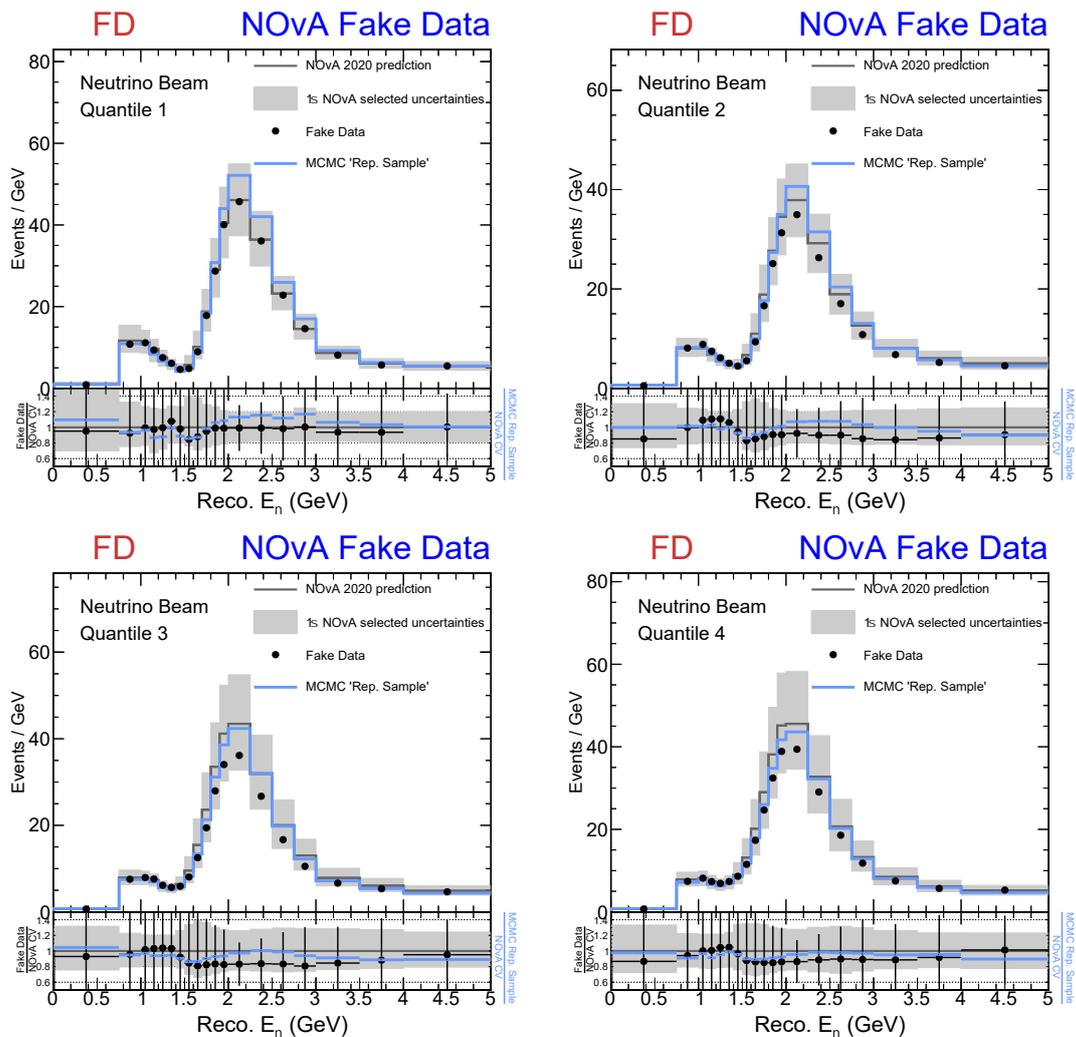


Figure 7.18: FHC FD ν_μ Representative Sample predictions (blue) with a priori 1σ error of the 57 systematic uncertainties in the grey band. The Rep. Sample does not produce strong agreement with the FD FHC ν_μ fake data (black points). It is noticeably worse in Quantile 1, especially at the peak of the distribution.

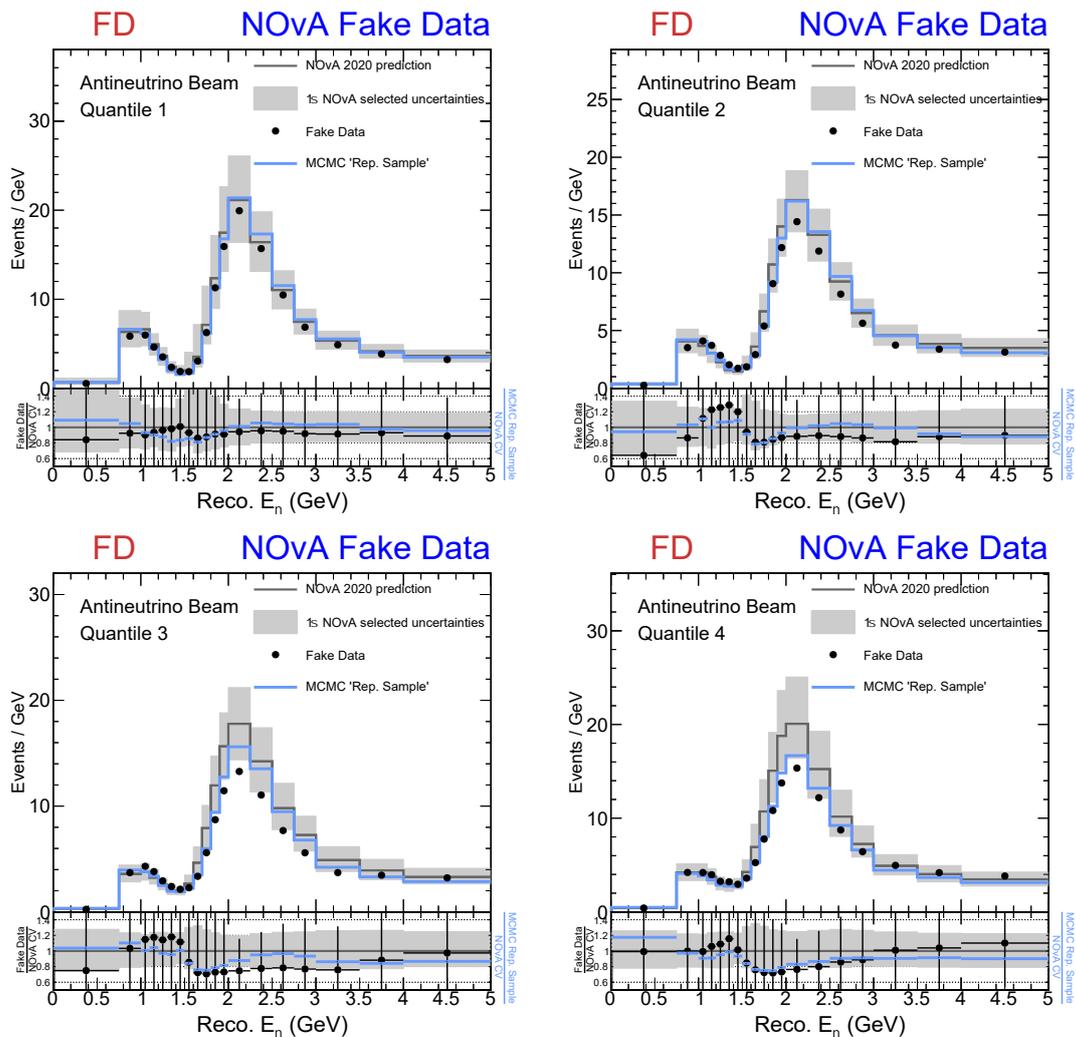


Figure 7.19: RHC FD ν_μ Representative Sample predictions (blue) with a priori 1σ error of the 57 systematic uncertainties in the grey band. The Rep. Sample does not produce strong agreement with the FD RHC ν_μ fake data, especially in the most populated bins of Quantile 1.

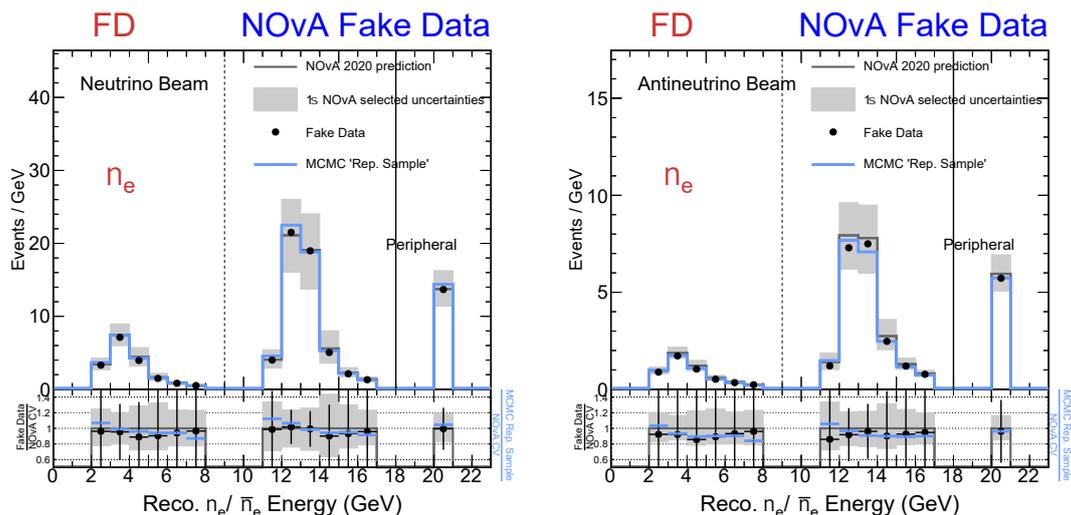


Figure 7.20: FD ν_e Representative Sample predictions for FHC and RHC (blue) with a priori 1σ error of the 57 systematic uncertainties in the grey band. The Rep. Sample prediction is not qualitatively different from the nominal prediction (grey) nor fake data (black points), indicating this parameter is not strongly constrained due to low ν_e statistics.

Predictions of Random MCMC Samples

Thus far we have seen the values of the oscillation parameters in the highest probability (i.e. modes for 1D marginals and the high density regions for 2D marginals) are not the same as the values from the Rep. Sample, especially for $\sin^2\theta_{23}$. We have also seen in the previous section the FD ν_μ and ν_e Rep. Sample prediction does not produce strong agreement to the fake data, in particular for the Quantile 1 distributions. This motivates us to examine predictions for the *densest* region of MCMC samples from the 1D and 2D marginals. In this Subsection we plot 200 *random* samples from the 1.9 million samples collected. Because we are selecting samples randomly to plot, we select samples proportional to the probability and therefore will plot predictions from the denser regions of the 1D and 2D marginal space.

Figures 7.21 and 7.22 show the FD ν_μ FHC and RHC distributions in the quantiles of predictions from 200 randomly drawn samples. Again, we draw these samples randomly and therefore are more likely to select samples from the denser regions of the marginal distributions. We see these 200 random predictions produce a narrow band of predictions that encapsulate the fake data. This narrow band produces

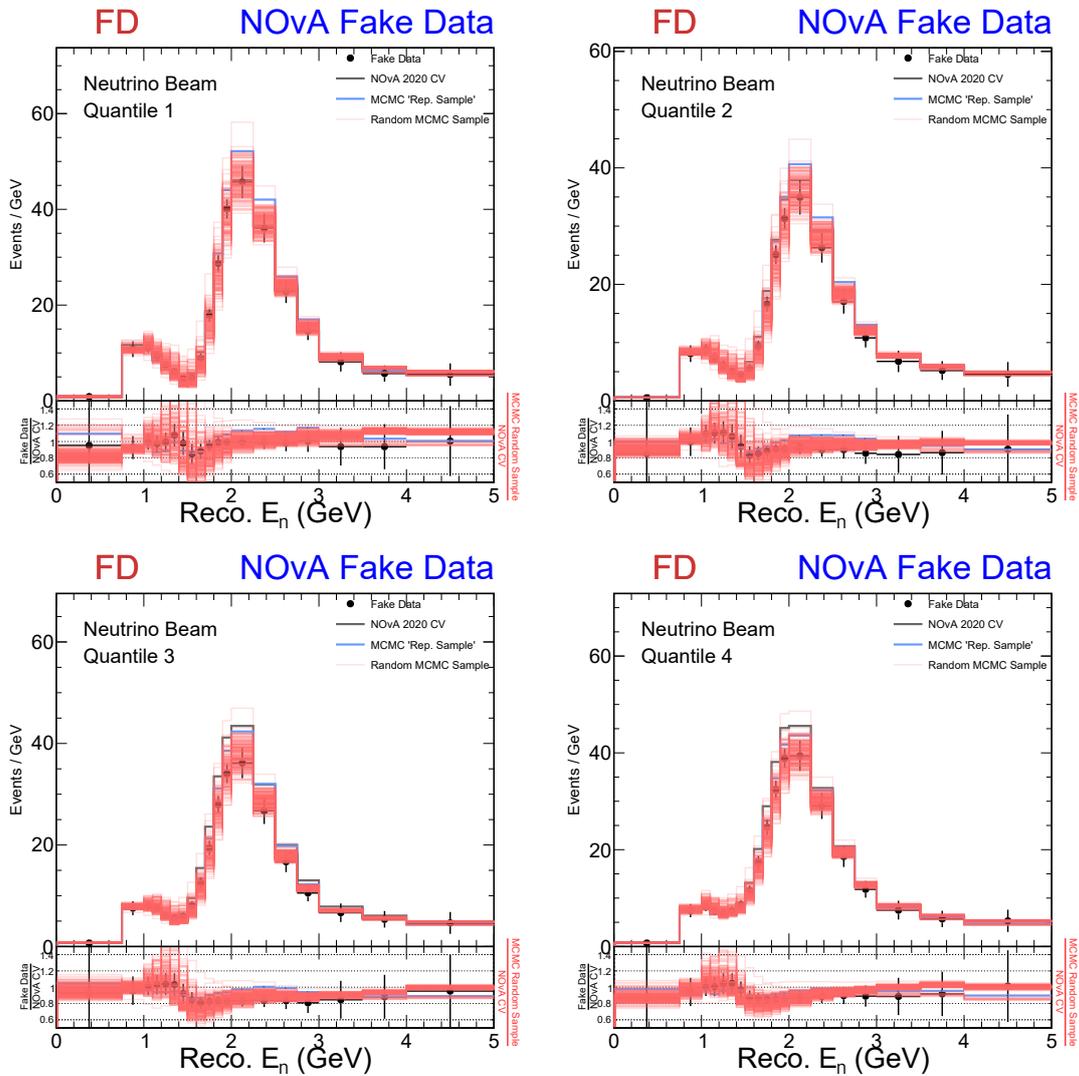


Figure 7.21: FD FHC ν_μ predictions from 200 random MCMC samples drawn in red. We see the red predictions encapsulate the fake data. Many of the predictions from the random samples appear in better agreement with the fake data than the Rep. Sample prediction.

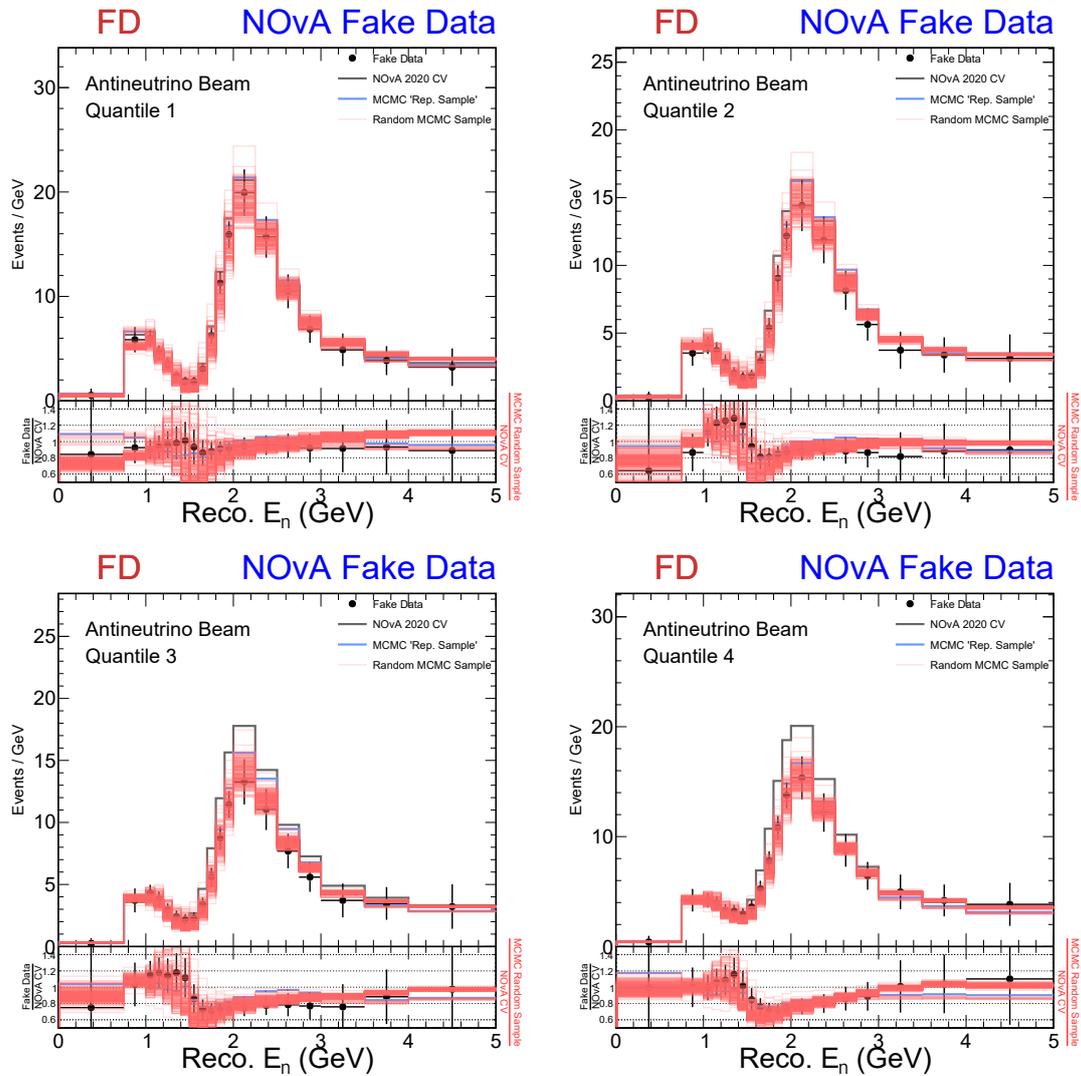


Figure 7.22: FD RHC ν_μ predictions from 200 random MCMC samples drawn in red. We see the red predictions encapsulate the fake data. Many of the predictions from the random samples appear in better agreement with the fake data than the Rep. Sample prediction.

