Measurements of Electron Antineutrino Disappearance in the Double Chooz Experiment

Rachel Carr

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2015

ABSTRACT

Measurements of Electron Antineutrino Disappearance in the Double Chooz Experiment

Rachel Carr

This thesis presents complementary measurements of the neutrino oscillation parameter $\sin^2 2\theta_{13}$ made by observing the electron antineutrino flux from two nuclear reactors in Chooz, France. Antineutrinos are identified through both products of the inverse beta decay interaction, $\bar{\nu}_e + p \rightarrow e^+ + n$, in a high-precision liquid scintillator detector located approximately 1 km from the reactors. The most sensitive signal channel involves neutron captures by a gadolinium dopant, while a search for neutron captures on hydrogen provides a high-statistics validation. In both channels, the value of $\sin^2 2\theta_{13}$ is revealed by an energy- and reactor power-dependent deficit of antineutrino-like events, compared to a reactor simulation. All analyses produce results consistent with one another and with findings of other experiments. These datasets also expose features of the antineutrino spectrum not predicted in reactor flux models. Data from the newly inaugurated near detector, anticipated in the final part of this thesis, holds unique potential to clarify these features, pursue signals from sterile neutrinos, and contribute to global knowledge of three-neutrino mixing.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

ν-DAQ	neutrino DAQ
ADC	analog-to-digital converter
ANN	artificial neural network
CC	charged current
CKM	Cabibbo-Kobayashi-Maskawa
CL	confidence level
CMB	cosmic microwave background
CP	charge-parity
d.o.f.	degrees of freedom
DAQ	data acquisition system
EDF	Électricité de France
EW	electroweak
FADC	flash ADC
FD	far detector
FEE	front-end electronics
FN+SM	fast neutron and stopping muon

GC	gamma catcher
HV	high voltage
IBD	inverse beta decay
ID	inner detector
IV	inner veto
LED	light emitting diode
MC	Monte Carlo
MLE	maximum likelihood estimation
NC	neutral current
ND	near detector
NT	neutrino target
o-PXE	ortho-phenylxylylethane
OV	outer veto
PDF	probability density function
PE	photo electron
PMT	photomultiplier tube
PPO	2,5-diphenyloxazole
PWR	pressurized water reactor
RRM	Reactor Rate Modulation
SM	Standard Model
UV	ultraviolet
VEV	vacuum expectation value

Acknowledgments

This thesis owes many numbers and ideas to my Double Chooz colleagues, and the final results belong to the entire analysis team. I have been fortunate to learn from and collaborate with such a dedicated and generous group. I am also grateful to have been a member of the Columbia and Nevis communities, particularly as a student of Mike Shaevitz. In this and all my pursuits, I am thankful to my parents, siblings, and husband for supporting me and for sharing with me their curiosity about the world.

Chapter 1

Introduction

Neutrinos are the most abundant massive particles in the universe, but because of their extraordinarily low interaction cross section, many of their basic features remain undiscovered. Observations over the past half-century demonstrate that these particles have mass, but its magnitude and underlying nature are unknown. Whether neutrinos behave identically to their antiparticles is also unestablished. Even the number of light neutrino species is an open question, as some evidence seems to defy the conventional three-neutrino framework. Resolutions of these and other problems in the neutrino sector may inform deeper issues in fundamental physics, from the origin of flavor to the primordial matter-antimatter asymmetry of the universe.

Important insight into the nature of neutrinos comes from the phenomenon of flavor oscillations. Through this effect, unique among fundamental particles, a neutrino created in one flavor has a nonzero probability of being detected in a different flavor. This probability varies periodically as a function of distance traveled and, over a neutrino spectrum, as a function of energy. Oscillations arise from the non-identity of the neutrino flavor and mass eigenstates and are governed by six independent parameters: three mixing angles, two mass squared differences, and a CP-violating phase. By the mid-2000s, the absolute magnitudes of both mass splittings were known, and two mixing angles had been well measured. Attention turned to the last mixing angle, called θ_{13}. Earlier studies showed this parameter was small, but determining its precise size, including whether it might be zero, required a new generation of experiments.

Double Chooz and two similar reactor-based experiments were designed to measure or place an upper limit on θ_{13}, or more directly, $\sin^2 2\theta_{13}$. Tighter constraints on this parameter may help to

explain the structure of neutrino flavor mixing, which departs from the mixing pattern in quarks. Establishing that θ_{13} is not zero, and learning its magnitude, are critical steps towards observing CP violation in neutrino oscillations. A precise θ_{13} measurement is also key to determining the neutrino mass state ordering. In turn, knowledge of the mass ordering and CP symmetry, or lack thereof, in the neutrino sector may shape theories of neutrino mass generation, flavor, the baryon asymmetry, and other aspects of particle physics and cosmology.

In 2011, Double Chooz became the first reactor-based experiment to find evidence that $\sin^2 2\theta_{13}$ is nonzero. That observation, like all Double Chooz measurements to date, was made with a single liquid scintillator detector located approximately 1 km from the reactors of the Chooz Nuclear Power Plant in northeastern France. An energy-dependent deficit of electron antineutrino-like events, compared to a reactor flux simulation, indicated that $\sin^2 2\theta_{13}$ was not far below the previous experimental limits. In this analysis, as in all Double Chooz analyses, antineutrinos were detected through inverse beta decay. Both products of this interaction are observed: a positron, which annihilates promptly, and a neutron, which thermalizes before being absorbed by a nucleus. To elevate the neutron-capture gamma signal above background radioactivity, the central part of the detector is doped with Gd. In this first analysis, and in a 2012 update, all antineutrino candidates included a Gd capture–like signal.

Beginning in 2013, Double Chooz pioneered a new antineutrino detection channel involving neutron captures on hydrogen. This technique allows inverse beta decays to be observed in undoped regions of the detector, roughly tripling the signal that can be collected in a given data-taking period. With this statistical enhancement come systematic challenges. The H-based analysis involves regions of the detector that are harder to calibrate, along with potentially enormous contamination from accidental-coincidence backgrounds. Increasingly powerful strategies have been developed to overcome these obstacles.

This thesis presents the final Gd- and H-channel $\sin^2 2\theta_{13}$ measurements made using only the Double Chooz far detector. These analyses include about two years of observation, approximately twice the period analyzed in the previous set of publications and four times the period included in the first Double Chooz measurement. New techniques improve signal efficiency, background rejection, and detection-related uncertainties in both signal channels. For each dataset, a pair of Monte Carlo–based fits extract $\sin^2 2\theta_{13}$ from complementary dimensions of the signal candidates.

Beyond confirming θ_{13}-driven oscillation, these analyses reveal features of the reactor antineutrino spectrum not predicted by traditional models.

The continuous progression of Double Chooz measurements is now approaching a systematic limit. In both the Gd- and H- channel results presented here, $\sin^2 2\theta_{13}$ uncertainty is heavily dominated by uncertainty on the reactor flux model. That quantity is irreducible in a single-detector experiment. Fortunately, as its name implies, Double Chooz will not remain a single-detector experiment. The near detector, which began operation in the last weeks of 2014, will soon provide a superior constraint on the flux shape and normalization. The complete, two-detector configuration of Double Chooz will bring rapid improvements in $\sin^2 2\theta_{13}$ precision. At the same time, it will enable sensitive measurements of the antineutrino spectrum and searches for new physics.

The following chapters are divided into four parts. The first part reviews the theory of neutrino interactions and oscillations, discusses some unresolved questions, and motivates the use of reactor-generated antineutrinos as an experimental tool. The second part describes the layout of the Double Chooz experiment, relevant aspects of the reactors, and the composition of the detectors. The third part traces each step in the most recent Gd- and H-channel measurements of $\sin^2 2\theta_{13}$, all cooperative efforts of the Double Chooz collaboration. Particular focus is placed on the motivation, implementation, and outcome of the Rate+Shape oscillation fit, the main contribution of this author. Finally, the fourth part looks toward future results achievable in Double Chooz.

Part I

Theoretical and experimental context

Chapter 2

Neutrinos in the Standard Model

Several important neutrino properties are described by the Standard Model (SM) of particle physics, currently the most complete theory of fundamental particles and forces. This chapter describes the incorporation of neutrinos into the SM and reviews relevant features of the model.

2.1 Origins of neutrino theory

The existence of a light, neutral, scarcely interactive particle was first proposed by Wolfgang Pauli in 1930. In a letter to colleagues, Pauli suggested that a particle with these properties could resolve the apparent violation of energy conservation observed in nuclear beta decays [1]. Three years later, Enrico Fermi incorporated such a particle in his beta decay theory [2, 3]. Fermi's four-fermion model, involving the direct coupling of nucleons with an electron and neutrino, was an excellent match for beta decay data. However, the theory led to cross section divergences at high energies. It also neglected the possibility of parity-violating weak interactions, theorized by Tsung-Dao Lee and Chen-Ning Yang [4] and experimentally identified by Chien-Shiung Wu in the 1950s [5]. By the late 1960s, Fermi theory had evolved into the more complete Glashow-Weinberg-Salam model of weak interactions, also known as electroweak theory (reviewed in, for example, [6]). This theory, along with quantum chromodynamics (QCD), forms the SM.

2.2 Electroweak theory

Electroweak (EW) theory gives the modern description of neutrino[1] interactions. It is a unified quantum field theory of the electromagnetic and weak interactions, which become a single interaction at energies above ~ 100 GeV, corresponding to the temperature of the universe in the first $\sim 10^{-36}$ seconds after the Big Bang.

The mathematical foundation of EW theory is the SU(2) $\times$ U(1) gauge group. The SU(2) group represents weak isospin symmetry, with gauge bosons W^i $(i = 1, 2, 3)$, while the U(1) group represents weak hypercharge symmetry, with the gauge boson B. Local gauge invariance requires all of these bosons to be massless.

Particles in EW theory are described by fermion fields, ψ_α, organized in three generations. Each generation contains two flavors of quarks[2] (u_α, d_α) and two flavors of leptons, one charged[3] (l_α) and one neutral (ν_α, the neutrino). It is often convenient to decompose these fields into their left- and right-handed components: $\psi_{L(R)} = \frac{1}{2}(1 \overset{(+)}{-} \gamma_5)\psi$, where γ_5 is the product of i and the four Dirac matrices. Under weak isospin, the left-handed quarks and leptons transform as doublets, $\Psi_\alpha = (\nu_\alpha, l_\alpha)_L$ and $(u^L_\alpha, d'_\alpha)_L$. Here, $d'_\alpha = \sum_j V_{\alpha\beta} d_\beta$, where $\mathbf{V}$ is the Cabibbo-Kobayashi-Maskawa (CKM) matrix. The right-handed quarks and charged leptons are weak isospin singlets. At the time of EW formulation, neutrinos (antineutrinos) had been observed exclusively in left-handed (right-handed) helicity states, and they were assumed to be massless. As a result, only left-handed neutrino fields were included in the model. Weak isospin invariance requires all fermions to be massless.

At energies below ~ 100 GeV, the SU(2) $\times$ U(1) symmetry of EW theory is spontaneously broken by the Higgs mechanism [7, 8, 9]. In this scheme, the Higgs field, $\phi = (\phi^+, \phi^0)$, acquires a vacuum expectation value (VEV) of $\langle\phi\rangle = (0, v/\sqrt{2})$, with $v \approx 246$ GeV. The degrees of freedom of the Higgs and EW gauge bosons mix, leaving the massive $W^\pm$ and Z, massless photon, and a

[1] Hereafter, unless a distinction is made, "neutrino" indicates both neutrinos and antineutrinos.

[2] Instead of subscripts, this chapter also uses the common notation of (u, d) for the first generation of quarks, (c, s) for the second generation, and (t, b) for the third generation.

[3] Instead of subscripts, this and subsequent chapters also use the common notation of (e, ν_e) for the first generation of leptons, (μ, ν_μ) for the second generation, and (τ, ν_τ) for the third generation. In some cases, leptons antileptons are distinguished by $-$ and $+$ superscripts.

new scalar field, the Higgs boson (H). A consequence of EW symmetry breaking is the ability for fermions to acquire mass via their Yukawa couplings to the Higgs field.

After EW symmetry breaking, the EW Lagrangian can be written (adopting some of the notation of [10])[4]:

$$\mathcal{L}_{EW} = \sum_{\alpha} \overline{\psi}_{\alpha} \left(i\partial\!\!\!/ - m_{\alpha} - \frac{m_{\alpha} H}{v} \right) \psi_{\alpha} \tag{2.1a}$$

$$- e \sum_{\alpha} Q_{\alpha} \overline{\psi}_{\alpha} \gamma^{\mu} \psi_{\alpha} A_{\mu} \tag{2.1b}$$

$$- \frac{g}{2\sqrt{2}} \sum_{\alpha} \overline{\Psi}_{\alpha} \gamma^{\mu} (1 - \gamma^5)(T^{+} W_{\mu}^{+} + T^{-} W_{\mu}^{-}) \Psi_{\alpha} \tag{2.1c}$$

$$- \frac{g}{2\cos\theta_W} \sum_{\alpha} \overline{\psi}_{\alpha} \gamma^{\mu} \left((g_V)_{\alpha} - (g_A)_{\alpha} \gamma^5 \right) \psi_{\alpha} Z_{\mu} \tag{2.1d}$$

where α sums over the three generations.

Line 2.1a includes the kinetic terms ($i\overline{\psi}_{\alpha} \partial\!\!\!/ \psi_{\alpha}$), Dirac mass terms ($-m_{\alpha} \overline{\psi}_{\alpha} \psi_{\alpha}$), and Yukawa interactions of the fermion and Higgs fields ($-m_{\alpha}/v \ \overline{\psi}_{\alpha} H \psi_{\alpha}$), where m_{α} are the fermion mass parameters. The fermion masses can be expressed as $m_{\alpha} = h_{\alpha} v/\sqrt{2}$, indicating their proportionality to the Higgs VEV and coupling constants, h_{α}, whose values must be determined empirically.

Line 2.1b describes the interactions of quantum electrodynamics (QED), the domain of the unbroken U(1) group that emerges from EW symmetry breaking. The charge of the electron is given by $-e = g \sin\theta_W$, where g is the SU(2) coupling constant. Q_{α} is the electric charge of ψ_{α} divided by e. The photon field is defined as $A = B\cos\theta_W + W^3 \sin\theta_W$. Evidently, the weak mixing angle, θ_W, describes the mixing of the W^3 and B fields.

Line 2.1c describes the charged current weak interactions. Here, $T^{\pm}$ are the raising and lowering operators for weak isospin. The charged weak boson field is defined as $W^{\pm} = (W^1 \mp iW^2)/\sqrt{2}$. The appearance of the left-handed projection operator, $(1 - \gamma^5)/2$, indicates that the charged current interaction acts only on left-handed fermion fields and right-handed antifermion fields. The interaction is thus maximally parity (P) violating. The product of γ^{μ} and $(1 - \gamma^5)$ gives the interaction its characteristic $V - A$ (vector minus axial vector current) nature. In charged current quark interactions, the involvement of the CKM matrix permits flavor-changing decays. A phase in the CKM matrix allows the interaction to violate charge-parity (CP) symmetry.

[4]This chapter and the following four chapters use natural units of $\hbar = c = 1$ unless otherwise specified.

Line 2.1d describes the neutral current weak interactions. The neutral weak boson field is defined in a manner complementary to the photon, $Z = -B\sin\theta_W + W^3\cos\theta_W$. In the neutral current interaction, the $V-A$ structure is more elaborate than in the charged current case, involving ψ_α-dependent couplings: the vector coupling, $(g_A)_\alpha = T^3_\alpha - 2Q_\alpha \sin^2\theta_W$, and axial vector coupling $(g_V)_\alpha = T^3_\alpha$, where T^3 is the z-projection of weak isospin.

Table 2.1 lists selected properties of the fermions in EW theory.

2.3 Neutrino properties in EW theory

Several neutrino properties are immediately visible from the EW Lagrangian. Neutrinos participate in two basic types of interactions: charged current (CC) interactions mediated by the $W^\pm$ bosons and neutral current (NC) interactions mediated by the Z boson. The flavor of a neutrino is defined by the charged lepton participating in its CC interactions. A neutrino which undergoes CC interactions with l_α is ν_α; for example, the decay $W^- \to e^- + \nu$ produces an electron neutrino, ν_e.

Calculating the tree-level matrix element for NC and CC interactions, with the assumption that the mass of the boson, M_b, is large compared to its momentum, yields a common factor of $(\sqrt{2}/8)(g/M_b)^2$. For the W bosons, this factor is called the Fermi constant:

$$G_F = \frac{\sqrt{2}}{8}\left(\frac{g}{M_W}\right)^2 \approx 1.166 \times 10^{-5}\ \text{GeV}^{-2} \tag{2.2}$$

This expression reveals the inverse relationship between the mass of the weak gauge bosons and the strength of the weak interaction. It is the large mass of the $W^\pm$ and Z that give neutrino interactions, at least at energies typically probed in experiments, their characteristically small cross sections.