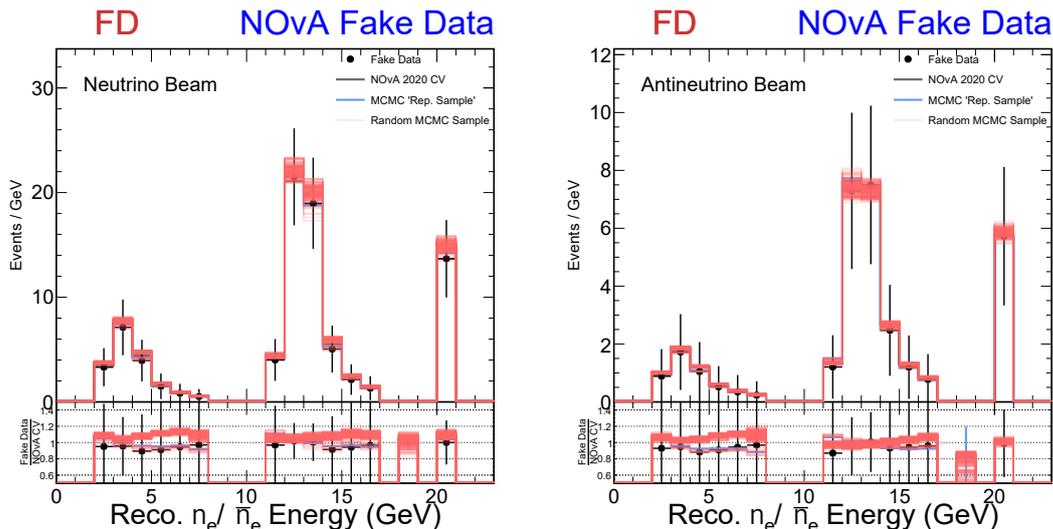


Figure 7.23: FD ν_e predictions from 200 random MCMC samples drawn in red in the NOvA analysis binning (low PID, high PID, and Peripheral). We see the narrow band of predictions agree with the fake data and the Rep. Sample prediction.

a much stronger agreement to the Quantile 1 fake data as well, where the Rep. Sample prediction did not agree. This is evidence the highest probability regions of the parameter space may be more effective in describing a “best fit” to the true oscillation parameter values, better than the Rep. Sample prediction.

This observation is true for the FD ν_e samples as well, Figure 7.23. The narrow band from 200 random samples agrees with the fake data. However, in the Low PID bins of the Core sample, we see that there are two sets of red bands (visible in the ratio plot). This is likely because the δ_{CP} parameter is not well constrained (seen in the 1D marginal), so 200 random values of this parameter likely produces a wide range of values of δ_{CP} .

The primary conclusion from examining the Rep. Sample prediction and the 200 random sample predictions is the posterior space is very complex. As we have seen in the ν_μ and ν_e samples, the Rep. Sample and 200 random sample predictions provide two different metrics to assess MCMC samples; they both constrain the oscillation parameters and may have different implications for how to interpret MCMC samples.

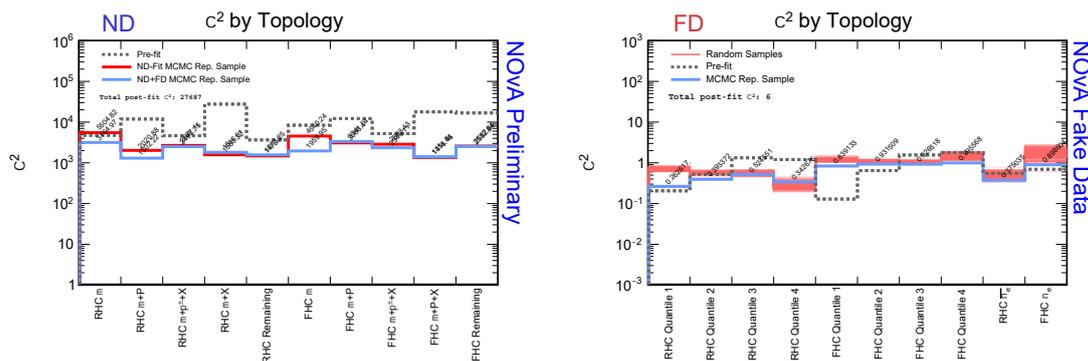
(a) The χ^2 for each ND topology.(b) The χ^2 for each FD selection.

Figure 7.24: **Left:** The χ^2 is consistent between the ND-only and ND+FD joint fit, showing stability in the result. The ND+FD χ^2 is improved slightly in the ND μ topologies because the muon energy scale systematic is constrained only by the FD ν_μ samples; it was pulled to positive value thereby indirectly improving these two topologies. **Right:** The Rep. Sample, in this metric, produces the “best fit” when compared to the 200 random MCMC samples. However, in Quantile 1 the Rep. Sample χ^2 is still worse than the nominal prediction.

χ^2 metric

At this stage we have looked at the results of the joint fit from a qualitative perspective. Figure 7.24 plots the χ^2 for the ND topologies (left) and the FD samples (right). In the left plot, the χ^2 values of the ND topologies from the ND+FD joint fit (blue) are shown, and the ND-only fit χ^2 s from Chapter 6 (red) are shown for reference. We see the results are consistent between the fits, suggesting the Rep. Sample is likely the same point in the posterior space from these two fits, indicating stability in the result. However, the two μ topologies are most discrepant between these two MCMC fits. This is largely due to the impact of the `kCorrMuEScaleSyst2020` systematic. In the ND-only fit we removed the parameter from the results, as we saw it adversely affected the $Reco E_\nu$ prediction in the ND quantiles because E_μ is not well constrained from our ND fit variable $Reco |\vec{q}|$. However, in the ND+FD joint fit, the `kCorrMuEScaleSyst2020` parameter can be constrained by the FD ν_μ samples. It was constrained by the FD fake data; the Rep. Sample pull value for `kCorrMuEScaleSyst2020` = $+1.2\sigma$. This means when we apply the Rep. Sample pull values to all ND and FD predictions, we see the χ^2 for the ND μ topologies have improved. This effect can be seen further in the 1D projections onto the $Reco |\vec{q}|$

variable in the Appendix A.18 & A.19.

The right plot shows the χ^2 for the FD samples for the data-Rep. Sample (blue) and the 200 random MCMC samples (red). In this χ^2 metric we observe that the Rep. Sample produces the best χ^2 agreement over the 200 random samples, of which some are from the densest region in the parameter space. This appears contradictory to what we see in the FD ν_μ and ν_e Rep. Sample predictions, however, this may be due to the effects of marginalization. Again, we see different results in the agreement to the fake data when analyzing the max \mathcal{LL} sample and the highest marginal probability samples.

Correlations between oscillation parameter & uncertainty

With marginalization we can look into any correlations that may exist between oscillation parameter and systematic uncertainty. A few are selected here in Figure 7.25. In the left two plots are 2D marginals of the muon energy scale systematic and CCQE normalization systematic paired with Δm_{32}^2 . In the top left plot we see a slight anti-correlation between these two parameters. This is an interesting conclusion as many of the pairs of parameters we investigated do not have an obvious relationship. The bottom left plot shows a high density region of the CCQE normalization parameter at -1.8σ , but the Rep. Sample values from the ND and ND+FD MCMC fits are near -0.5σ , far from this dense region. This very likely an illustration of the effects of marginalizing a multi-dimensional parameter space down to two parameters. In the top right plot of the RES Δ scale systematic created in Subsection 6.5.2 and $\sin^2 \theta_{23}$, a similar conclusion is true. The dense region and Rep. Sample values are in two separate, distinct regions of the parameter space. In the bottom right plot of the MEC Double Gaussian `SigmaQ3_1` systematic, similar to the top right plot, we see densest regions are for from the ND+FD and ND-only Rep. Sample, however, these values are extremely similar suggesting stability in the two fits.

The ability to marginalize over any pair of parameters allows us to learn about

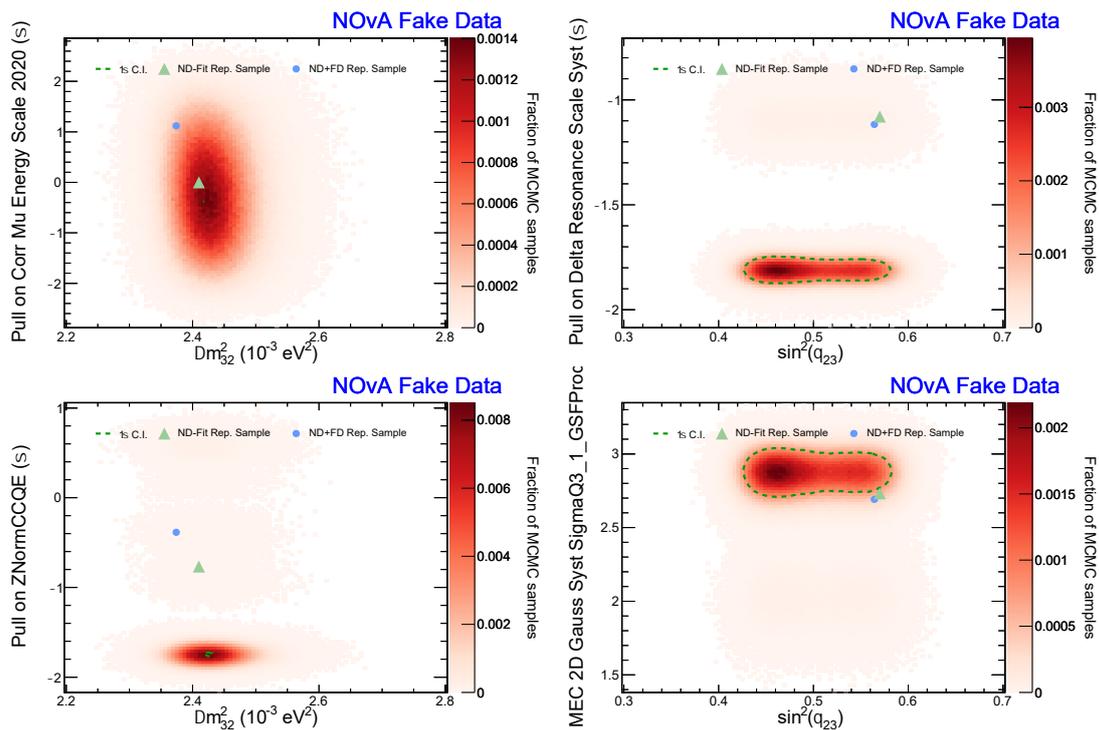


Figure 7.25: 2D marginal distributions of various uncertainties and oscillation parameters. The blue point is the ND+FD Rep. Sample value. The green triangle represents the ND-only Rep. Sample pull value for the *uncertainty*, while for the oscillation parameters it is the true value from the FD fake data in the ND+FD joint fit. In these pairs of parameters there are no obvious correlations between uncertainty and oscillation parameter. However, there is a slight anti-correlation between the muon energy scale systematic and Δm_{32}^2 .

correlations between them. Again, there are countless numbers of pairs of 2D distributions that can be made from this analysis, however, we report most oscillation-systematic parameter pairs do not have any significant correlations. This is a new level of analysis that can be gained from MCMC sampling brought forward in this thesis.

7.5 Future Improvements

The work to produce a ND-only fit to the NOvA data and a ND+FD joint fit of the ND data & FD fake data has taken a large amount of effort. Throughout this effort, we have learned many things about NOvA’s interaction modeling, oscillation parameter constraints, and MCMC sampling. From what we have learned, we have identified some future developments that would improve the robustness of the analysis. Some possible improvements to the fitting are listed here grouped together.

7.5.1 ND Improvements

- **Incorporate E_μ as a third variable in the ND fit.** We have learned that we do not constrain the muon systematic uncertainties well with $Reco |\vec{q}|$. This – or some other muon kinematic variable – should be included in future analysis iterations. Moreover, muon kinematics are what NOvA measures best, and any muon kinematic variable would ideally be well constrained from NOvA ND data. With respect to the MCMC sampling **quantity**, introducing a third variable for the ND fit would slow the rate of MCMC sampling, and likely require further adjustments to the input into MCMC, such as additional phase space cuts on $Reco |\vec{q}|$ and E_{had}^{vis} .
- **An improved FSI systematic uncertainty set.** There is strong evidence in NOvA the current tune of the hN FSI model is insufficient. We see the mean free path uncertainty pulled to values larger than $+3.5\sigma$ for the ND

data and ND+FD joint fits. This pull effectively removes NOvA tune to external pion scattering data. This indicates a problem with the model. One option is to assess other FSI models accessible in GENIE, or to update the hN tune to incorporate nucleon intranuclear interactions, though this is a second order effect to pion re-interactions. We believe this may also be associated with an insufficient neutron model/uncertainty in NOvA; the large MFP pull reduces pion activity, and by extension, fewer neutrons are produced. NOvA's simulation over-predicts neutrons and could help to explain the large pull. Investigating neutron-specific variables with these large FSI pull values applied to the NOvA prediction could reveal more about the connection between these two physics effects.

- **Introduce a ND ν_e constraint.** The ND fit does not explicitly constrain ν_e events. To create additional topologies of ν_e events would likely improve the constraint of several uncertainties and δ_{CP} . This would also drastically slow the MCMC sampling and reduce the **quantity** of samples due to the added bins. Moreover, a choice of variable(s) to use for constraining the ND ν_e sample(s) would be necessary. An improved constraint on the ν_e sample may also resolve the octant bias in $\sin^2 \theta_{23}$.

7.5.2 ND+FD Improvements

Following the analysis of the joint fit in the previous Section, there are several direct items this analysis could followup on:

- **Construct an “out-of-model” ND and FD fake data set with known inputs.** This would allow a more focused constraint on the sensitivity/bias of the oscillation parameters. The challenge here is to create a fake data set that resembles the ND Residual Difference as best as possible with known distortions to the uncertainty model. This would allow the analyzer to replicate a type of Residual Difference fit, while understanding what models are being stressed in the fit. (We acknowledge that to create a fake data set with

known inputs that exactly matches the Residual Difference is a nontrivial task, discussed in Sec. 7.1.)

- **Compare joint probability distributions between the ND-only and ND+FD fit for the systematics.** Due to the overwhelming statistical power of the ND data, we believe the incorporation of the FD data would only modestly changes these distributions. This would be useful to confirm or deny.
- **Investigate another high probability point from the posterior space.** For example, how does selecting the next-highest max \mathcal{LL} sample change the result on the bias of $\sin^2 \theta_{23}$? How would data-MC agreement change if using the highest probability point from a marginalized distribution instead? These are simply two possibilities that can be explored.

One of the primary items to followup on is the bias in the octant of $\sin^2 \theta_{23}$. In the marginalized space of this parameter, the densest region of MCMC samples show a bias towards the lower octant of $\sin^2 \theta_{23}$, despite the true value in the upper octant. We are aware this sensitivity to octant is driven by ν_e events. Some additional followup items include a robust check of the construction of the FD ν_e fake data, including the oscillated and unoscillated components. A second followup would be to perform a FD ν_μ -only Residual fit. A ν_μ -only fit is insufficient to break the octant degeneracy, and therefore we should find the marginalized distribution to be symmetric in $\sin^2 \theta_{23}$. Investigating this bias is beyond the scope of this thesis, however such an investigation is not limited to these suggestions.

7.5.3 MCMC Improvements

From the MCMC perspective, there are a number of ways that MCMC sampling can be improved. We understand the objective is to maximize the number of MCMC samples and explore the posterior as efficiently as possible. We list some suggestions here:

- **Introduce different priors for different systematics.** For our uncertainty

model we used the same Gaussian-like “Custom” prior. However, there are parameters that could be limited further and other parameters that have special cases.

- **Introduce a correlation matrix of the systematic parameters.** This could originate from a prior MCMC fit or from MCMC warmup alone. This information would likely drastically improve the number of samples produced from a fit as MCMC would step through the posterior more efficiently, rather than using CPU time to complete MCMC warmup.

This is by no means exhaustive of how MCMC sampling can be improved, but these are ideas that arose during this analysis.

7.6 Conclusions

This analysis has two distinct components: the **ND-only fit** focusing on interaction modeling and the **ND+FD joint fit** to constrain the oscillation parameters. Each of these components are required to tell the full story of constraining the oscillation and interaction parameters within NOvA. This work uses Bayesian inference to estimate these parameters via the Hamiltonian Monte Carlo algorithm within Markov Chain Monte Carlo to explore the posterior parameter space. This thesis required a significant amount of development to achieve these results, and establishes how to best maximize NOvA’s information from its two detectors simultaneously.