The EW Lagrangian also conveys an important point about neutrino mass in the SM. The first line of Eq. 2.1a contains the Dirac mass terms for the quarks and charged leptons:

$$\begin{aligned}\mathcal{L}_{EW} &\supset -\sum_\alpha m_\alpha \overline{\psi}_\alpha \psi_\alpha \\ &\supset -\sum_\alpha \frac{h_\alpha v}{2}\left(\overline{(\psi_\alpha)_L}(\psi_\alpha)_R + \overline{(\psi_\alpha)_R}(\psi_\alpha)_L\right)\end{aligned} \tag{2.3}$$

Because EW theory was formulated without right-handed neutrino states, $(\nu_\alpha)_R$, it is not possible to construct such a term for neutrinos. Thus, although no theoretical principle requires this condition, neutrinos are massless in the SM.

Table 2.1: Selected properties of the fermions in EW theory, including the quantum numbers for weak isospin (T) and its z-projection (T^3), weak hypercharge (Y), and electric charge ($Q = T^3 + Y/2$); approximate mass after EW symmetry breaking; and occasion of first observation [10, 11].

Particle	Chiral states	T	T^3	Y	Q	Mass (MeV)	First observation
First generation							
Up quark	u_L	$1/2$	$1/2$	$1/3$	$2/3$	2.3	SLAC
	u_R	0	0	$4/3$			(1968)
Down quark	d_L	$1/2$	$-1/2$	$1/3$	$-1/3$	4.8	SLAC
	d_R	0	0	$-2/3$			(1968)
Electron	e_L	$1/2$	$-1/2$	-1	-1	0.51	Cambridge
	e_R	0	0	-2			(1897)
Electron neutrino	$(\nu_e)_L$	$1/2$	$1/2$	-1	0	$< 2 \times 10^{-6}$	Savannah R. (1956)
Second generation							
Charm quark	c_L	$1/2$	$1/2$	$1/3$	$2/3$	1300	Brookhaven/SLAC
	c_R	0	0	$4/3$			(1974)
Strange quark	s_L	$1/2$	$-1/2$	$1/3$	$-1/3$	95	Manchester
	s_R	0	0	$-2/3$			(1947)
Muon	μ_L	$1/2$	$-1/2$	-1	-1	110	Caltech/Harvard
	μ_R	0	0	-2			(1937)
Muon neutrino	$(\nu_\mu)_L$	$1/2$	$1/2$	-1	0	$< 2 \times 10^{-6}$	Brookhaven (1962)
Third generation							
Top quark	t_L	$1/2$	$1/2$	$1/3$	$2/3$	170000	Fermilab
	t_R	0	0	$4/3$			(1995)
Bottom quark	b_L	$1/2$	$-1/2$	$1/3$	$-1/3$	4200	Fermilab
	b_R	0	0	$-2/3$			(1977)
Tau	τ_L	$1/2$	$-1/2$	-1	-1	1800	SLAC
	τ_R	0	0	-2			(1976)
Tau neutrino	$(\nu_\tau)_L$	$1/2$	$1/2$	-1	0	$< 2 \times 10^{-6}$	Fermilab (2000)

Chapter 3

Neutrino oscillations

In the 1960s, experiments began to suggest that the SM assumption of massless neutrinos is incorrect. Evidence came in the form of neutrino flavor oscillations, a phenomenon possible only if the mass of at least one neutrino state is nonzero. This chapter lays out the formalism of neutrino oscillations and summarizes the experimental indications of their existence.

3.1 Vacuum oscillations in a three-neutrino model

Bruno Pontecorvo proposed the idea of neutrino oscillations in 1957 [12], in analogy to the recently studied mixing of neutral kaons [13]. Ziro Maki, Masami Nakagawa, and Shoichi Sakata developed the concept with more formality in 1962 [14]. Oscillations follow from supposing that the neutrino mass states (the eigenstates of the vacuum propagation Hamiltonian) are not identical to the flavor states that participate in weak interaction. The mixing of the flavor states, $|\nu_\alpha\rangle$ ($\alpha = e, \mu, \tau$), and mass states, $|\nu_i\rangle$ ($i = 1, 2, 3$), can be described by a unitary matrix $\mathbf{U}$. The α-flavored state of a neutrino or antineutrino can then be written:

$$\begin{aligned} |\nu_\alpha\rangle &= \sum_{i=1}^{3} U^*_{\alpha i} |\nu_i\rangle \\ |\bar{\nu}_\alpha\rangle &= \sum_{i=1}^{3} U_{\alpha i} |\bar{\nu}_i\rangle \end{aligned} \tag{3.1}$$

The matrix $\mathbf{U}$ is often called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. It describes the rotation between the $n = 3$ dimensional mass and flavor bases, so it is conveniently

parametrized by $n(n-1)/2 = 3$ angles and $n(n+1)/2 = 6$ phases. Not all of these phases are physically meaningful, since some can be absorbed into the definitions of the lepton fields. The number of physical phases depends on the nature of neutrino mass, to be discussed in Ch. 4. If neutrinos are Dirac particles, the number of physical phases is $(n-1)(n-2)/2 = 1$. An additional $n-1 = 2$ physical phases appear in $\mathbf{U}$ if neutrinos are Majorana particles, but as shown below, these phases do not affect oscillations.

With the mixing angles labeled θ_{12}, θ_{13}, and θ_{23}, the Dirac phase denoted δ, and the Majorana phases denoted α_1 and α_2, $\mathbf{U}$ can be written as the product of three rotations and a diagonal matrix containing the Majorana phases:

$$\begin{aligned}\mathbf{U} = &\begin{bmatrix}1 & 0 & 0\\ 0 & \cos\theta_{23} & \sin\theta_{23}\\ 0 & -\sin\theta_{23} & \cos\theta_{23}\end{bmatrix}\begin{bmatrix}\cos\theta_{13} & 0 & \sin\theta_{13}e^{-i\delta}\\ 0 & 1 & 0\\ -\sin\theta_{13}e^{i\delta} & 0 & \cos\theta_{13}\end{bmatrix}\begin{bmatrix}\cos\theta_{12} & \sin\theta_{12} & 0\\ -\sin\theta_{12} & \cos\theta_{12} & 0\\ 0 & 0 & 1\end{bmatrix}\\ &\times\begin{bmatrix}e^{i\alpha_1/2} & 0 & 0\\ 0 & e^{i\alpha_2/2} & 0\\ 0 & 0 & 1\end{bmatrix}\end{aligned} \tag{3.2}$$

or, explicitly identifying each component, suppressing the Majorana phases, and abbreviating $\sin\theta_{ij}$ ($\cos\theta_{ij}$) as s_{ij} (c_{ij}):

$$\mathbf{U} = \begin{bmatrix}U_{e1} & U_{e2} & U_{e3}\\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3}\\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3}\end{bmatrix} = \begin{bmatrix}c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta}\\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13}\\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13}\end{bmatrix} \tag{3.3}$$

Flavor state mixing has profound impact on neutrino propagation. A complete discussion requires wave packet formalism or a field theory approach including details of the neutrino production and detection. The following paragraphs present a simplified treatment, in the style of [10], which arrives at the correct result.

Consider an initially pure α-flavored neutrino state traveling in the vacuum. Using Eq. 3.1 and the standard time propagator of relativistic quantum mechanics, the time evolution of the this state can be expressed as:

$$|\nu_\alpha(t)\rangle = \sum_i^3 e^{-i(E_i t - \mathbf{p}_i\cdot\mathbf{x})} U^*_{\alpha i}|\nu_i\rangle \tag{3.4}$$

where E_i and $\mathbf{p}_i$ are, respectively, the energy and momentum of the i^{th} mass state. Using the complex conjugate of Eq. 3.1, Eq. 3.4 can be written:

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_i^3 e^{-i(E_i t - \mathbf{p}_i \cdot \mathbf{x})} U^*_{\alpha i} U_{\beta i} \right) |\nu_\beta\rangle \tag{3.5}$$

From basic quantum mechanics, the probability that the neutrino will be observed in the β flavor at some particular time T is:

$$P_{\nu_\alpha \to \nu_\beta}(T) = |\langle \nu_\beta | \nu_\alpha(T) \rangle|^2 \tag{3.6a}$$

$$= \left| \sum_{\beta=e,\mu,\tau} \left(\sum_i^3 e^{-i(E_i T - p_i L)} U^*_{\alpha i} U_{\beta i} \right) \right|^2 \tag{3.6b}$$

where $L = \hat{\mathbf{k}} \cdot (\mathbf{x}(T) - \mathbf{x}(0))$ is the distance traveled by the neutrino, along the unit vector $\hat{\mathbf{k}}$, between $t = 0$ and $t = T$. Evaluating the square in Line 3.6b introduces phases of the form:

$$\delta\varphi_{ij} = (E_i - E_j)T - (p_i - p_j)L \tag{3.7a}$$

$$= (E_i - E_j)\left(T - \frac{E_i + E_j}{p_i + p_j} L \right) + \frac{m_i^2 - m_j^2}{p_i + p_j} L \tag{3.7b}$$

where Line 3.7b makes use of the relationship $E_i^2 = p_i^2 + m_i^2$. Line 3.7b can be simplified in a variety of ways, based on different assumptions about the propagation of the mass states. A simple assumption is that these states have a common energy, so that $E_\nu = E_i = E_j$ for all i, j. This assumption eliminates the first term of Line 3.7b. At the ultrarelativistic energies observed in current experiments, $E_\nu \gg m$ allows the approximation $p_i = p_j = E_\nu$. Line 3.7b then becomes:

$$\delta\varphi_{ij} = \frac{m_i^2 - m_j^2}{2E_\nu} L = \frac{\Delta m_{ij}^2 L}{2E_\nu} \tag{3.8}$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$, often called the mass splitting between the i^{th} and j^{th} states. Manifestly, in the three-neutrino model, only two of the three mass splittings are independent.

Following the stated assumptions, Eq. 3.6 for neutrinos (antineutrinos) evaluates to:

$$P_{\nu_\alpha \to \nu_\beta (\bar{\nu}_\alpha \to \bar{\nu}_\beta)} = \delta_{\alpha\beta} - 4 \sum_{i>j} \mathrm{Re}(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E_\nu} \right) \tag{3.9a}$$

$$\overset{+}{(-)}\, 2 \sum_{i>j} \mathrm{Im}(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}) \sin \left(\frac{\Delta m_{ij}^2 L}{2E_\nu} \right) \tag{3.9b}$$

This expression demonstrates that a neutrino of energy E_ν, created in the flavor α, has a nonzero probability of being detected as flavor β at a distance L from its creation, provided that:

1. **U** differs from the identity matrix; that is, the neutrino flavor and mass states are not identical.

2. At least one Δm^2_{ij} is nonzero, implying that at least one m_i is nonzero.

In the case of $\alpha \neq \beta$, Eq. 3.9 may be called an appearance probability, as it represents the probability of ν_β appearing from a source of ν_α. For $\alpha = \beta$, Eq. 3.9 is often called a survival probability. A disappearance probability is the complement of a survival probability: $P_{\nu_\alpha \to \nu_x} = 1 - P_{\nu_\alpha \to \nu_\alpha}$, where $x \neq \alpha$.

The sign variation in Line 3.9b indicates that the probabilities of neutrinos and antineutrinos can differ if **U** is complex. In particular, **U** must contain $\delta \neq 0$, as Majorana phases drop out of Eq. 3.9. Note that a neutrino-antineutrino difference can occur only in appearance probabilities. For disappearance probabilities, the neutrino and antineutrino cases are identical. This equivalence is a consequence of charge-parity-time (CPT) invariance.

Although Eq. 3.9 is widely applicable to neutrinos propagating in the vacuum, it is not universal. The validity of this and other oscillation formulas depends on the quantum coherence of the neutrino mass states. If the uncertainty on the neutrino momentum is too small at production or detection ($\Delta p \lesssim \hbar/2L$), or if the neutrino travels too great a distance ($L \gtrsim 2E_\nu\sigma/|\Delta m^2|$, where σ describes the spatial resolution of the detection process), this coherence is lost (see, for example, [15]). Virtually all neutrino detection schemes and experimentally accessible neutrino sources escape these conditions. Neutrinos produced in supernovae are a possible exception [16]. The remainder of this work discusses only neutrino experiments which meet the coherence criteria. The vacuum oscillation probabilities derived here receive some modifications for neutrinos propagating in dense material.

3.2 Two-neutrino approximation

If the absolute values of the two mass splittings differ greatly, the formalism in Sec. 3.1 can often be approximated by a two-neutrino model. Such an approximation includes only one mass splitting and one mixing angle. The operative mass splitting is the one for which $|\Delta m^2_{ij}| \sim \langle E_\nu \rangle / \langle L \rangle$, where $\langle L \rangle$ and $\langle E_\nu \rangle$ are the average baseline and energy of the neutrinos under consideration. A much smaller mass splitting, $|\Delta m^2_{ij}| \ll \langle E_\nu \rangle / \langle L \rangle$, can be neglected because the associated oscillations are too low in frequency to significantly affect probabilities. In some cases, a much larger mass splitting can be neglected: if $\theta_{ji} \ll 1$, mass splittings of magnitude $|\Delta m^2_{ij}| \gg \langle E_\nu \rangle / \langle L \rangle$ can be

ignored because the associated oscillations average out to a small, L/E_ν–independent change in the oscillation probability equal to $\frac{1}{2}(\sin^2 2\theta_{ji}) \ll 1$.

Experiments indicate $|\Delta m^2_{21}| \ll |\Delta m^2_{31}|$ (see values in Tab. 3.1), so the two-neutrino model is frequently useful. For instance, using Eq. 3.9 and Eq. 3.3, the survival probability of electron neutrinos or antineutrinos with $\langle E_\nu \rangle \sim 1$ MeV at baselines of $\langle L_{short} \rangle \sim 10^3$ m can be approximated by:

$$P_{\nu_e \to \nu_e (\bar{\nu}_e \to \bar{\nu}_e)} = 1 - \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m^2_{31} L}{4E_\nu} \right) \tag{3.10}$$

because $|\Delta m^2_{31}| \sim \langle E_\nu \rangle / \langle L_{short} \rangle$, while $|\Delta m^2_{21}| \ll \langle E_\nu \rangle / \langle L_{short} \rangle$. Meanwhile, the survival of electron neutrinos or antineutrinos of the same energy at baselines of $\langle L_{long} \rangle \sim 10^5$ m is effectively approximated by:

$$P_{\nu_e \to \nu_e (\bar{\nu}_e \to \bar{\nu}_e)} = 1 - \sin^2 2\theta_{12} \sin^2 \left(\frac{\Delta m^2_{21} L}{4E_\nu} \right) \tag{3.11}$$

because $|\Delta m^2_{21}| \sim \langle E_\nu \rangle / \langle L_{long} \rangle$, $|\Delta m^2_{31}| \gg \langle E_\nu \rangle / \langle L \rangle$, and $\theta_{13} \ll 1$. Figure 3.1 illustrates the regimes in which Eq. 3.10 and Eq. 3.11 are reasonable approximations.

Similarly, at baselines of $\langle L_{long} \rangle \sim 10^5$ m, the disappearance probability of muon neutrinos or antineutrinos is approximately:

$$P_{\nu_\mu \to \nu_\mu (\bar{\nu}_\mu \to \bar{\nu}_\mu)} = 1 - \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m^2_{32} L}{4E_\nu} \right) \tag{3.12}$$

If factors of $\hbar$ and c are restored, the expression $\sin^2 \left(\Delta m^2_{ij} L / 4E_\nu \right)$ can be written as $\sin^2 \left(1.267 \Delta m^2_{ij} L / E_\nu \right)$, where Δm^2_{ij} has units of eV^2, L has units of meters, and E_ν has units of MeV.

3.3 Observations of neutrino oscillations

In the last five decades, multiple experiments have shown that the scheme of Sec. 3.1 is realized in nature. These experiments observed neutrinos from a variety of sources: solar neutrinos, originating in fusion reactions in the core of the sun; atmospheric neutrinos, from the decay of mesons created when cosmic rays strike nuclei in the upper atmosphere of the earth; accelerator neutrinos, produced by meson decays from a high-energy proton beam impinging upon a target; and reactor neutrinos, generated in the beta decay of nuclear fission products. In many cases, the two-neutrino model of Sec. 3.2 has proved a good approximation, allowing the mixing matrix $\mathbf{U}$ to be explored through

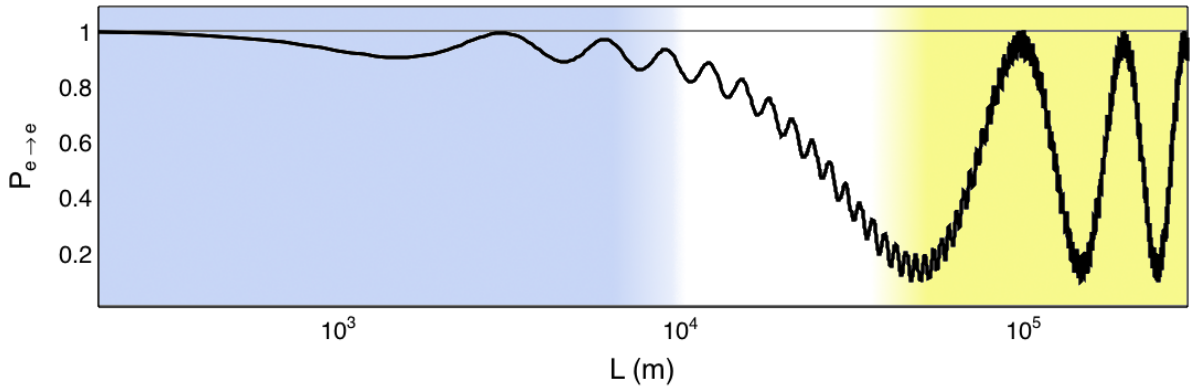

Figure 3.1: The survival probability for electron neutrinos or antineutrinos with $E_\nu = 3$ MeV (close to the mean energy of electron antineutrinos from reactors) over a range of baselines covering $L \sim 10^3 - 10^5$ m, computed in the three-neutrino model. The blue (yellow) region roughly indicates the range of L in which a two-neutrino model involving $\sin^2\theta_{13}$ and $|\Delta m^2_{31}|$ ($\sin^2\theta_{12}$ and $|\Delta m^2_{21}|$) is a reasonable approximation. The curve is drawn with the global best-fit oscillation parameters from [10].

the three sub-matrices shown in Eq. 3.2. The following sections highlight some of the major experimental results in each of the three sectors.