ND conclusions

Much development has been made in the goal of fitting the ND data. Numerous fake data fits were performed to demonstrate MCMC can effectively constrain NOvA’s uncertainty model. Moreover, we have learned there are systematics within NOvA’s uncertainty model that cannot be constrained and have removed many of these from MCMC sampling. Much work has been made to improve the quality and quantity of MCMC samples, such as: adjusting the ND topological binning, phase

space cuts, and omitting systematics that cannot be constrained from NOvA's ND data. In the efforts of the ND fitting, much was learned about the quality of neutrino interaction modeling in NOvA's simulation. To this end, additional degrees of freedom relating to the RES and DIS cross section models were introduced to help improve the data-MC agreement. These efforts were successful, but a lingering discrepancy still remains in the form of the ND Residual Difference. Outlined below are three important findings relative to interaction modeling and the ND fitting:

- **low- Q^2 events are over-predicted.** These are predominately RES and DIS events that appear in the ND topologies along the $Reco |\vec{q}| = E_{had}^{vis}$ kinematic boundary and prominently in the $\mu + \pi^\pm + X$ and $\mu + (P+)X$ topologies.
- **The MEC Double Gaussian uncertainty model is degenerate** in describing a NOvA fake data or data set. This was made clear from NOvA ND fake data fits that show different pull values in a number of the Double Gaussian parameters produce the same prediction. This is likely due to the 13 parameters that weight MEC interactions in a region of phase space that largely overlaps with the QE model.
- **NOvA's hN FSI tune is insufficient.** The large pull value of the mean free path systematic (from both the ND and ND+FD fits) effectively removes the MFP tune we applied to obtain agreement between the hN model and external pion scattering data. This particular systematic significantly changes the ND topological predictions, therefore improving this tune should improve the agreement to the ND data. In addition, we believe this large pull is associated with the neutron systematic as well, where the large MFP pull reduces neutron activity to improve the ND agreement.

Taking a step back, the ND fitting is an attempt to take seriously the NOvA ND data and interaction models in the context of constraining oscillation parameters in NOvA.

ND + FD conclusions

From the perspective of fitting the FD predictions, a full ND+FD joint fit has been developed that constrains 57 NOvA systematic uncertainties and 3 oscillation parameters with 1.9 million MCMC samples. What's more, a procedure for testing and investigating the ND Residual Difference with MCMC has been developed, which provides consideration of the ND and FD distributions all at once; an entirely new means of interrogating the NOvA data.

For this analysis, we opted to assess the impact of the Residual Difference on the oscillation parameters by use of a fake data set. Specifically, we selected to use a FD fake data set with *no known inputs*, in the form of the ND Residual Difference. The advantage to this choice is to capture the ND Residual Difference exactly in our FD fake data set. This is simply one choice of assessment.

From this analysis we have learned the 60-dimensional posterior space is highly complex. This has implications for how one should interpret the results. We report an apparent inconsistency in the marginal distributions of Rep. Sample (max \mathcal{LL}) values and the high probability region values for the oscillation parameters. This inconsistency, however, simply requires careful consideration to analyze the MCMC samples. We outline two important conclusions from this joint MCMC sampling:

- **The Rep. Sample (max \mathcal{LL}) value may not be identical to the highest probability region in the posterior.** The posterior is a complex space. There is no mathematical guarantee the Rep. Sample point lies near the region of high probability samples in the posterior (if the Rep. Sample point was a Dirac delta function, for example). To this end, we plot 200 random MCMC samples in Fig. 7.21 – which are selected proportional to the high probability region in the posterior – and find the band of these predictions are qualitatively different than the Rep. Sample prediction, suggesting the Rep. Sample point may not be identical to the highest density region in our posterior space. A useful tool to learn more about the shape of the posterior would be the use of a clustering algorithm. Another useful development would be

tools for visualizing these high-dimensionality spaces. While we are unaware of any such tools at present, we note the MCMC samples contain all the information necessary for such a visualization, as they capture the entire joint probability distribution. Both of these tools would allow us to draw more concrete conclusions and shed light on the shape of the posterior, for example, to determine if the Rep. Sample point is in fact near the highest probability regions in the posterior.

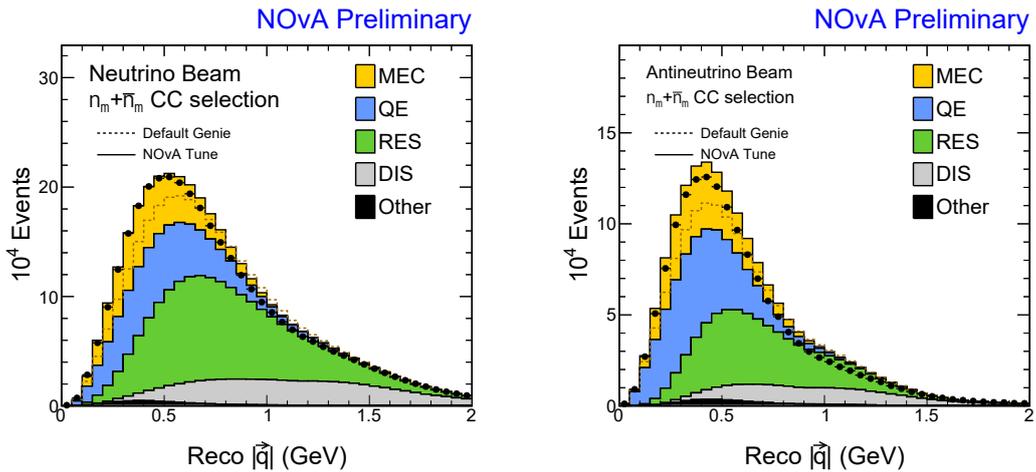
- **The Rep. Sample ($\max \mathcal{LL}$) value may not be identical to the highest probability region in a *marginalized* distribution.** We first see this effect from Figure 5.3, where marginalizing over arbitrary parameters may move the region of high probability in the lower dimension space. In our analysis, Figures 7.17, 7.11, and 7.25 represent the start of this “deep dive”. They clearly show there are distinct high-probability regions, at least in many of these 2-D comparisons, with the max likelihood value lying in the region with overall lower probability. We note here one of the weaknesses of the Residual fake data fit – there are no “true values” for the systematics, just the max likelihood values from the ND fit. In principle the FD data (although small) is adding additional information and could refine/improve the determination of these parameters.

To place this work in the larger context, this analysis sets the ground work for the next iteration of joint fit analyses in NOvA. There are numerous opportunities to strengthen and improve the fit outlined in this thesis. NOvA is well positioned to produce a more robust joint analysis in the next round of joint fitting. Looking further into the future, we hope the Deep Underground Neutrino Experiment can maximize the utility of this analysis for a joint fit analysis within the next generation of oscillation experiments.

Appendix A

Appendix

A.1 MEC Tune in NOvA



(a) NOvA ND data and MC predictions of neutrinos in variable $Reco |\vec{q}|$. (b) NOvA ND data and MC predictions of antineutrinos in variable $Reco |\vec{q}|$.

Figure A.1: $Reco |\vec{q}|$ in the NOvA ND for the data (black dots) and the MC prediction, broken down by true neutrino interaction. Before the MEC tuning, the MC predicts a small amount of MEC interactions – up to the dashed line (“Default GENIE”). After tuning in true $(|\vec{q}|, q_0)$, (“NOvA Tune”, solid), the MEC tuned prediction produces a stronger agreement with the data.

A.2 ND Topological Samples

A.2.1 ND MC 1D Projections onto $Reco |\vec{q}|$

A.2.2 ND Representative Sample 1D Projections onto $Reco |\vec{q}|$
from Fake Data Fit

A.2.3 ND Systematics Ranked by χ^2 in ND RHC Topologies

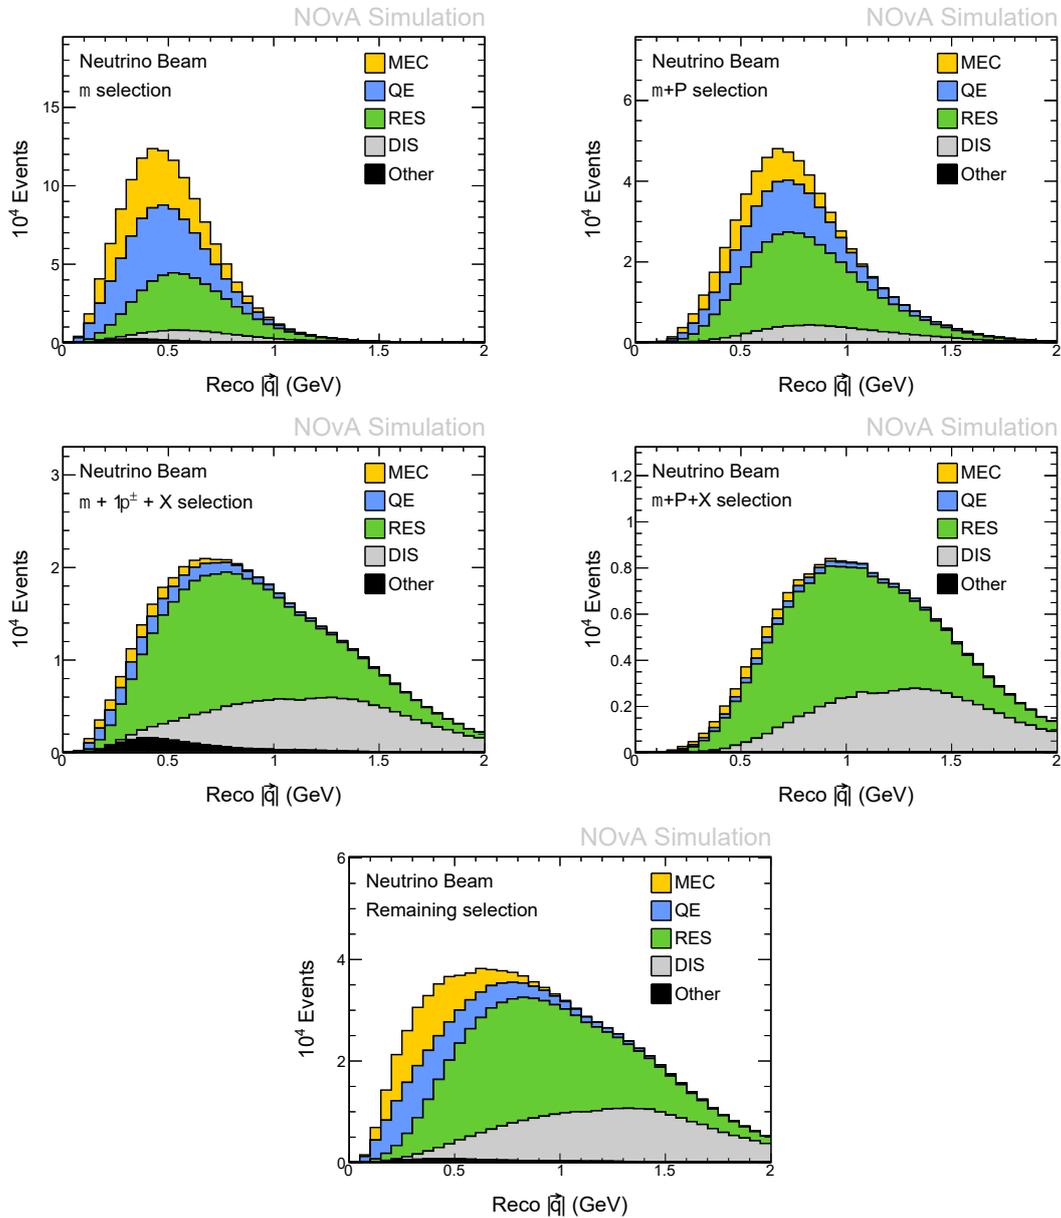


Figure A.2: Projection of the 2D $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space onto the $Reco |\vec{q}|$ axis, projecting out E_{had}^{vis} . These plots of the FHC topologies break down neutrino events by true interaction type in the simulation. Complementary plots to Figure 4.8.

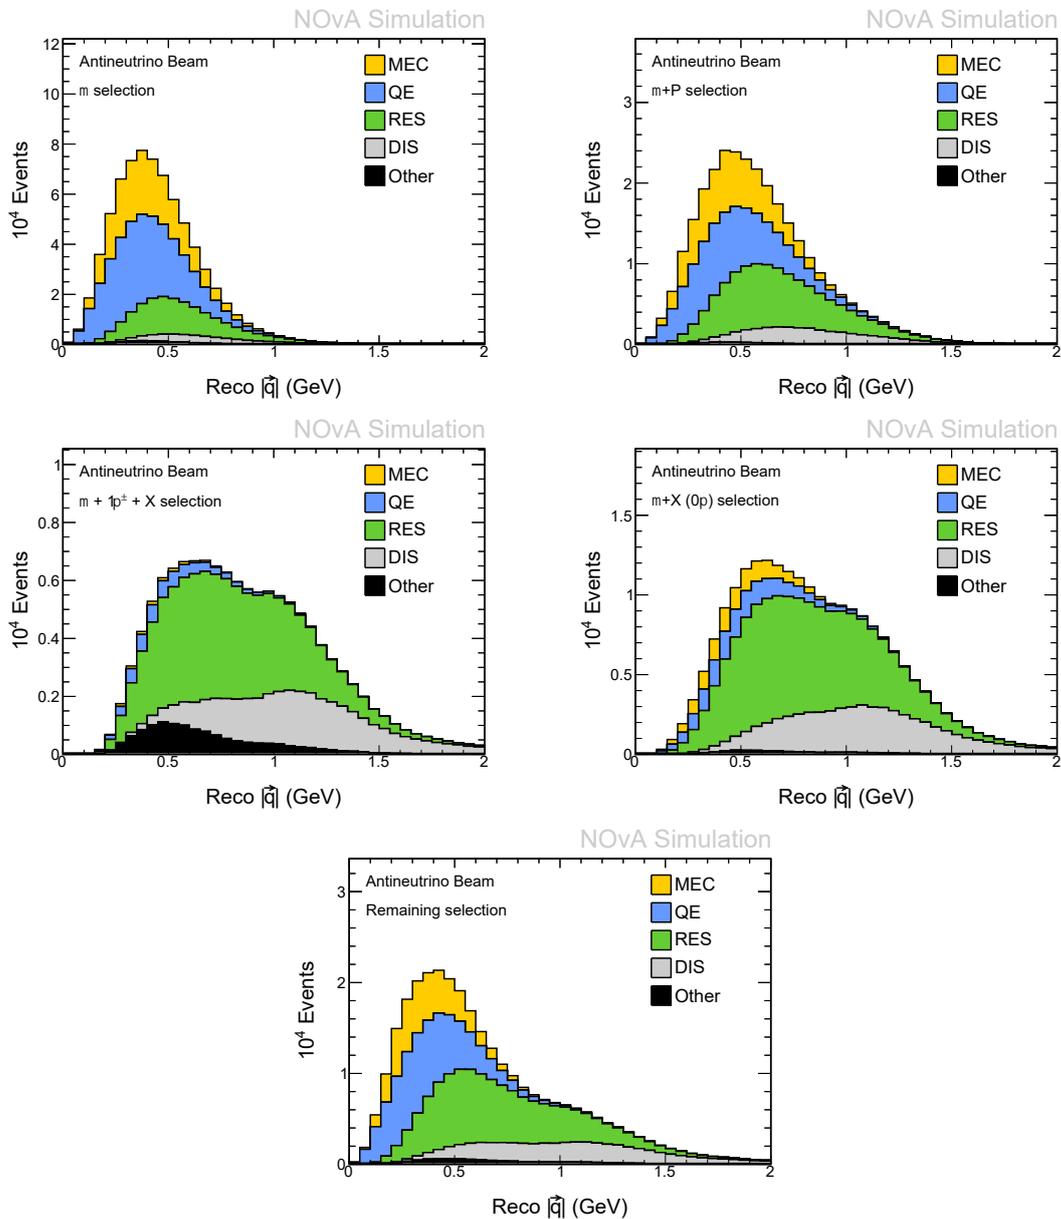


Figure A.3: Projection of the 2D $(E_{had}^{vis}, Reco |\vec{q}|)$ phase space onto the $Reco |\vec{q}|$ axis, projecting out E_{had}^{vis} . These plots of the the RHC topologies break down neutrino events by true interaction type in the simulation. Complementary plots to Figure 4.9

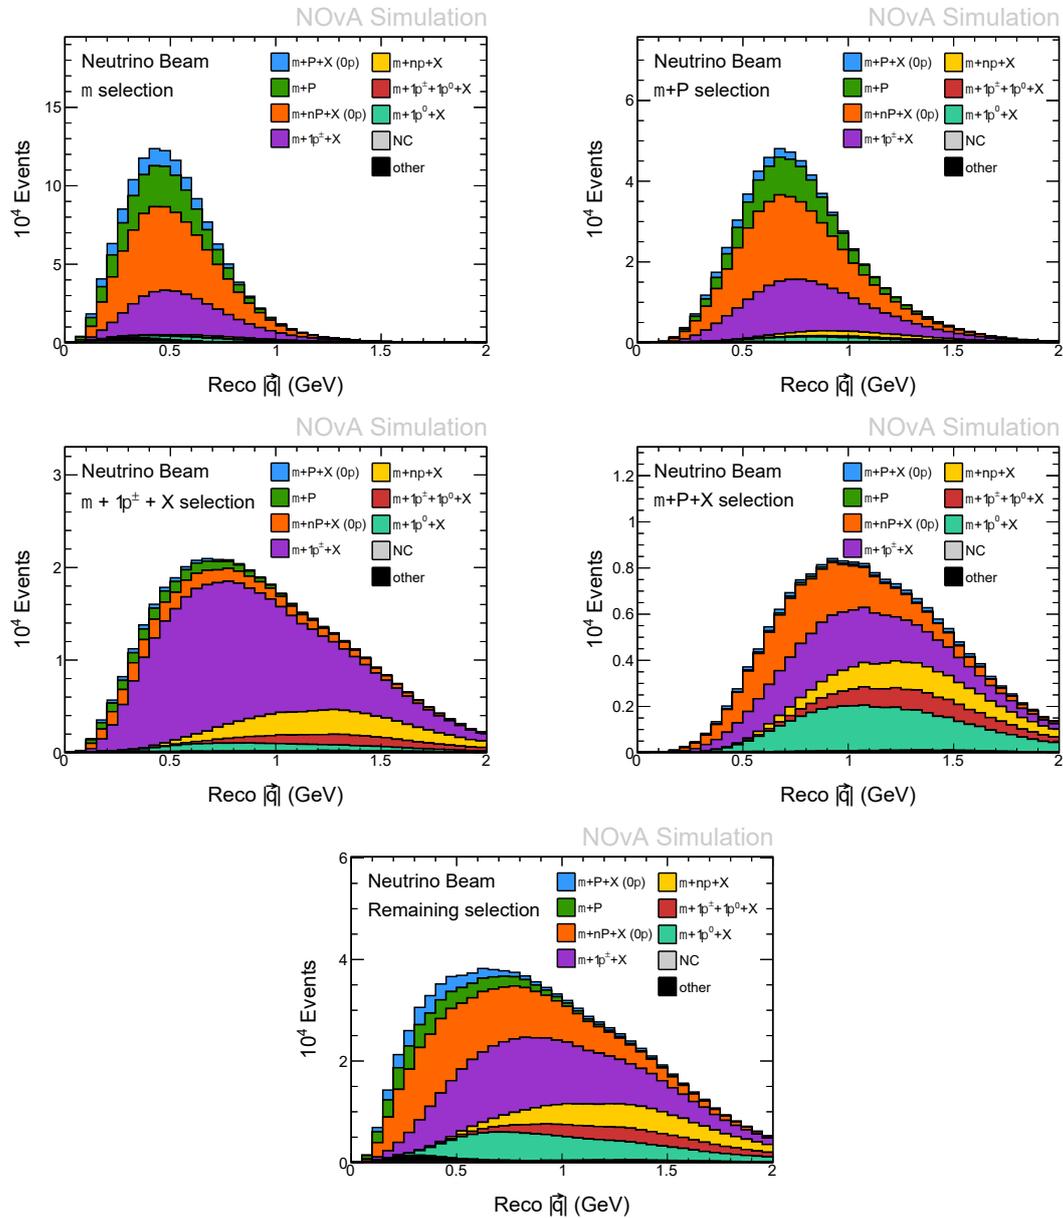


Figure A.4: Projection of the 2D ($E_{had}^{vis}, Reco |\vec{q}|$) phase space onto the $Reco |\vec{q}|$ axis. These plots of the FHC topologies break down neutrino events by the final state particles from the interaction. Complementary plots to Figure 4.10.

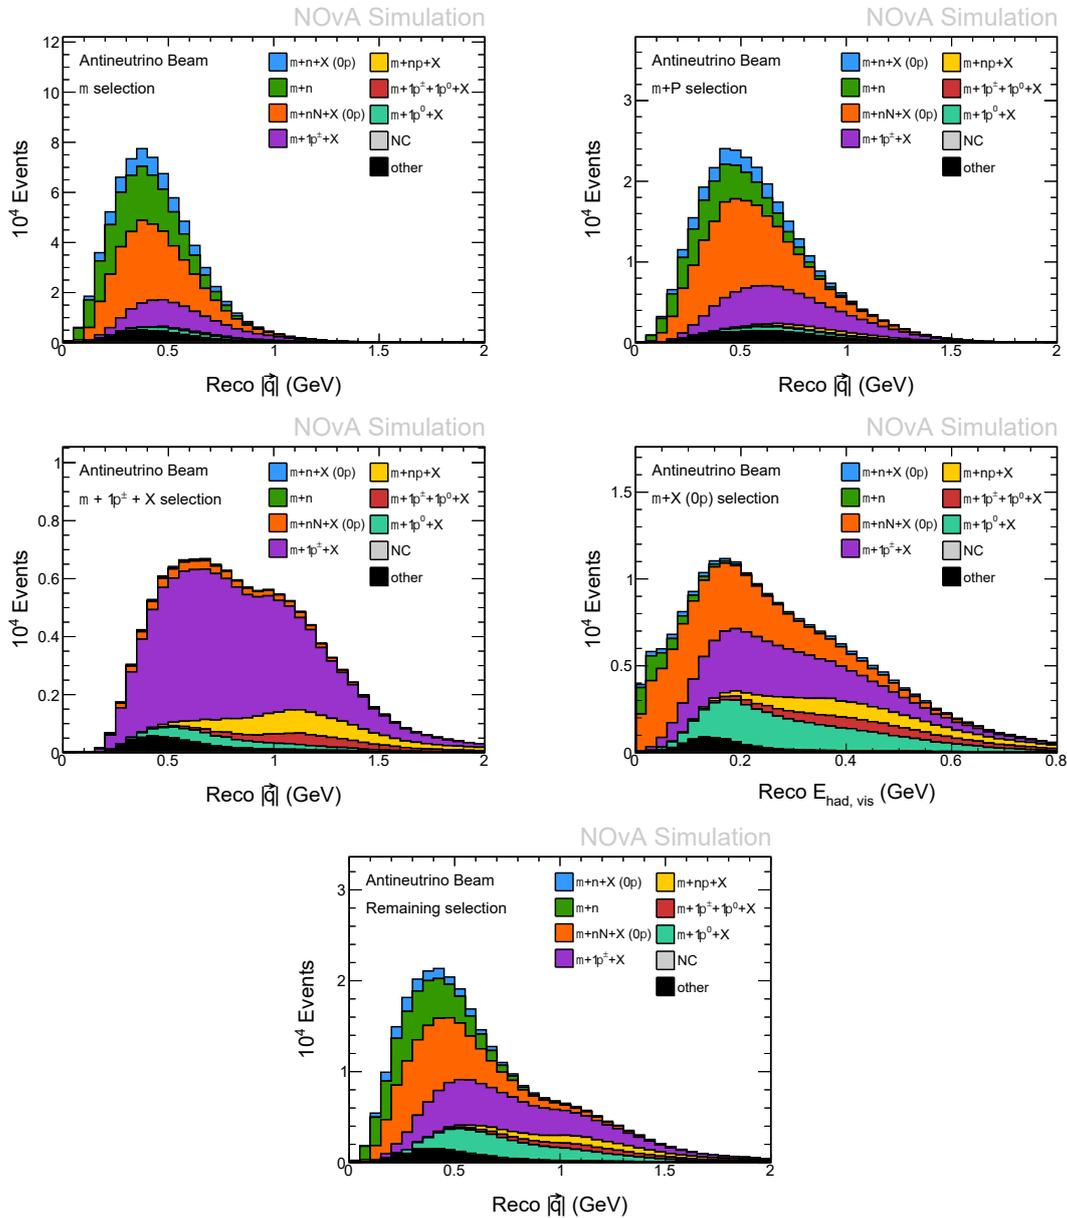


Figure A.5: Projection of the 2D ($E_{had}^{vis}, Reco |\vec{q}|$) phase space onto the $Reco |\vec{q}|$ axis. These plots of the RHC topologies break down neutrino events by the final state particles from the interaction. Complementary plots to Figure 4.11.

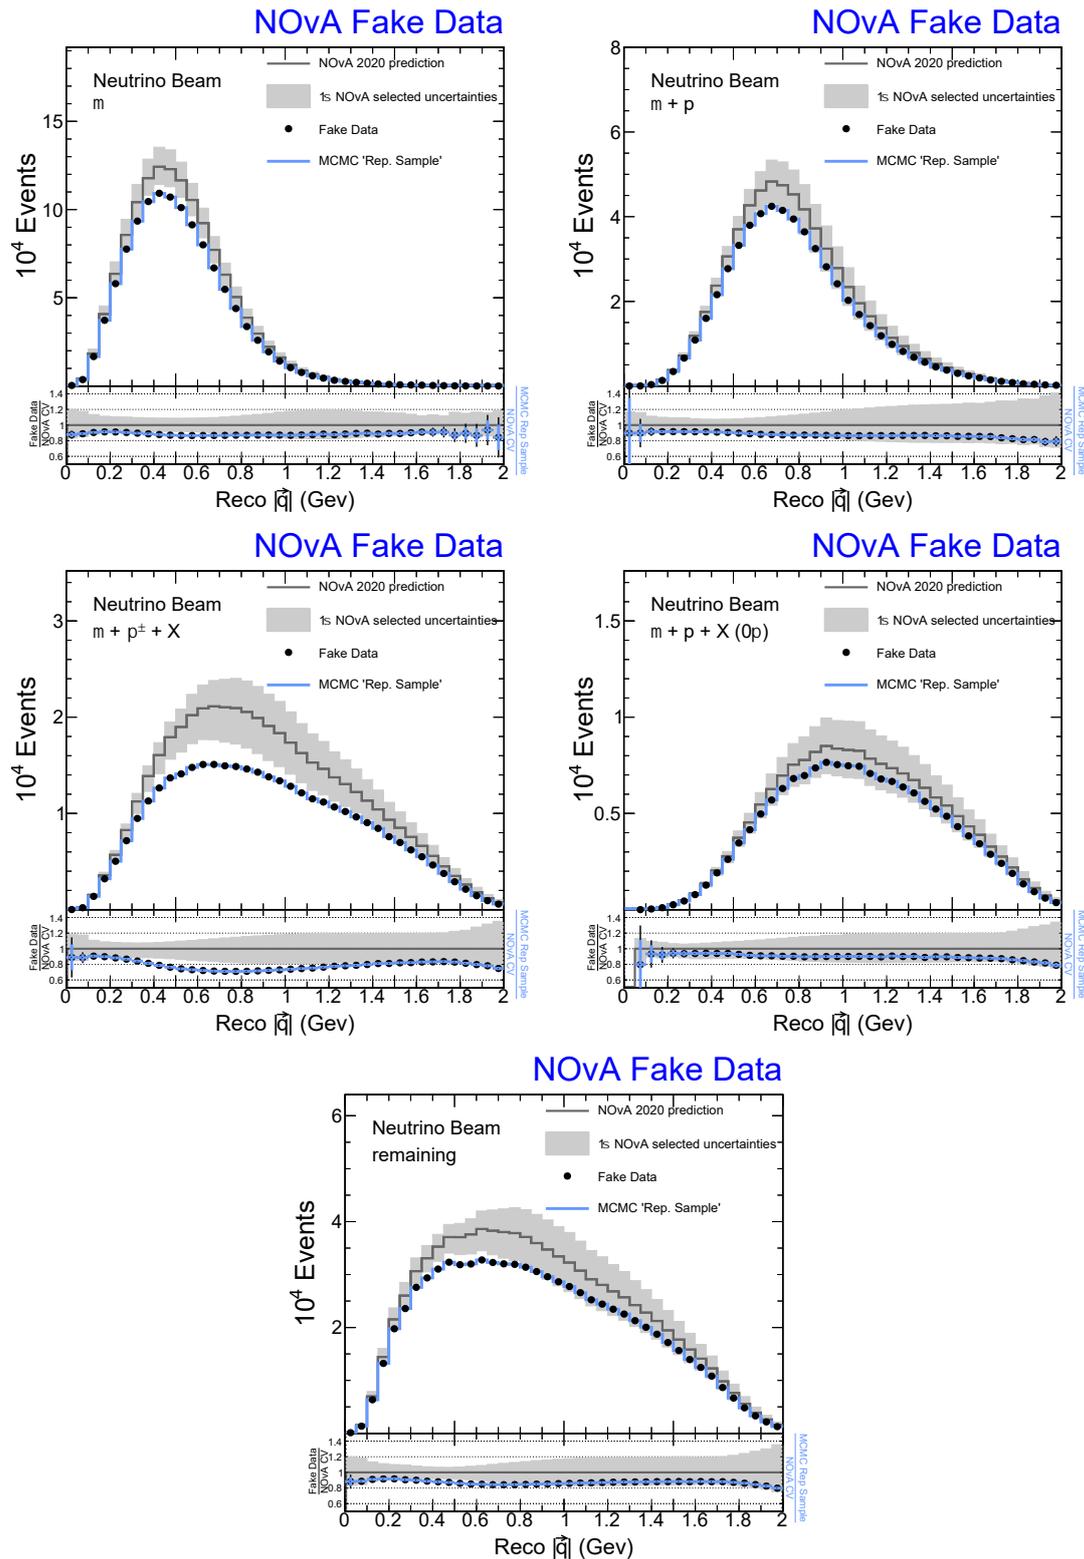


Figure A.6: Representative Sample predictions, in blue, on $Reco |\vec{q}|$ projection on FHC topologies. The grey histogram represents the nominal prediction and the grey error band represents the $\pm 1\sigma$ range for the 16 selected uncertainties. The fake data are the black dots. Note the Rep. Sample prediction (largest $\mathcal{L}\mathcal{L}$ MCMC sample) produces perfect agreement with the fake data, confirming MCMC can constrain NOvA uncertainties effectively. Complimentary to 6.3.

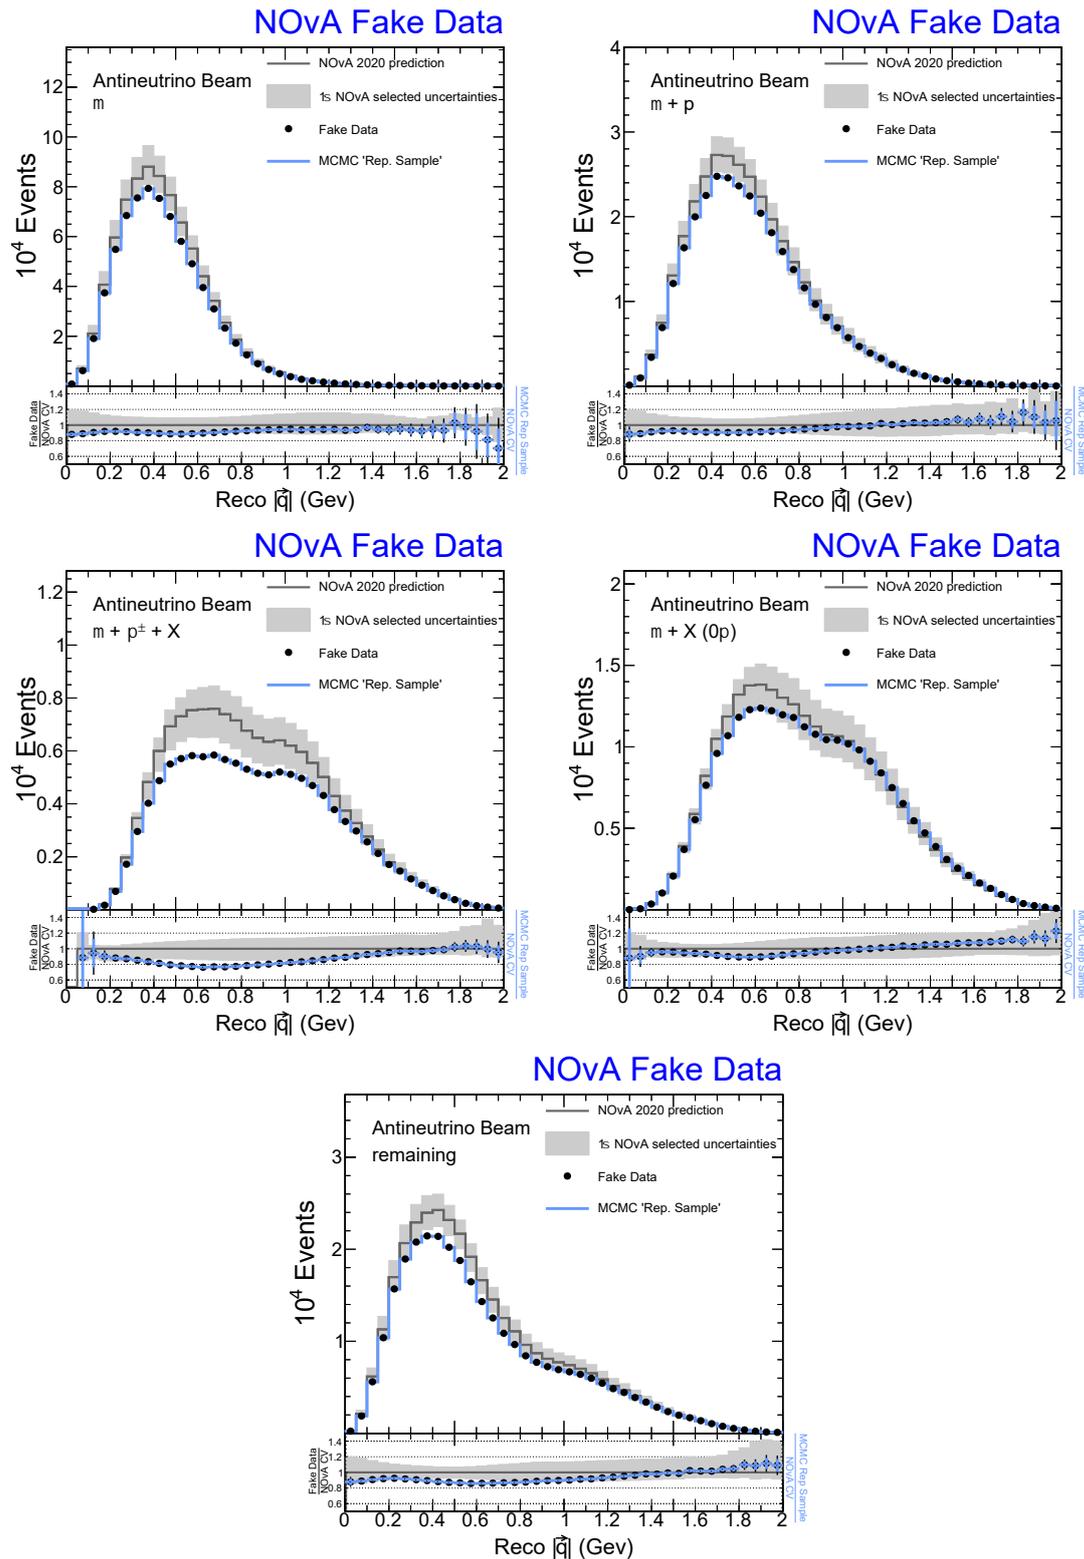


Figure A.7: Representative Sample predictions, in blue, projected onto $Reco |q|$ for the RHC topologies. The grey histogram represents the nominal prediction and the grey error band represents the 1 priori $\pm 1\sigma$ range for the 16 selected uncertainties. The fake data are the black dots. Note the Rep. Sample prediction (largest $\mathcal{L}\mathcal{L}$ MCMC sample) produces perfect agreement with the fake data, confirming MCMC can constrain NOvA uncertainties effectively. Complimentary to 6.4.