3.3.1 Measurements in the θ_{12} sector

Neutrino oscillations were first observed in the late 1960s in the Homestake experiment led by Raymond Davis, Jr. and John Bahcall. The experiment consisted of a tank of tetrachloroethene installed deep in a South Dakota gold mine in order to limit cosmic ray backgrounds. Atoms of ^{37}Cl in the fluid served as a target for CC interactions of electron neutrinos. The resulting ^{37}Ar atoms could be collected and counted. Davis and Bahcall sought to measure the flux of electron neutrinos generated by fusion reactions in the the sun and to compare that measurement to the prediction of the Standard Solar Model. They measured about one third of the predicted flux, an anomaly which persisted through decades of observation [17, 18].

In the 1990s, the Sudbury Neutrino Observatory (SNO) made complementary measurements of the solar neutrino flux. Using heavy water as the detection medium, SNO was able to identify both

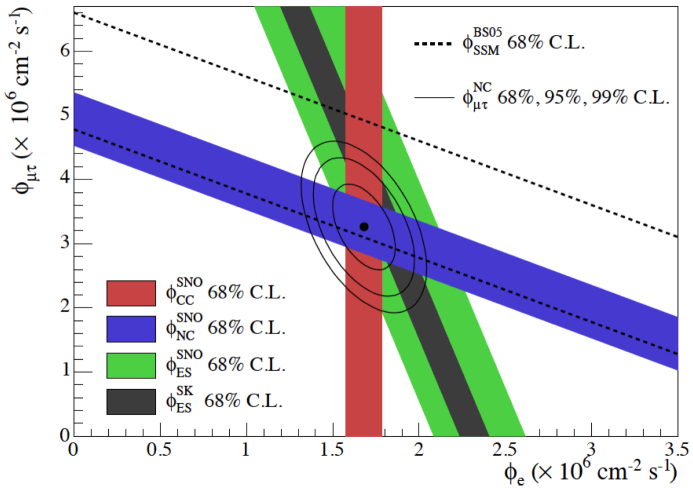

Figure 3.2: Measurements of the the flux of muon and tau neutrinos ($\phi_{\mu\tau}$) versus the flux of electron neutrinos (ϕ_e) made by the Sudbury Neutrino Observatory (SNO) and Super-Kamiokande (SK) via the CC, NC, and neutrino-electron elastic scattering (ES) channels [20]. The point marks the SNO best fit from the CC and NC channels. The dashed lines indicate the total neutrino flux predicted by the Standard Solar Model. The nonzero value of $\phi_{\mu\tau}$ is evidence for solar neutrino oscillations.

CC and NC interactions. At solar neutrino energies, $E_\nu \lesssim 10$ MeV, only electron neutrinos can participate in CC interactions, while all flavors can participate in thresholdless NC interactions. Through these two channels, SNO showed that the total flux of solar neutrinos agrees with theoretical predictions and that the ratio of electron neutrinos to all flavors is approximately one third [19]. This 2001 discovery established neutrino oscillations as the explanation for the discrepancy observed in the Homestake experiment.

Solar neutrino oscillations are more complicated than the vacuum oscillations described in Sec. 3.1. The extremely dense environment of the solar interior leads to significant neutrino-electron interactions, and these couplings alter the energy levels of the mass states. The involvement of this process, the Mikheyev-Smirnov-Wolfenstein (MSW) effect [21, 22], introduced some degeneracy to the interpretation of the Homestake and SNO observations. Later experiments showed that the survival probability for solar electron neutrinos with $E_\nu \gtrsim 2$ MeV is $P_{\nu_e \to \nu_e} \approx \sin^2\theta_{12}$. Consequently, the flavor composition measurements from SNO and the Super-Kamiokande (Super-K) experiment, shown in Fig. 3.2, provide constraints on θ_{12}. Because of its centrality in solar neutrino oscillations,

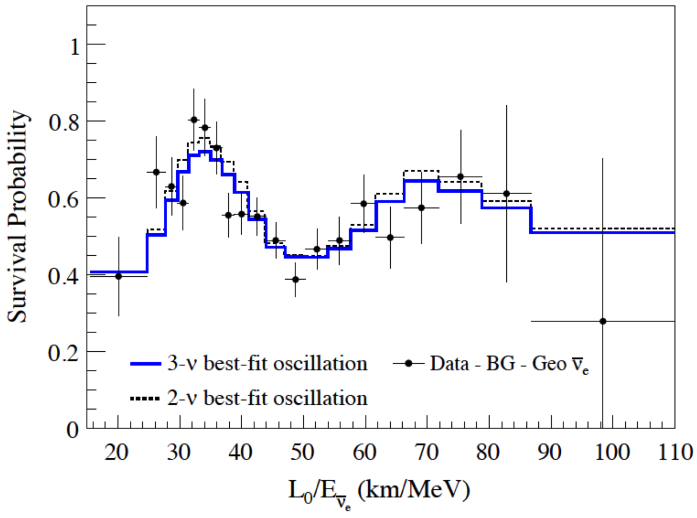

Figure 3.3: The ratio of background-subtracted electron antineutrino candidates observed in KamLAND to the no-oscillation prediction, as a function of L_0/E_ν, where $L_0 = 180$ km is the average reactor-detector distance and E_ν is the neutrino energy [23]. The L_0/E_ν–dependent deficit is due to θ_{12}-driven oscillation. The dashed black curve is the best fit to the data in a two-neutrino model; the solid blue curve is the best fit in a three-neutrino model.

θ_{12} is often called the solar mixing angle.

Oscillations driven by θ_{12} have also been studied in the Kamiokande Liquid Scintillator Antineutrino Detector (KamLAND), which shares many features with the Double Chooz detector. KamLAND is located in a mine in central Japan, an average of $L \sim 180$ km from the approximately 50 nuclear reactors spread along the Japanese coastline. For electron antineutrinos at typical reactor energies, $E_\nu \sim$ 1–10 MeV, the disappearance probability at that baseline is well described by Eq. 3.11. Figure 3.3 illustrates the prominent effect of this disappearance on the electron antineutrino flux observed in KamLAND. Measurements from KamLAND constrain Δm^2_{21} as well as θ_{12}.

3.3.2 Measurements in the θ_{23} sector

In addition to the solar neutrino measurements shown in Fig. 3.2, Super-K made the first observation of atmospheric neutrino oscillations. Atmospheric neutrinos are created in the electron and muon flavors through cosmogenic pion and kaon decays, with energies ranging from the MeV to multi-TeV scale. Because they are produced all around the earth, which is essentially transparent to their

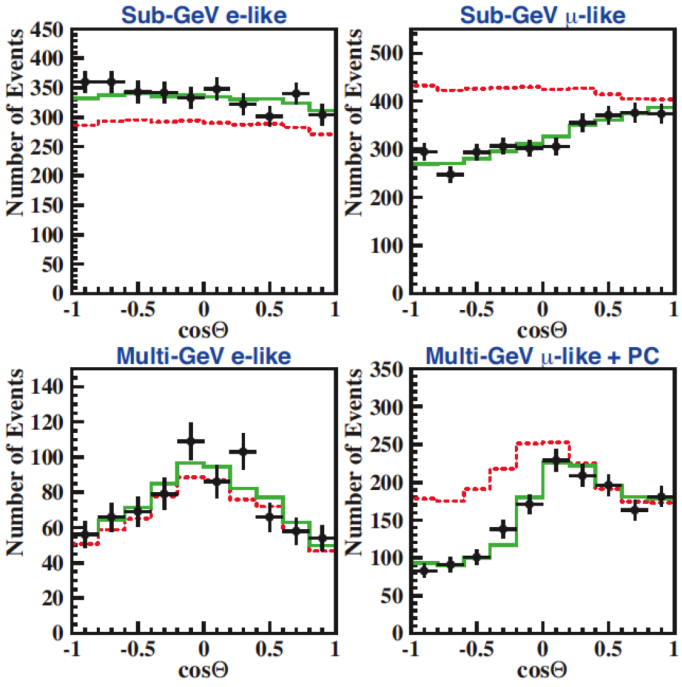

Figure 3.4: Super-K observations of electron (left) and muon (right) neutrino events as a function of zenith angle. The dashed red histogram is the nominal prediction, assuming no oscillations. The green histogram is the best fit to the data. Zenith angle-dependent disappearance is clearly visible in the muon neutrino distributions. The angle-independent shift in the electron neutrino distributions comes from a fit for systematic normalization effects. Figure from [10], as provided by the Super-K collaboration.