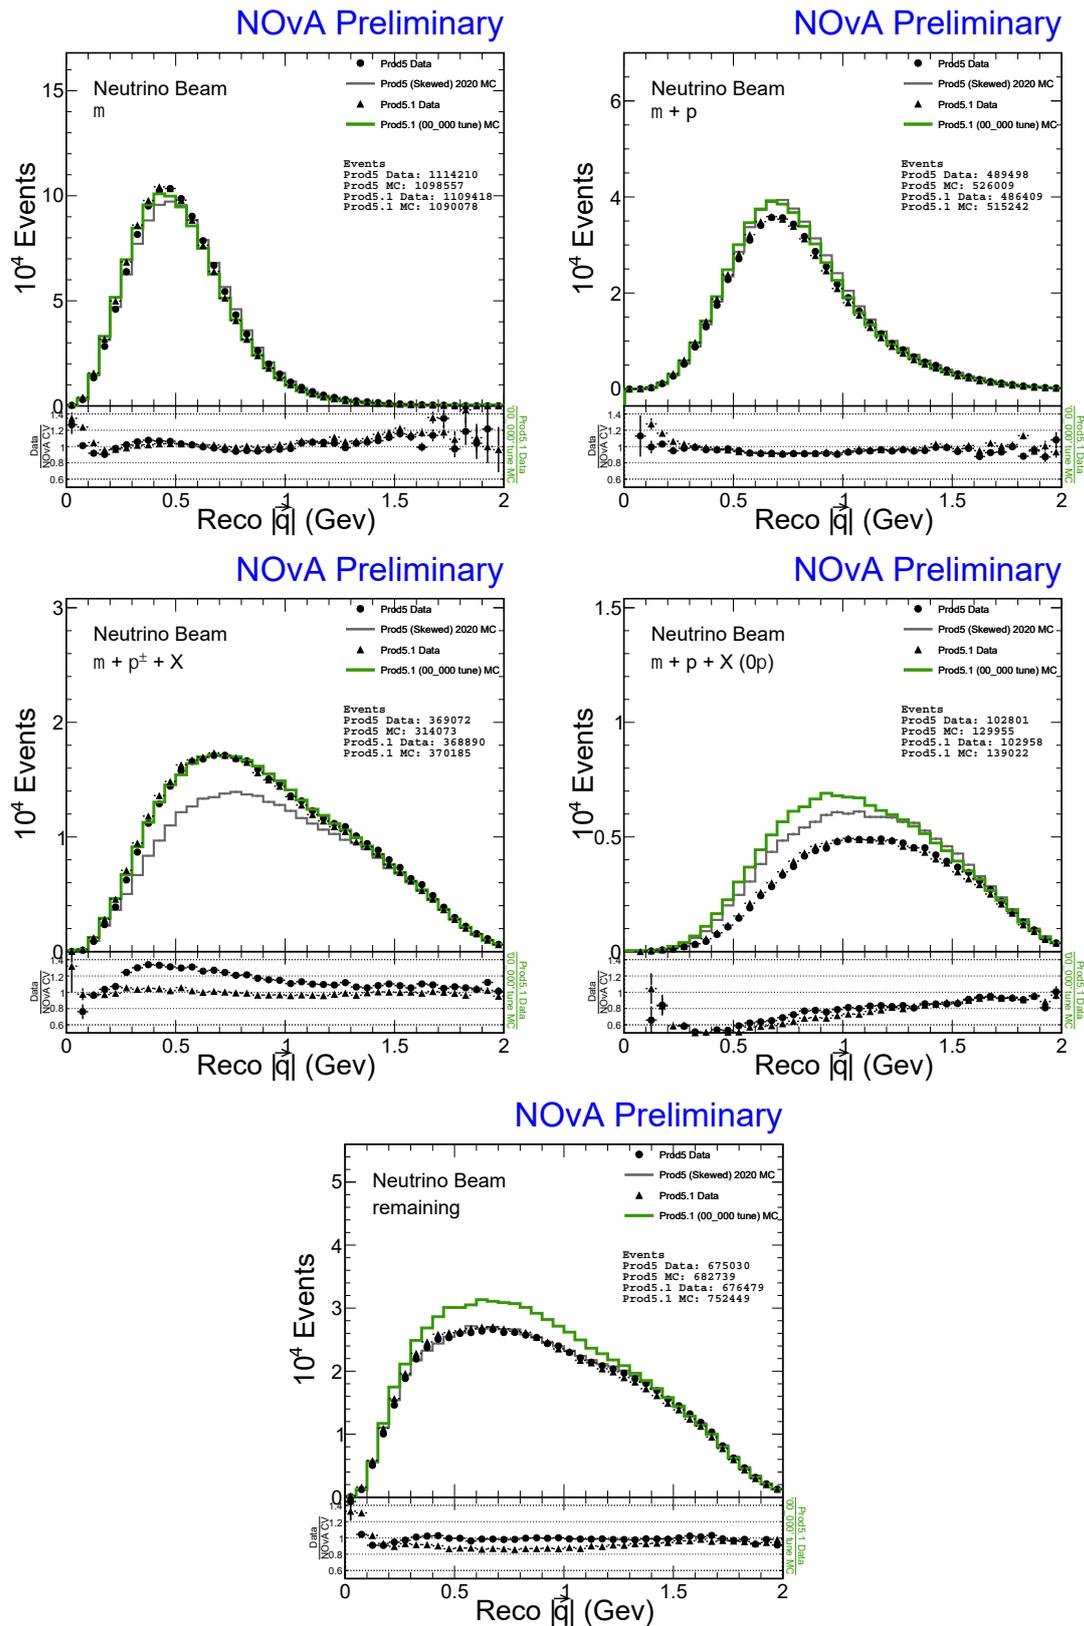


Figure A.8: Comparisons of Prod5 data (black dots) and Prod5 MC (grey) and Prod5.1 data (black triangles) and Prod5.1 MC (green) predictions projected onto $Reco |\vec{q}|$ in the FHC topological samples. This plot highlights the changes in the underlying physics models from Prod5 to Prod5.1. Note the disagreement between the Prod5.1 data and Prod5.1 MC, as this is what is used in this analysis. Complementary 1D projections in E_{had}^{vis} variable can be seen in the Fig. 6.9.

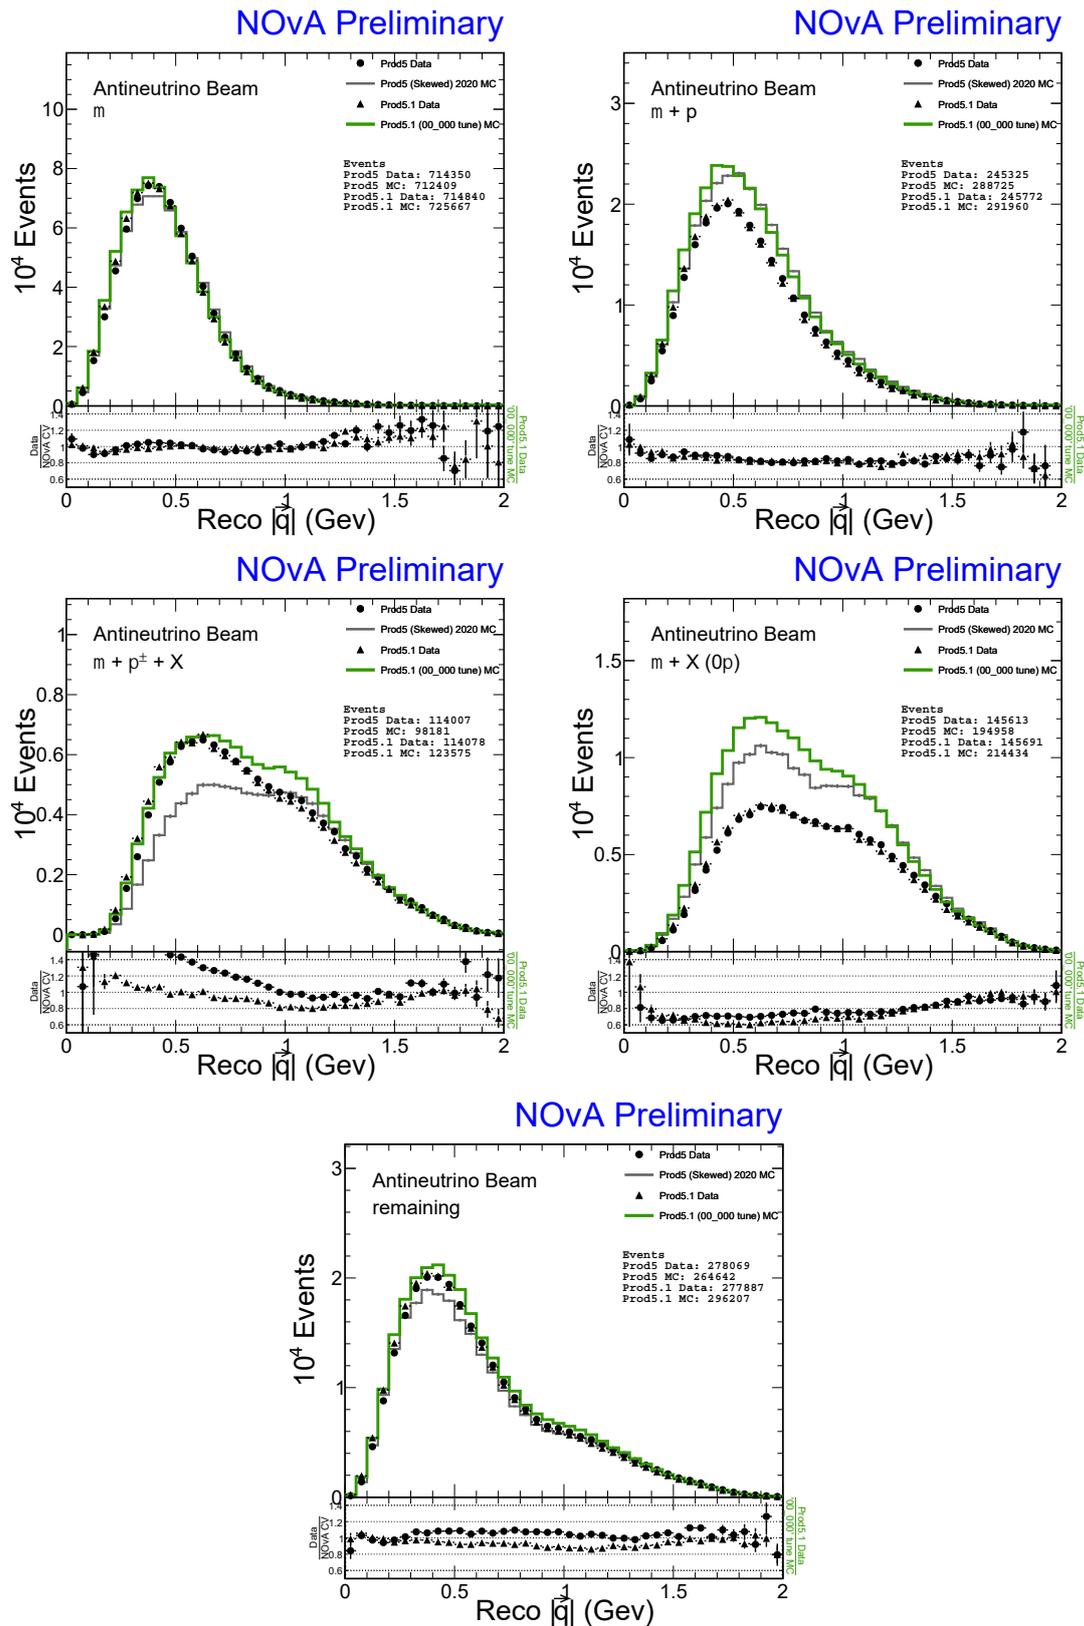


Figure A.9: Comparisons of Prod5 data (black dots) and Prod5 MC (grey) and Prod5.1 data (black triangles) and Prod5.1 MC (green) predictions projected onto $Reco |\vec{q}|$ in the RHC topological samples. This plot highlights the changes in the underlying physics models from Prod5 to Prod5.1. Note the disagreement between the Prod5.1 data and Prod5.1 MC, as this is what is used in this analysis. Complementary 1D projections in E_{had}^{vis} variable can be seen in the Fig. 6.10.

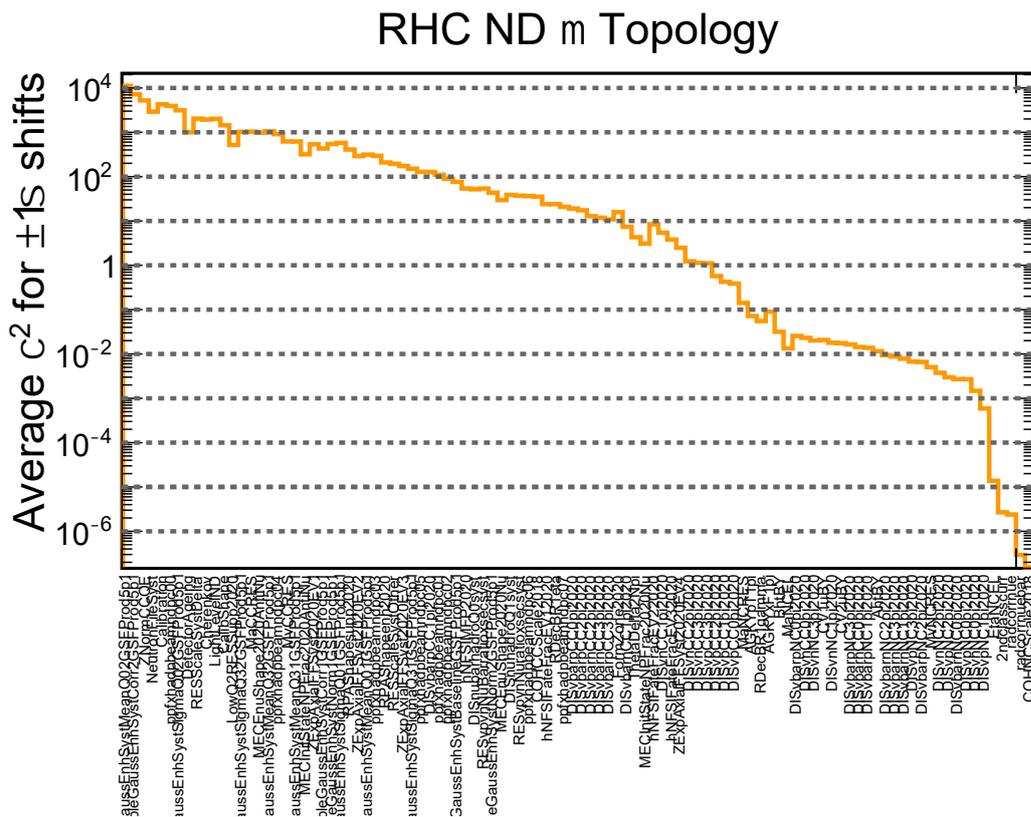


Figure A.10: ND RHC $\mu \chi^2$ systematic ranking between the nominal prediction and the $+1\sigma$ shifted prediction. Note the log-scale Y-axis, which illustrates the range of χ^2 values for NOvA's systematic uncertainties. The primary conclusion is there are 10-15 systematics that alter the ND prediction substantially, and there are 15-20 systematics that alter the ND predictions very little.

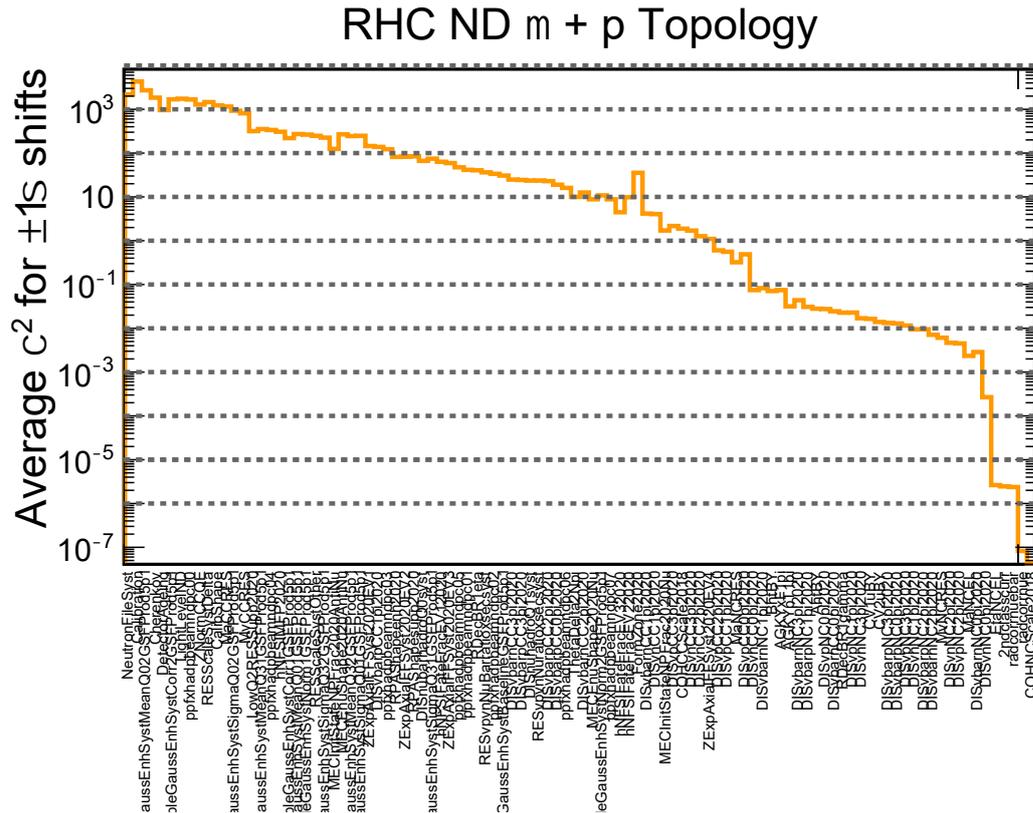


Figure A.11: ND RHC $\mu + P$ χ^2 systematic ranking between the nominal prediction and the $+1\sigma$ shifted prediction. Note the log-scale Y-axis, which illustrates the range of χ^2 values for NOvA’s systematic uncertainties. The primary conclusion is there are 10-15 systematics that alter the ND prediction substantially, and there are 15-20 systematics that alter the ND predictions very little.

A.3 ND Fit to Data

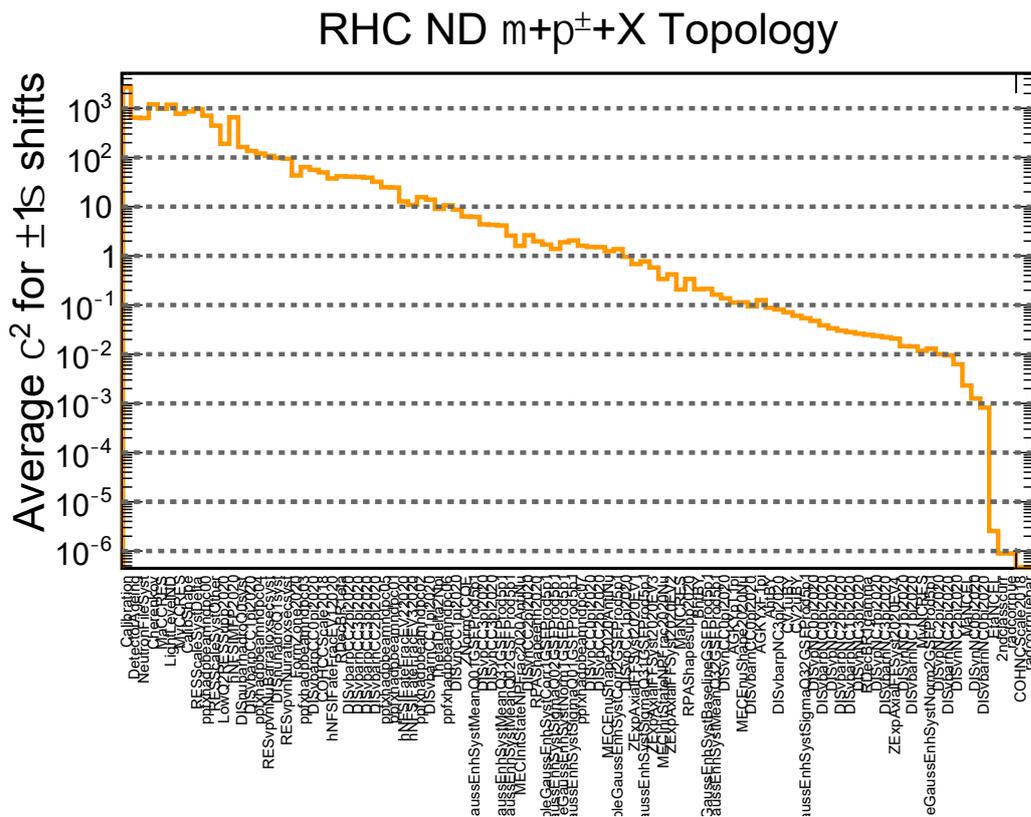


Figure A.12: ND RHC $\mu + \pi^\pm + X$ χ^2 systematic ranking between the nominal prediction and the $+1\sigma$ shifted prediction. Note the log-scale Y-axis, which illustrates the range of χ^2 values for NOvA's systematic uncertainties. The primary conclusion is there are 10-15 systematics that alter the ND prediction substantially, and there are 15-20 systematics that alter the ND predictions very little.

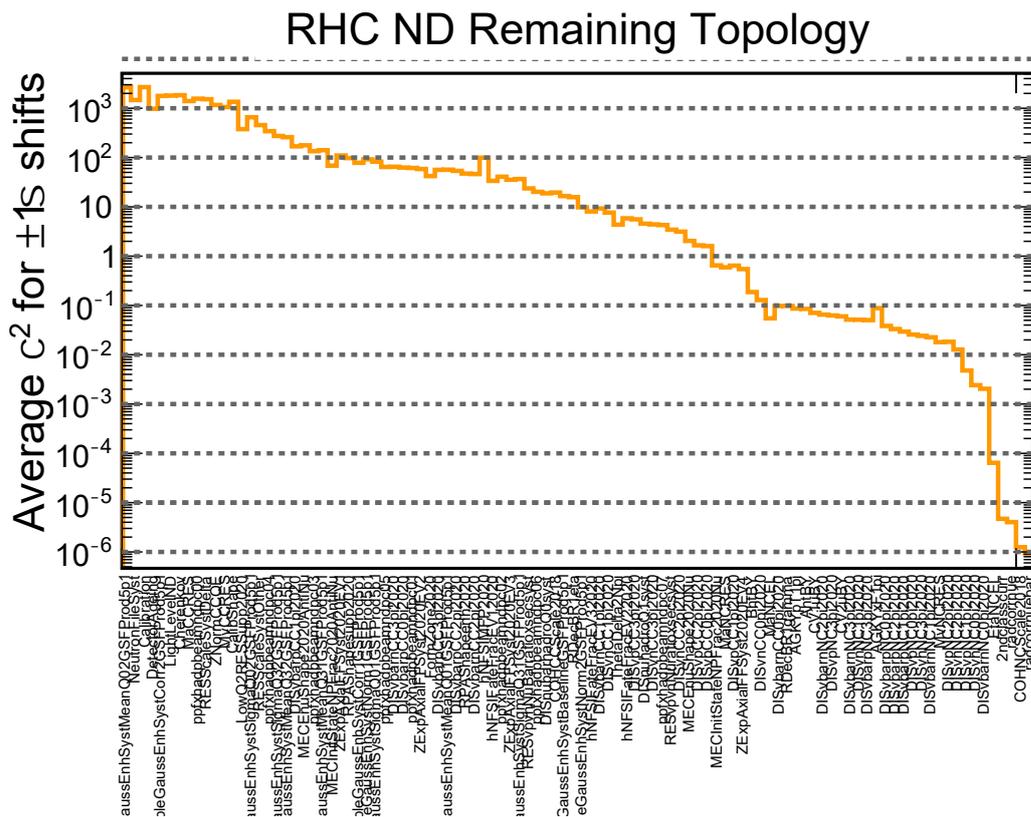


Figure A.14: ND RHC EvElse χ^2 systematic ranking between the nominal prediction and the $+1\sigma$ shifted prediction. Note the log-scale Y-axis, which illustrates the range of χ^2 values for NOvA’s systematic uncertainties. The primary conclusion is there are 10-15 systematics that alter the ND prediction substantially, and there are 15-20 systematics that alter the ND predictions very little.

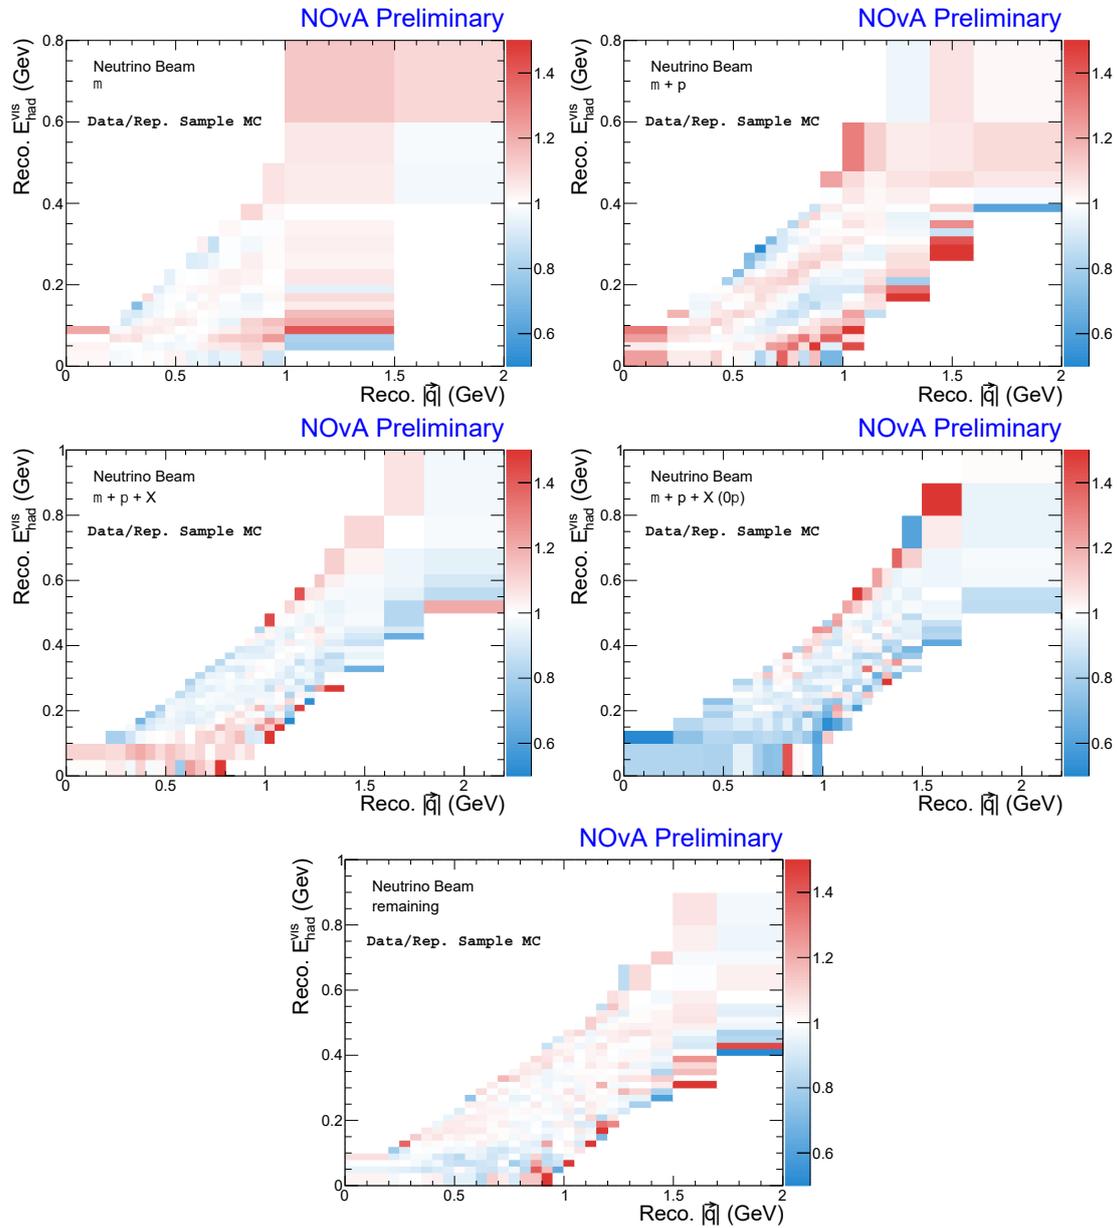


Figure A.15: 2D Data-Rep. Sample MC ratio distributions in ND FHC Topologies. Complimentary to the 2D data-MC ratio distributions from Fig. 6.11 and the 2D χ^2 distributions from Fig. 6.30.

A.4 ND+FD Joint Fit

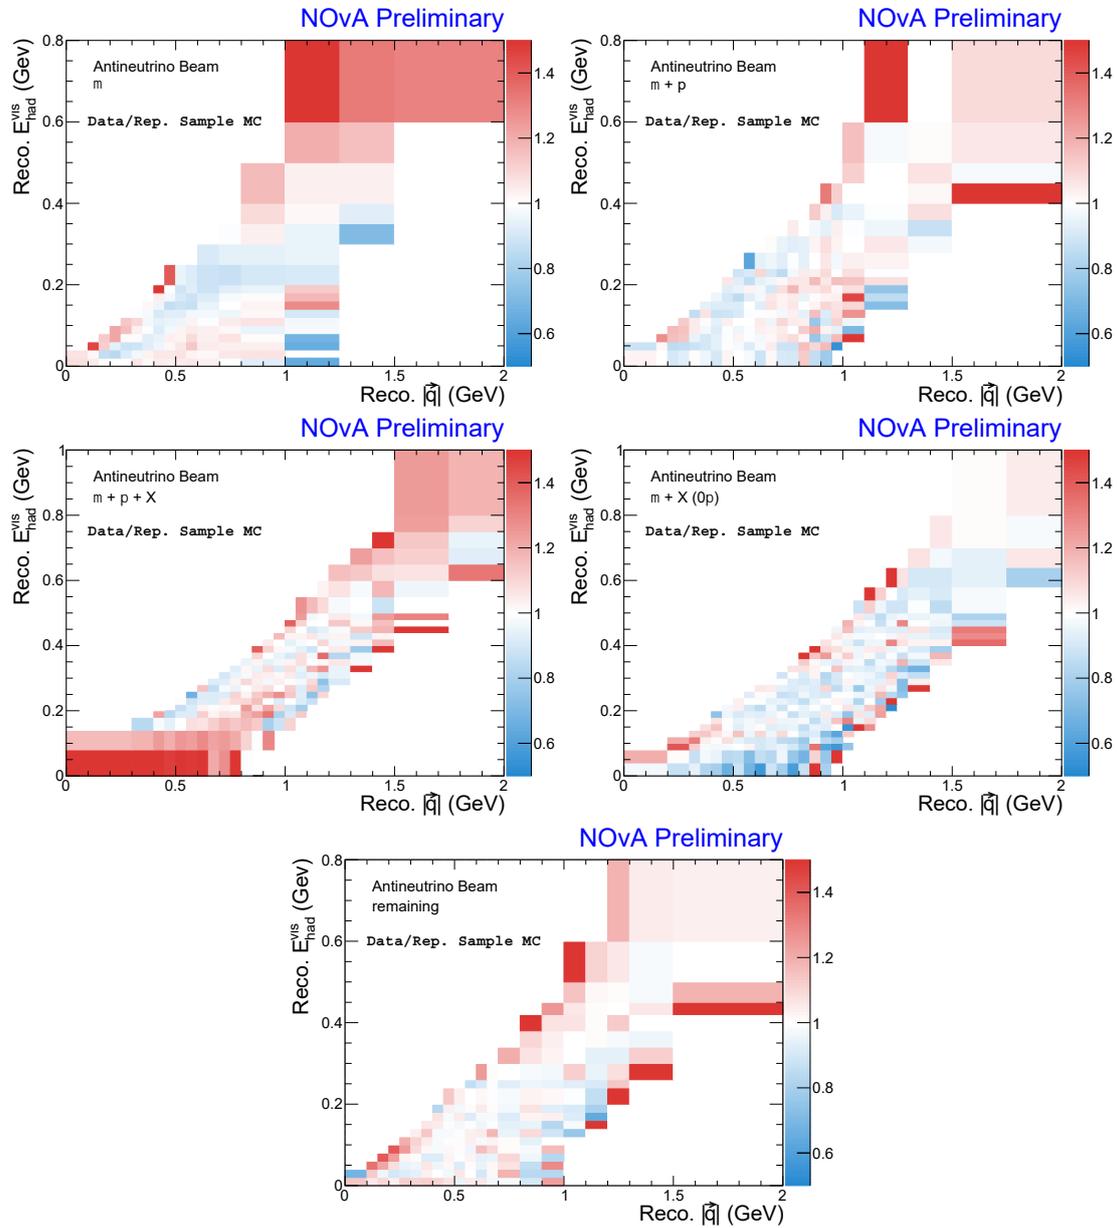


Figure A.16: 2D Data-Rep. Sample MC ratio distributions in ND RHC Topologies. Complimentary to the 2D data-MC ratio distributions from Fig. 6.12 and the 2D χ^2 distributions from Fig. 6.32.