propagation, atmospheric neutrinos arriving at a particular location have traveled a wide range of baselines. Super-K detected atmospheric neutrinos in the $E_\nu \sim 0.1$–10 GeV range, covering baselines of $L \sim 15$–13000 km. By detecting the Cherenkov light from the charged lepton products of CC interactions, the experiment was able to reconstruct the incident neutrino direction for each event. Figure 3.4 shows the Super-K observations of muon and electron neutrinos as a function of zenith angle, a proxy for L. The distribution exhibits a pronounced, zenith angle-dependent deficit of muon neutrinos. Along with the SNO measurements, this 1998 discovery provided conclusive evidence that neutrinos oscillate [24].

Super-K analyzed oscillations in the two-neutrino model of Eq. 3.12. The parameters θ_{23} and Δm^2_{32} are often called the atmospheric mixing parameters. Accelerator-based experiments, including

KEK to Kamiokande (K2K) and Tokai to Kamioka (T2K) in Japan and the Main Injector Neutrino Oscillation Search (MINOS) at Fermilab, have also measured these parameters. These experiments project beams consisting primarily of muon neutrinos or antineutrinos, with $E_\nu \sim 1$ GeV, over baselines of $L \sim$ 300–700 m. Oscillations driven by θ_{23} cause energy-dependent deficits in the detected muon neutrino spectrum.

3.3.3 Measurements in the θ_{13} sector

In the late 1990s, the Palo Verde and CHOOZ experiments began to observe electron antineutrinos at $L \sim 1$ km from commercial nuclear reactors in Arizona and France, respectively. Oscillations at their baselines would be described by 3.10. In fact, neither experiment found a significant deficit of events. The left plot of Fig. 3.5 depicts the near-unity ratio of observed to predicted candidates in CHOOZ. Failure to observe an L/E_ν–dependent disappearance led Palo Verde and CHOOZ to place upper limits on the corresponding oscillation amplitude, $\sin^2 2\theta_{13}$ [25, 26]. CHOOZ produced the stronger limit: $\sin^2 2\theta_{13} < 0.15$ at 90% confidence level (CL) near the presently known value of Δm^2_{31} [26]. The right plot of Fig. 3.5 shows the oscillation parameter space excluded by CHOOZ.

Detecting oscillations below the CHOOZ limit required more precise experiments. The current generation of reactor experiments continues to probe the CHOOZ channel of electron antineutrino disappearance, now with far greater sensitivity. Along with Double Chooz, these initiatives include the Daya Bay experiment in China and the Reactor Electron Antineutrino Observatory (RENO) in South Korea. Accelerator-based projects like T2K and the NuMI Off-axis ν_e Appearance (NOνA) experiment at Fermilab pursue a second signal channel, the appearance of electron neutrinos in a muon neutrino beam. The probability of that appearance is more complicated than the simple two-neutrino formula of reactor antineutrino disappearance. In full, the muon-to-electron neutrino oscillation probability is:

$$\begin{aligned} P_{\nu_\mu \to \nu_e} = & \sin^2\theta_{23}\ \sin^2 2\theta_{13}\ \frac{\sin^2\left((1-x)\Delta_{31}\right)}{(1-x)^2} \\ & + y^2 \cos^2\theta_{23} \sin^2 2\theta_{12} \frac{\sin^2(x\Delta_{31})}{x^2} \\ & + y \sin 2\theta_{12}\ \sin 2\theta_{13} \sin 2\theta_{23} \frac{\sin^2(x\Delta_{31})}{x^2} \frac{\sin^2\left((1-x)\Delta_{31}\right)}{(1-x)^2} \\ & \times (\cos\Delta_{31}\cos\delta - \sin\Delta_{31}\sin\delta) \end{aligned} \tag{3.13}$$

where $y = \Delta m^2_{21}/\Delta m^2_{31}$ and $\Delta_{31} = \Delta m^2_{31} L/4E_\nu$. For long baselines, the MSW effect enters in a manner parametrized by $x = 2\sqrt{2} G_F n_e E_\nu/\Delta m^2_{31}$, where n_e is the average electron density of the earth that the beam crosses. This parameter-rich oscillation probability presents many measurement opportunities, but it also creates ambiguity in measurements of $\sin^2 2\theta_{13}$. Substantial neutrino production and detection uncertainties, combined with low signal statistics, further limit the sensitivity of the electron appearance channel.

Nonetheless, T2K provided the first positive indication of the value of $\sin^2 2\theta_{13}$. In June 2011, the experiment announced 2.5σ evidence of electron neutrino appearance [27]. This measurement is consistent with $0.03 < \sin^2 2\theta_{13} < 0.28$ at 90% CL, assuming $\delta = 0$ and the normal mass hierarchy, as defined in Sec. 3.4. Given the moderate baseline of T2K, this measurement is not significantly affected by degeneracies from the MSW effect.

During the following year, the three reactor-based $\sin^2 2\theta_{13}$ experiments supplied further evidence of a nonzero θ_{13}. Because these experiments are insensitive to δ, the mass hierarchy, and mixing angles beyond θ_{13}, and because their systematic uncertainties can be tightly controlled, they can achieve much greater precision than accelerator-based measurements. Double Chooz reported its first oscillation measurement, $\sin^2 2\theta_{13} = 0.086 \pm 0.041(\text{stat}) \pm 0.030(\text{syst})$, in December 2011 [28]. Daya Bay and RENO released consistent measurements in March 2012 and April 2012, respectively [29, 30]. Since then, all three reactor experiments and T2K have reported increasingly precise measurements of $\sin^2 2\theta_{13}$. The global best fit now stands at $\sin^2 2\theta_{13} = 0.093 \pm 0.008$, with the no-oscillation hypothesis disfavored at more than 10σ [10]. Chapter 6 describes reactor-based experiments in more detail.

3.4 Experimental status of the three-neutrino model

Figure 3.6 summarizes the results of many neutrino oscillation experiments. Remarkably, all but one of these findings are consistent with the same trio of solutions in (θ, Δm^2) space. Evidently, neutrinos from a variety of sources, across at least four decades of energy, provide strong support for the three-neutrino model of flavor oscillations. Table 3.1 gives the current global best-fit values for the mixing angles and mass splittings in this model. The single anomalous result in Fig. 3.6 may hint at new physics beyond the three-neutrino scheme, a possibility introduced in Ch. 5.

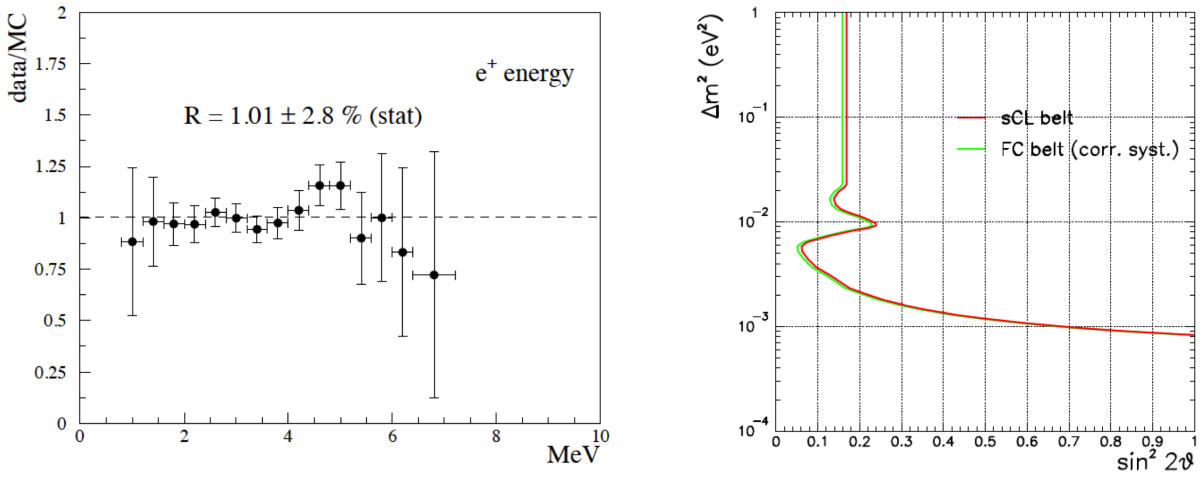

Figure 3.5: Left: The ratio of background-subtracted electron antineutrino candidates to the no-oscillation prediction in CHOOZ. The consistency with unity shows that $\sin^2 2\theta_{13}$-driven oscillations, if present, are very small in magnitude. Right: The region (above and to the right of the curves) in ($\sin^2 2\theta_{13}$, Δm^2_{31}) space excluded at 90% CL in CHOOZ. The two curves use different frequentist procedures to obtain limits [26].