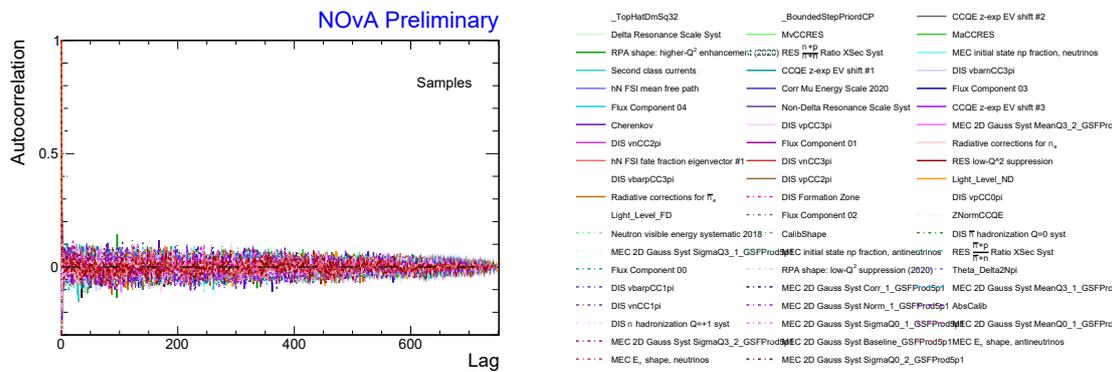


Figure A.17: Plot of the autocorrelations for a single, random MCMC chain from the ND+FD joint fit. There are 750 samples within a single MCMC chain. The autocorrelations are approximately the 10% suggesting the MCMC chain is efficiently moving through the posterior.

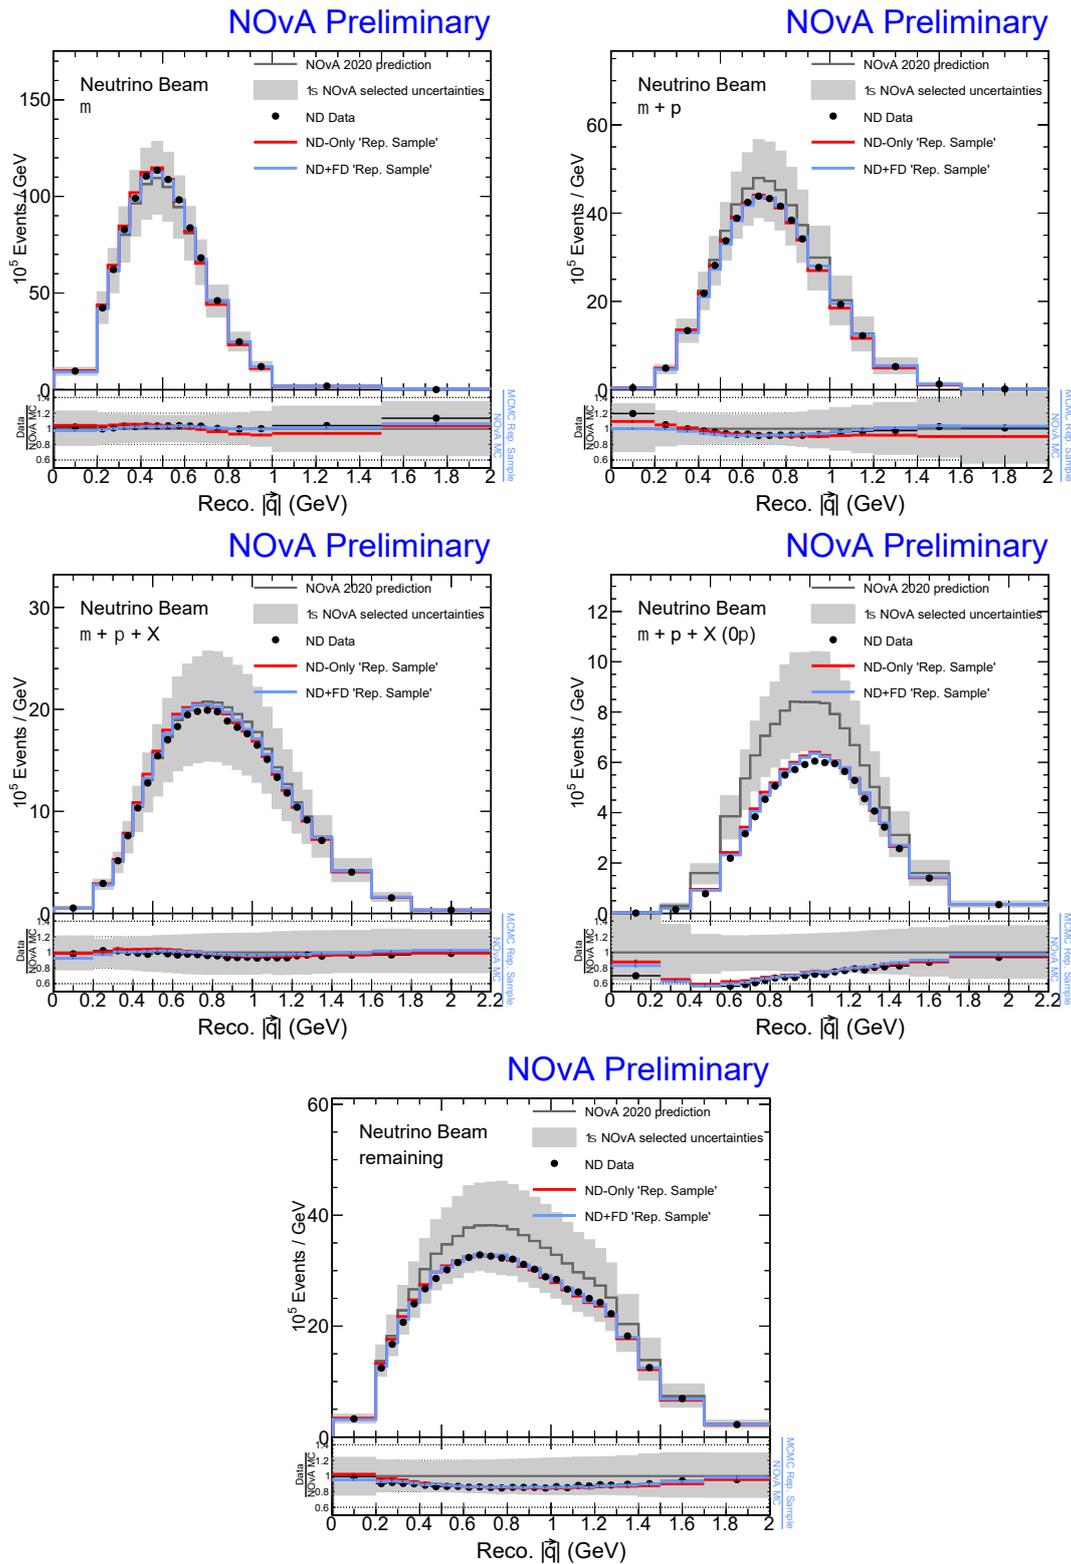


Figure A.18: 1D Projections onto $Reco. |\vec{q}|$ of the Rep. Sample from the ND-only fit and ND+FD joint fit in ND FHC topologies. Note how removing the muon energy scale systematic in the ND-only fit (red) worsens the agreement in this variable.

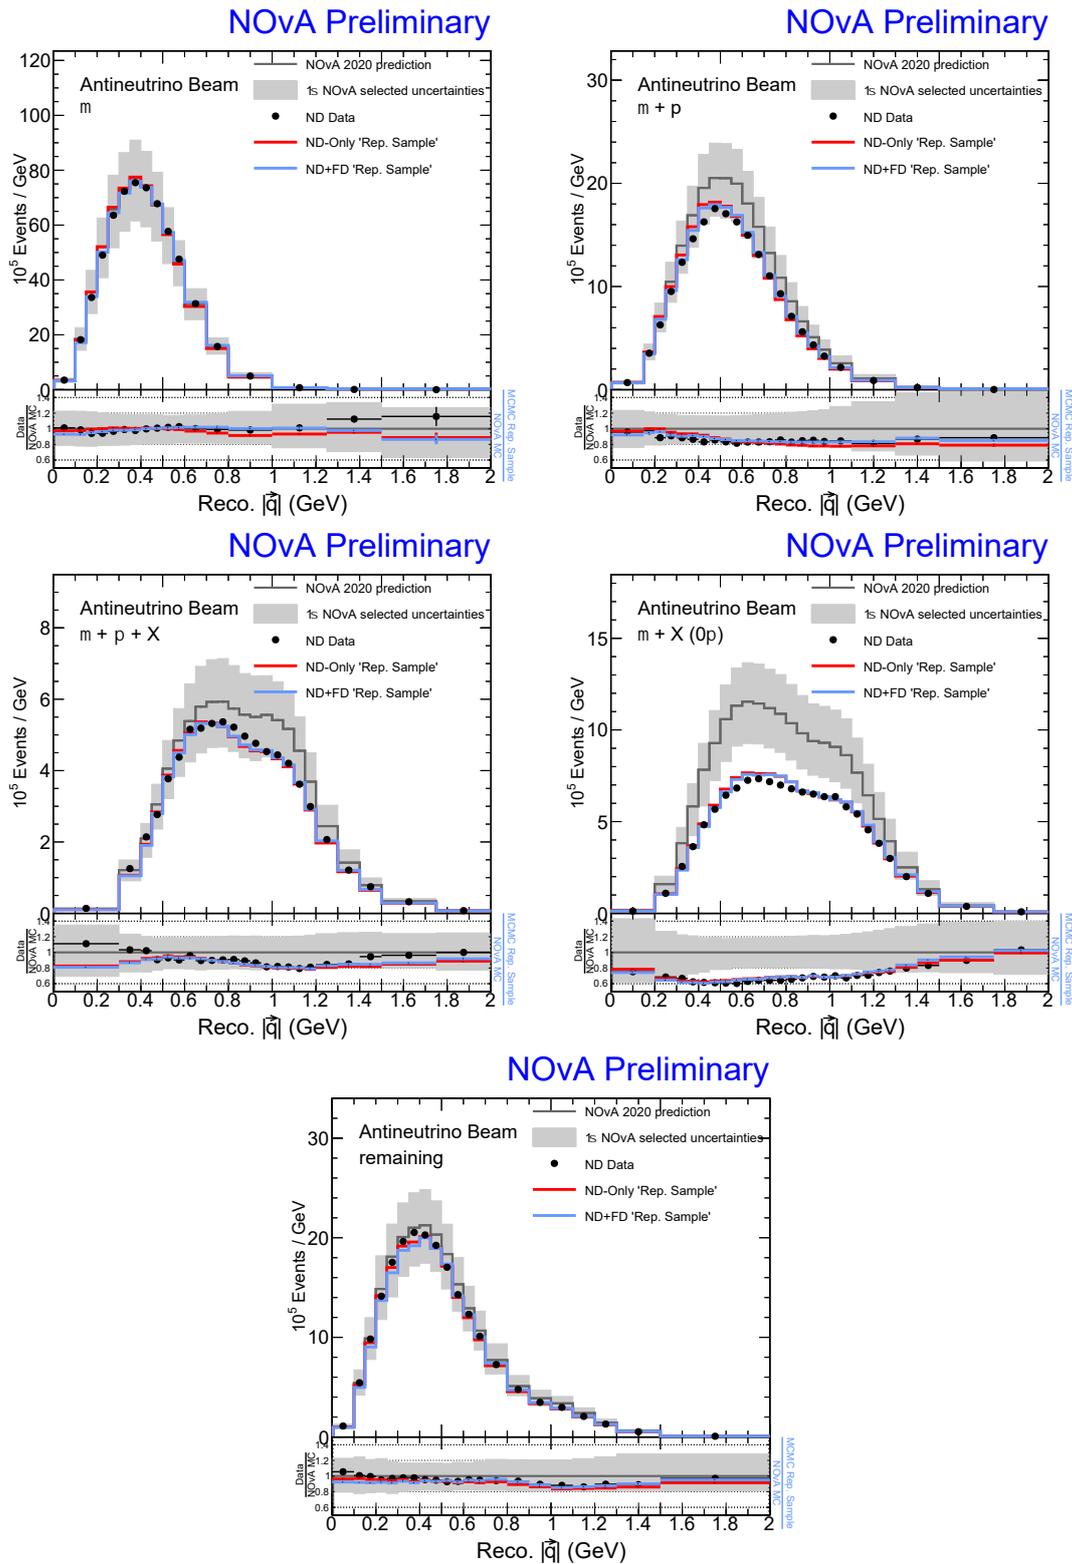
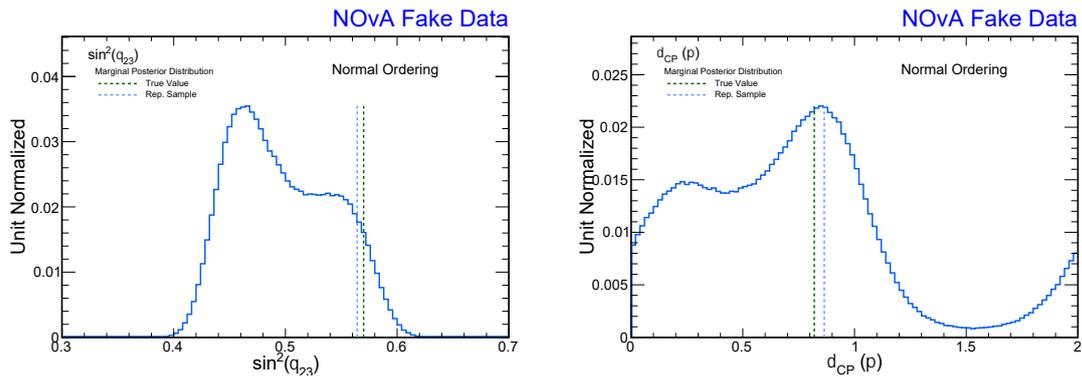
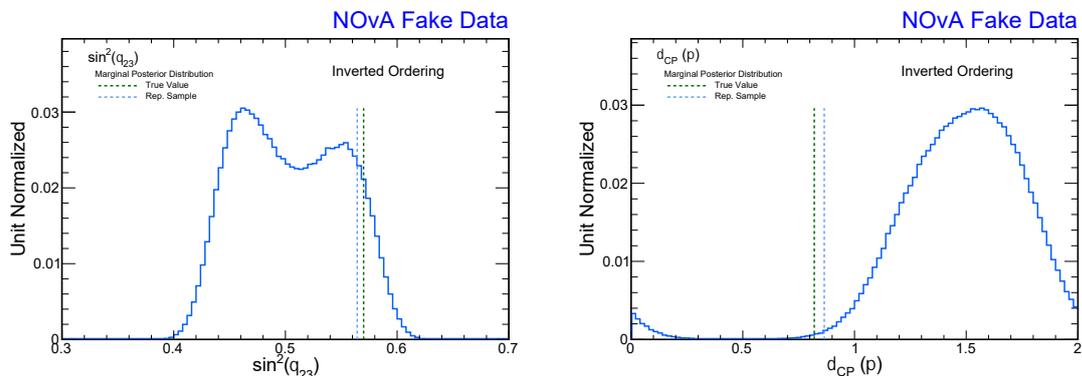


Figure A.19: 1D Projections onto $Reco. |\vec{q}|$ of the Rep. Sample from the ND-only fit and ND+FD joint fit in ND RHC topologies. Note how removing the muon energy scale systematic in the ND-only fit (red) worsens the agreement in this variable.

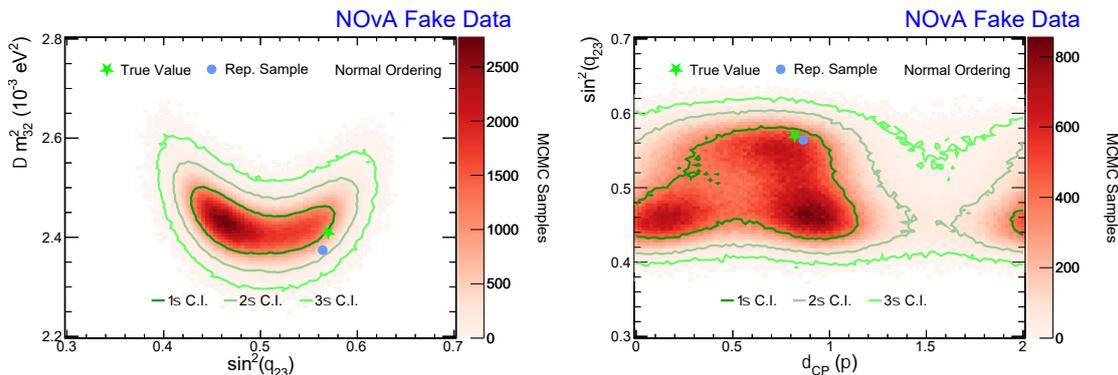


(a) Posterior probability distribution of $\sin^2 \theta_{23}$ marginalized over the Normal Ordering samples only. (b) Posterior probability distribution of δ_{CP} marginalized over the Normal Ordering samples only.

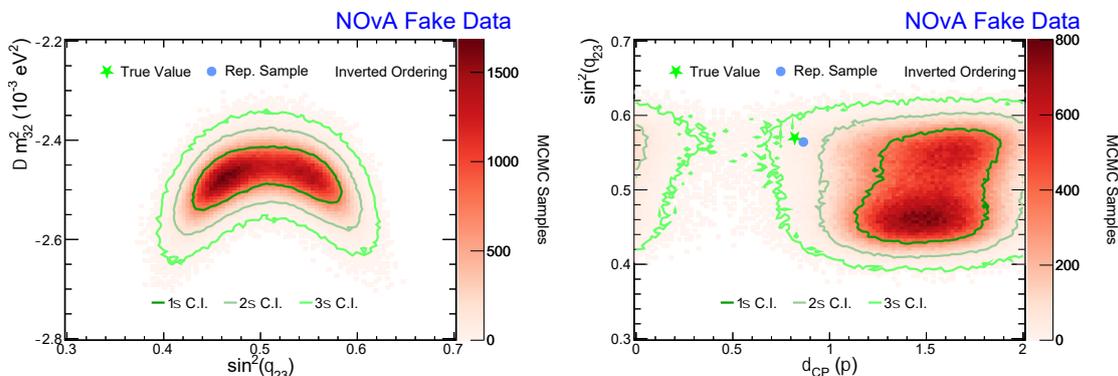


(c) Posterior probability distribution of $\sin^2 \theta_{23}$ marginalized over the Inverted Ordering samples only. (d) Posterior probability distribution of δ_{CP} marginalized over the Inverted Ordering samples only.

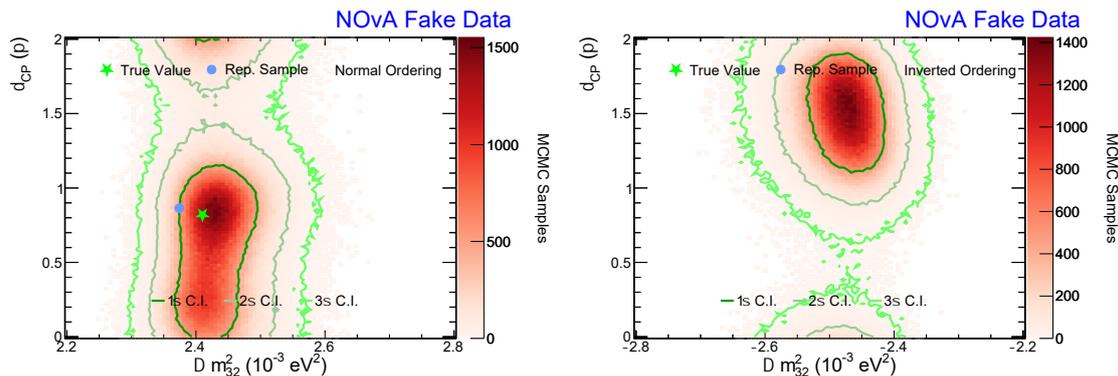
Figure A.20: 1D posterior probability distributions marginalized over the Normal and Inverted Ordering samples separately.



(a) Posterior probability distribution of Δm_{32}^2 and $\sin^2 \theta_{23}$ marginalized over the Normal Ordering samples only. (b) Posterior probability distribution of δ_{CP} and $\sin^2 \theta_{23}$ marginalized over the Normal Ordering samples.



(c) Posterior probability distribution of Δm_{32}^2 and $\sin^2 \theta_{23}$ marginalized over the Inverted Ordering samples only. (d) Posterior probability distribution of δ_{CP} and $\sin^2 \theta_{23}$ marginalized over the Inverted Ordering samples only.



(e) Posterior probability distribution of Δm_{32}^2 and δ_{CP} marginalized over the Normal Ordering samples only. (f) Posterior probability distribution of Δm_{32}^2 and δ_{CP} marginalized over the Inverted Ordering samples only.

Figure A.21: 2D posterior probability distributions marginalized over the Normal and Inverted Ordering samples separately. The green lines represent the Bayesian 1σ , 2σ , and 3σ credible intervals.

Bibliography

- [1] IceCube Neutrino Observatory. *Pauli presents hypothetical “neutron” particle*. <https://icecube.wisc.edu/neutrino-history/1931/01/1931-pauli-presents-hypothetical-neutron-particle/>. Accessed: 2023-06-23 (cit. on p. 2).
- [2] P. Hernandez. “Neutrino physics”. en. In: *CERN Yellow Reports* (2016). DOI: 10.5170/CERN-2016-005.85. URL: <https://e-publishing.cern.ch/index.php/CYR/article/view/434> (cit. on p. 2).
- [3] Mark Thomson. *Modern Particle Physics*. 4th ed. India: Cambridge University Press, 2013 (cit. on pp. 3, 4).
- [4] M. Tanabashi et al. “Review of Particle Physics”. In: *Phys. Rev. D* 98 (3 Aug. 2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001> (cit. on pp. 4, 98).
- [5] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. In: *Sov. Phys. JETP* 26 (1968), pp. 984–988 (cit. on p. 6).
- [6] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the unified model of elementary particles”. In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870 (cit. on p. 6).
- [7] K. Abe et al. “Combined Analysis of Neutrino and Antineutrino Oscillations at T2K”. In: *Phys. Rev. Lett.* 118.15 (2017), p. 151801. DOI: 10.1103/PhysRevLett.118.151801. arXiv: 1701.00432 [hep-ex] (cit. on p. 7).
- [8] F. P. An et al. “Observation of electron-antineutrino disappearance at Daya Bay”. In: *Phys. Rev. Lett.* 108 (2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803. arXiv: 1203.1669 [hep-ex] (cit. on p. 7).
- [9] S. Abe et al. “Precision Measurement of Neutrino Oscillation Parameters with KamLAND”. In: *Physical Review Letters* 100.22 (June 2008). DOI: 10.1103/physrevlett.100.221803. URL: <https://doi.org/10.1103%5C%2Fphysrevlett.100.221803> (cit. on p. 7).
- [10] Ivan Esteban et al. “Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23} , δ_{CP} , and the mass ordering”. In: *Journal of High Energy Physics* 2019.1 (Jan. 2019). DOI: 10.1007/jhep01(2019)106. URL: <https://doi.org/10.1007%5C%2Fjhep01%5C%282019%5C%29106> (cit. on p. 8).
- [11] A. Yu. Smirnov. “The MSW effect and solar neutrinos”. In: *10th International Workshop on Neutrino Telescopes*. May 2003, pp. 23–43. arXiv: hep-ph/0305106 (cit. on p. 11).

- [12] Erika Catano-Mur. “Recent results from NOvA”. In: *56th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. June 2022. arXiv: 2206.03542 [hep-ex] (cit. on p. 14).
- [13] Alexander Craig Booth. “Electron Neutrino Appearance at the NOvA Experiment”. In: (Jan. 2021). URL: <https://www.osti.gov/biblio/1827395> (cit. on pp. 15, 160).
- [14] Alex Himmel. *New Oscillation Results from the NOvA Experiment*. July 2020. DOI: 10.5281/zenodo.4142045. URL: <https://doi.org/10.5281/zenodo.4142045> (cit. on pp. 17, 19, 20, 41).
- [15] A. Abed Abud et al. *Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report*. 2021. arXiv: 2103.13910 [physics.ins-det] (cit. on p. 20).
- [16] Luis Alvarez-Ruso et al. “Recent highlights from GENIE v3”. In: *Eur. Phys. J. ST* 230.24 (2021), pp. 4449–4467. DOI: 10.1140/epjs/s11734-021-00295-7. arXiv: 2106.09381 [hep-ph] (cit. on p. 22).
- [17] J. Nieves, J. E. Amaro, and M. Valverde. “Inclusive quasielastic charged-current neutrino-nucleus reactions”. In: *Phys. Rev. C* 70 (5 Nov. 2004), p. 055503. DOI: 10.1103/PhysRevC.70.055503. URL: <https://link.aps.org/doi/10.1103/PhysRevC.70.055503> (cit. on p. 22).
- [18] R. Gran et al. “Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV”. In: *Phys. Rev. D* 88 (11 Dec. 2013), p. 113007. DOI: 10.1103/PhysRevD.88.113007. URL: <https://link.aps.org/doi/10.1103/PhysRevD.88.113007> (cit. on pp. 22, 30).
- [19] Ch. Berger and L. M. Sehgal. “Lepton mass effects in single pion production by neutrinos”. In: *Phys. Rev. D* 76 (11 Dec. 2007), p. 113004. DOI: 10.1103/PhysRevD.76.113004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.76.113004> (cit. on pp. 22, 34).
- [20] A. Bodek and U. K. Yang. *Modeling Neutrino and Electron Scattering Inelastic Cross Sections*. 2003. arXiv: hep-ex/0308007 [hep-ex] (cit. on pp. 22, 35).
- [21] Steven Dytman. “Final State Interactions in Neutrino-Nucleus Experiments”. In: *Acta Physica Polonica B - ACTA PHYS POL B* 40 (Sept. 2009) (cit. on pp. 22, 37).
- [22] J L Hewett et al. “Fundamental physics at the intensity frontier. Report of the workshop held December 2011 in Rockville, MD.” In: (June 2012). DOI: 10.2172/1042577. URL: <https://www.osti.gov/biblio/1042577> (cit. on p. 23).
- [23] Arie Bodek and Tejin Cai. “Removal energies and final state interaction in lepton nucleus scattering”. In: *The European Physical Journal C* 79.4 (Apr. 2019). DOI: 10.1140/epjc/s10052-019-6750-3. URL: <https://doi.org/10.1140/epjc/s10052-019-6750-3> (cit. on pp. 26, 27).
- [24] Michael Dolce. “Removal Energy Investigation and Nuclear Potential Reweighting”. July 2020 (cit. on p. 26).

- [25] Aaron S. Meyer et al. “Deuterium target data for precision neutrino-nucleus cross sections”. In: *Phys. Rev. D* 93 (11 June 2016), p. 113015. DOI: 10.1103/PhysRevD.93.113015. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.113015> (cit. on p. 28).
- [26] Kirk Bays et al. “NOvA 2020 cross-section tune tech note”. In: *Internal Document NOvA DocDB 43962* (Sept. 2020) (cit. on pp. 28, 33, 34).
- [27] F. James and M. Roos. “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations”. In: *Comput. Phys. Commun.* 10 (1975), pp. 343–367. DOI: 10.1016/0010-4655(75)90039-9 (cit. on p. 32).
- [28] P. Stowell et al. “Tuning the GENIE Pion Production Model with MINERvA Data”. In: *Phys. Rev. D* 100.7 (2019), p. 072005. DOI: 10.1103/PhysRevD.100.072005. arXiv: 1903.01558 [hep-ex] (cit. on p. 35).
- [29] F. Sánchez. “Possibility of measuring Adler angles in charged current single pion neutrino-nucleus interactions”. In: *Phys. Rev. D* 93 (9 May 2016), p. 093015. DOI: 10.1103/PhysRevD.93.093015. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.093015> (cit. on p. 35).
- [30] E. S. Pinzon Guerra et al. “Using world charged π^\pm -nucleus scattering data to constrain an intranuclear cascade model”. In: *Phys. Rev. D* 99.5 (2019), p. 052007. DOI: 10.1103/PhysRevD.99.052007. arXiv: 1812.06912 [hep-ex] (cit. on p. 37).
- [31] L. L. Salcedo et al. “Computer Simulation of Inclusive Pion Nuclear Reactions”. In: *Nucl. Phys. A* 484 (1988), pp. 557–592. DOI: 10.1016/0375-9474(88)90310-7 (cit. on p. 37).
- [32] Jeremy Wolcott Michael Dolce and Hugh Gallagher. “Central value tuning and uncertainties for hN FSI model in GENIE 3, NOvA tech note”. Feb. 2020 (cit. on pp. 38–40).
- [33] C. Andreopoulos et al. “The GENIE Neutrino Monte Carlo Generator”. In: *Nucl. Instrum. Meth. A* 614 (2010), pp. 87–104. DOI: 10.1016/j.nima.2009.12.009. arXiv: 0905.2517 [hep-ph] (cit. on p. 41).
- [34] Adrian Cho. “NOvA’s shining moment”. In: *Science* 345.6204 (2014), pp. 1555–1557. DOI: 10.1126/science.345.6204.1555. eprint: <https://www.science.org/doi/pdf/10.1126/science.345.6204.1555>. URL: <https://www.science.org/doi/abs/10.1126/science.345.6204.1555> (cit. on p. 43).
- [35] Fermilab. *Concepts Rookie Book*. Operations Rookie Books (cit. on p. 44).
- [36] Ashley Back. “All-time FD-POT exposure, updated May 2022”. May 2022 (cit. on pp. 48, 49).
- [37] D S Ayres et al. “The NOvA Technical Design Report”. In: (Oct. 2007). DOI: 10.2172/935497. URL: <https://www.osti.gov/biblio/935497> (cit. on p. 51).
- [38] Fermilab Creative Services (cit. on pp. 52–54).

- [39] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688> (cit. on pp. 53, 55).
- [40] Neutron Task Force. “NOvA neutron systematic for 3-flavor analysis tech note”. Feb. 2020 (cit. on p. 55).
- [41] Oleg Samoylov v hewes. “NOvA Production 5 Detector Simulation Changes and Systematics”. Mar. 2020 (cit. on p. 55).
- [42] Lisa Koerner and Ryan Nichol. “Calibration Tech Note 2020 Analysis”. Mar. 2020 (cit. on p. 55).
- [43] Neutron Task Force. “NuMI flux systematic uncertainties for the NOvA third analyses tech note”. Mar. 2017 (cit. on p. 56).
- [44] L. Aliaga et al. “Neutrino flux predictions for the NuMI beam”. In: *Phys. Rev. D* 94 (9 Nov. 2016), p. 092005. DOI: 10.1103/PhysRevD.94.092005. URL: <https://link.aps.org/doi/10.1103/PhysRevD.94.092005> (cit. on p. 56).
- [45] Nitish Nayak. “Prod5 Flux PCA”. Feb. 2020 (cit. on p. 56).
- [46] Jeremy Wolcott. “Flux PCA for ND+FD fitting”. Feb. 2023 (cit. on p. 56).
- [47] A. Norman et al. “Performance of the NOvA Data Acquisition and Trigger Systems for the full 14 kT Far Detector”. In: *Journal of Physics: Conference Series* 664.8 (Dec. 2015), p. 082041. DOI: 10.1088/1742-6596/664/8/082041. URL: <https://dx.doi.org/10.1088/1742-6596/664/8/082041> (cit. on p. 58).
- [48] Rudolf E. Kálmán. “A new approach to linear filtering and prediction problems” transaction of the asme journal of basic”. In: 1960 (cit. on p. 58).
- [49] P V.C. Hough. “METHOD AND MEANS FOR RECOGNIZING COMPLEX PATTERNS”. In: (Dec. 1962) (cit. on p. 59).
- [50] Mattias Ohlsson. “Extensions and explorations of the elastic arms algorithm”. In: *Computer Physics Communications* 77.1 (1993), pp. 19–32. ISSN: 0010-4655. DOI: [https://doi.org/10.1016/0010-4655\(93\)90033-9](https://doi.org/10.1016/0010-4655(93)90033-9). URL: <https://www.sciencedirect.com/science/article/pii/0010465593900339> (cit. on p. 59).
- [51] R. Krishnapuram and J.M. Keller. “A possibilistic approach to clustering”. In: *IEEE Transactions on Fuzzy Systems* 1.2 (1993), pp. 98–110. DOI: 10.1109/91.227387 (cit. on p. 59).
- [52] Yibing Zhang. *Ana2020 Numu FHC Event Displays*. June 2020 (cit. on p. 59).
- [53] “Summary of new and retuned reconstruction and PID modules in production 5”. Mar. 2020 (cit. on p. 60).
- [54] Travis Grant Olson. “Measurement of $d^2\sigma/d\vec{q}dE_{avail}$ and 2p2h contribution using charged current ν_μ interactions in the NOvA Near Detector”. PhD thesis. Tufts U., Tufts U., 2021 (cit. on pp. 60, 61, 68).

-
- [55] Jeremy Wolcott Michael Baird Louise Suter. “Executive summary for 2020 3-flavor analysis review”. Mar. 2020 (cit. on pp. 61, 81).
- [56] Thomas Warburton. “Ana 20 NuMu Quantile Plots”. Sept. 2020 (cit. on p. 80).
- [57] Yibing Zhang. *Ana2020 Nue FHC Event Displays*. June 2020 (cit. on p. 84).
- [58] D. S. Sivia and John Skilling. *Data analysis : a Bayesian tutorial*. 1996 (cit. on p. 87).
- [59] Michael Betancourt. *A Conceptual Introduction to Hamiltonian Monte Carlo*. 2018. arXiv: 1701.02434 [stat.ME] (cit. on p. 91).
- [60] J E Gubernatis. “Marshall Rosenbluth and the Metropolis algorithm”. In: *Physics of Plasmas* 12.5 (May 2005). ISSN: 1070-664X. DOI: 10.1063/1.1887186. URL: <https://www.osti.gov/biblio/20736632> (cit. on p. 93).
- [61] Matthew D. Hoffman and Andrew Gelman. *The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo*. 2011. arXiv: 1111.4246 [stat.CO] (cit. on p. 96).
- [62] Chi Feng. *MCMC Interactive Gallery*. <https://chi-feng.github.io/mcmc-demo/> (cit. on p. 97).
- [63] Stan Development Team. *MC-Stan*. <https://mc-stan.org> (cit. on p. 100).
- [64] Chris Backhouse. *CAFAna Overview*. CAFAna Overview (cit. on p. 100).
- [65] John P. Boyd. “Asymptotic Fourier Coefficients for a C^∞ Bell (Smoothed-“Top-Hat”) & the Fourier Extension Problem”. In: *Journal of Scientific Computing* 29.1 (2006), pp. 1–24. DOI: 10.1007/s10915-005-9010-7. URL: <https://doi.org/10.1007/s10915-005-9010-7> (cit. on p. 160).
- [66] I. Antcheva et al. “ROOT — A C++ framework for petabyte data storage, statistical analysis and visualization”. In: *Computer Physics Communications* 180.12 (Dec. 2009), pp. 2499–2512. DOI: 10.1016/j.cpc.2009.08.005. URL: <https://doi.org/10.1016%5C%2Fj.cpc.2009.08.005> (cit. on p. 171).