Despite much experimental progress, some aspects of the three-neutrino model remain obscure. One question is the ordering of the neutrino mass states. The structure of solar neutrino oscillations establishes that $m_2 > m_1$. However, atmospheric and accelerator neutrino oscillations have not revealed the sign of Δm^2_{32}. Similarly, reactor experiments are insensitive to the sign of Δm^2_{31}. The two-neutrino examples of Eq. 3.10 and Eq. 3.11 illustrate the reason: the dominant terms in oscillation probabilities are invariant with respect to the sign of the mass splitting. If $m_3 > m_2$, the neutrino mass ordering would resemble the ordering of the quark and charged lepton masses; this arrangement is called the normal hierarchy. The inverted hierarchy has $m_3 < m_2$. Figure 3.7 portrays the two possibilities.

Details of the mixing matrix $\mathbf{U}$ are also incompletely known. None of the phases have been measured. The Dirac phase, δ, affects oscillation probabilities under conditions which are just beginning to become testable. Another outstanding question is the octant of θ_{23}, meaning whether the angle is smaller or larger than $\pi/4$. The nearly maximal magnitude of θ_{23} seems to hint at some deeper flavor structure, such as a hidden μ-τ symmetry. Precisely measuring this angle may

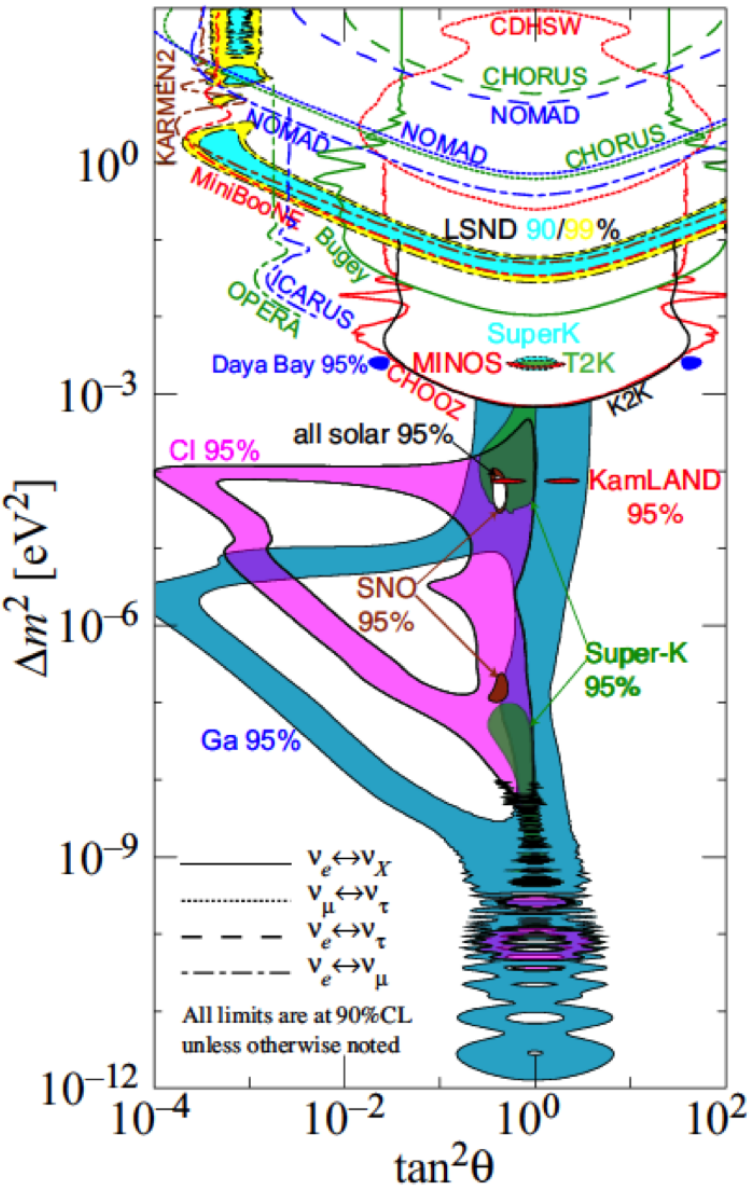

Figure 3.6: Regions in two-neutrino oscillation parameter space favored or excluded by a variety of experiments [10]. Color-filled regions are favored space; other curves are the boundaries of excluded space. Note that all but one experiment are consistent with a single set of three solutions: $(\tan^2\theta_{12},\ \Delta m^2_{21}) \sim (10^0, 10^{-4}\ \text{eV}^2)$, $(\tan^2\theta_{23},\ \Delta m^2_{32}) \sim (10^0, 10^{-3}\ \text{eV}^2)$, and $(\tan^2\theta_{13},\ \Delta m^2_{31}) \sim (10^{-2}, 10^{-3}\ \text{eV}^2)$.

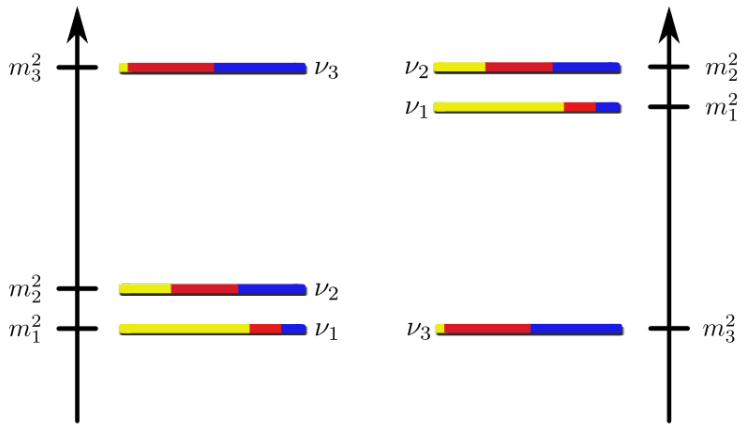

Figure 3.7: Two possibilities for the neutrino mass ordering: the normal hierarchy (left) and inverted hierarchy (right). The vertical axis is squared mass. Colors reflect the flavor content of each mass state: electron (yellow), muon (red), and tau (blue). Image from [31].

therefore be especially useful for understanding the pattern of neutrino masses and mixing. More concretely, uncertainty about the θ_{23} octant is a source of degeneracy in measurements of δ.

The discovery that $\sin^2 2\theta_{13}$ is nonzero, and is in fact large compared to the CHOOZ limit, makes all of these questions easier to answer. In general, experimental signatures of the mass hierarchy, CP violation, and the octant are all proportional to the magnitude of $\sin^2 2\theta_{13}$. On a more basic level, all of the mixing angles, including θ_{13}, must be nonzero for CP violation to exist in neutrino oscillations. These connections make measurements of $\sin^2 2\theta_{13}$ imperative for the continued exploration of neutrino properties.

Table 3.1: Current global best-fit values for the mixing angles and two independent mass splittings in the three-neutrino model [10]. Values marked NH result from assuming that the neutrino mass hierarchy is normal, as described in Ch. 4; values marked IH result from assuming that the mass hierarchy is inverted.

Parameter	Value	
$\sin^2 2\theta_{12}$	0.846 ± 0.021	
Δm^2_{21}	$(7.53 \pm 0.18) \times 10^{-5}$ eV^2	
$\sin^2 2\theta_{23}$	$0.999^{+0.001}_{-0.018}$	(NH)
	$0.1000^{+0.000}_{-0.017}$	(IH)
Δm^2_{32}	$(2.44 \pm 0.06) \times 10^{-3}$ eV^2	(NH)
	$-(2.52 \pm 0.07) \times 10^{-3}$ eV^2	(IH)
$\sin^2 2\theta_{13}$	0.093 ± 0.008	

Chapter 4

Neutrino mass

Flavor oscillations are a decisive sign that neutrinos have mass. Other observations, reviewed at the beginning of this chapter, indicate that the neutrino masses must be exceptionally small. To accommodate massive neutrinos of any magnitude, the SM requires an extension. Certain extensions are appealing because they naturally explain the extreme lightness of neutrinos. As noted at the end of this chapter, such models may also have broader consequences.

4.1 Experimental limits

Oscillations are sensitive to the squared differences between the neutrino masses, Δm^2_{ij}, but not to the absolute scale of those masses. The larger of the mass splittings does provide one lower bound: either m_2 or m_3 must exceed $\sqrt{|\Delta m^2_{32}|}$. The currently known value of $|\Delta m^2_{32}|$, as given in Tab. 3.1, implies at least one $m_i \gtrsim 0.05$ eV.

Two sources offer empirical upper bounds on the neutrino masses. Near their endpoints, beta decay spectra are minutely shaped by the effective mass of the electron antineutrino. Huge electron spectrometers have been built to examine this effect. Presently, the most restrictive limit comes from the Troitsk experiment: $m_{\bar{\nu}_e} < 2.05$ eV at 95% CL, where $m_{\bar{\nu}_e} = \sqrt{\sum_i^3 |U_{ei}|^2 m_i^2}$ [32]. Cosmology also provides limits on the sum of all neutrino masses, although these can be substantially model-dependent. Combining observations of the cosmic microwave background (CMB) and baryon acoustic oscillations, the Planck collaboration has estimated $\sum_i^3 m_i < 0.66$ eV at 95% CL [33]. Assumptions behind this estimate include the existence of exactly three massive neutrino species and

the Λ–cold dark matter (Λ–CDM) model.

4.2 Mechanisms of mass generation

4.2.1 Dirac mass

For the SM to describe nonzero neutrino mass, neutrino fields with right-handed chirality must be added to EW theory. Like the other right-handed fermions, these fields are weak isospin singlets; given $Q = 0$, the property $Y = 0$ follows. The addition of ν_R allows the formulation of a Dirac mass term for neutrinos, analogous to the terms that give mass to the charged leptons and quarks:

$$\mathcal{L}_{Dirac} = -m_D \bar{\nu}\nu = -\frac{h_\nu v}{2}(\overline{\nu_L}\nu_R + \overline{\nu_R}\nu_L) \tag{4.1}$$

For notational simplicity, the above expression and the following discussion involve only one neutrino species. Both can be generalized to include three or more.

Equation 4.1 allows neutrinos to acquire mass, and it preserves the apparent SM symmetry of lepton number conservation. Another of its features is less profitable. Identifying m_D with the largest possible m_i, roughly 2 eV, limits the strength of the neutrino-Higgs Yukawa coupling to $h_\nu \approx 2$ (2 eV)/(246 GeV) $\sim 10^{-11}$. The Yukawa coupling constant for the next-lightest fermion, the electron, is five orders of magnitude larger. This gap seems unnatural and may suggest that a different mechanism controls the neutrino masses.

4.2.2 Majorana mass

If right-handed neutrino states are added to the SM, and if lepton number conservation is not required, it becomes possible to construct a second type of mass term:

$$\mathcal{L}_{Majorana} = -\frac{m_M}{2}(\overline{\nu_R}(\nu_R)^C + \overline{(\nu_R)^C}\nu_R) \tag{4.2}$$

with $(\nu_R)^C = C\,\overline{\nu_R}^T$ is the charge conjugate of ν_R and has opposite chirality. Note that this term is invariant under weak isospin and weak hypercharge, which would not be true for the same term built from ν_L. Like the Dirac mass term, the Majorana mass term represents the coupling of left-handed and right-handed chiral states. Distinctly, the Majorana mass term is not generally invariant under the U(1) symmetry of electromagnetism. That feature precludes the construction of Majorana mass terms for the charged fermions.

It is possible to construct a Majorana neutrino, $\nu_M = \nu_R + (\nu_R)^C$, which is self-conjugate: $\nu_M^C = \nu_M$. In an SM extension that adopts this construction, neutrinos are their own antiparticles. In this case, ostensible neutrino-antineutrino differences in oscillation probabilities are merely differences between chiral states of the same field. The Majorana mass term mixes these states, changing lepton number by $\Delta L = \pm 2$.

Because Majorana mass generation is independent of EW symmetry breaking, m_M may relate to physics at a much different energy scale. Indeed, Majorana neutrinos are common in theories of physics beyond the SM. At experimentally accessible energies, however, few observables distinguish the Dirac and Majorana cases. A possible signature is neutrinoless double beta decay, in which a nucleus undergoes the transition ${}^{Z}_{A}N \rightarrow {}^{Z+2}_{A}N + 2e^-$ without the emission of neutrinos. A number of experiments currently search for this phenomenon, possible only if neutrinos are Majorana particles.

4.2.3 Seesaw mechanisms

If both Dirac and Majorana mass terms exist, their combination can provide a natural explanation for the extremely small neutrino masses. Summing Eqs. 4.1 and 4.2 produces:

$$\begin{aligned}
\mathcal{L}_{mass} &= -m_D \left(\overline{\nu_L}\nu_R + \overline{\nu_R}\nu_L\right) - \frac{m_M}{2}\left(\overline{\nu_R}(\nu_R)^C + \overline{(\nu_R)^C}\nu_R\right) \\
&= -\frac{1}{2}\begin{pmatrix}\overline{(\nu_L)^C} & \overline{\nu_R}\end{pmatrix}\begin{pmatrix}0 & m_D \\ m_D & m_M\end{pmatrix}\begin{pmatrix}\overline{\nu_L} \\ (\nu_R)^C\end{pmatrix} + h.c. \\
&= -\frac{1}{2}\,\overline{(\eta_L)^C}\,\mathbf{M}\,\eta_L + h.c.
\end{aligned} \tag{4.3}$$

where $\eta_L = \left(\nu_L, (\nu_R)^C\right)$. In a basis that diagonalizes $\mathbf{M}$,[1] Eq. 4.3 can be expressed as:

$$\begin{aligned}
\mathcal{L}_{mass} &= -\frac{1}{2}\,\overline{(\chi_L)^C}\,\widetilde{\mathbf{M}}\,\chi_L + h.c. \\
&= -\frac{1}{2}\,\overline{\chi}\,\widetilde{\mathbf{M}}\,\chi
\end{aligned} \tag{4.4}$$

where, in the last line, $\chi = \chi_L + (\chi_L)^C$. The eigenvalues of $\widetilde{\mathbf{M}}$ are $\frac{1}{2}\left(m_M \pm \sqrt{m_M^2 + 4m_D^2}\right)$. Since m_D derives from EW symmetry breaking, its value would naturally be close to the energy scale of that process, ~ 100 GeV. In contrast, m_M likely originates in physics at much higher energies, perhaps at the grand unified theory (GUT) scale of $\sim 10^{15}$ GeV. For these widely separated values of m_D and m_M, the eigenvalues of $\widetilde{\mathbf{M}}$ can be approximated by m_M and $-m_D^2/m_M$.

[1] In particular, for some unitary matrix $\mathbf{V}$: $\chi_L = \mathbf{V}^\dagger\eta_L$; $\overline{(\chi_L)^C} = \overline{(\eta_L)^C}\mathbf{V}^*$; $\widetilde{\mathbf{M}} = \mathbf{V}^*\mathbf{M}\mathbf{V}^\dagger$.

With the χ components labeled by $\chi = (n, N)$, and the negative sign of the smaller eigenvalue absorbed into the definition of n, Eq. 4.4 can be written:

$$\mathcal{L}_{mass} = -\frac{1}{2}(\bar{n} \;\; \bar{N}) \begin{pmatrix} m_D^2/m_M & 0 \\ 0 & m_M \end{pmatrix} \begin{pmatrix} n \\ N \end{pmatrix} \tag{4.5}$$

This term describes two Majorana fermions. The lighter of the pair, n, is mainly composed of the left-handed SM neutrino and has mass $m_D^2/m_M \sim 0.01$ eV. Meanwhile, N is dominated by the new right-handed neutrino and has mass at the GUT scale. The link between these particles lends this so-called seesaw mechanism its name: as the mass of N rises, the mass of n drops. This section has described the type I seesaw model, the simplest variation on a theme which surfaced around 1980 [34, 35, 36, 37, 38].

The light particle in Eq. 4.5 shares the properties of observable neutrinos, notably their small but nonzero mass. The GUT-scale partner can also serve an important function. In the hot early universe, this particle would have been produced copiously, followed by decays through lepton and antilepton channels. If the associated Yukawa couplings contained CP-violating phases, which may have some relation to δ in the neutrino mixing matrix, decays could favor the lepton channel. The net excess of leptons could have been converted to an excess of baryons via sphaleron processes in the SM. The ultimate result could be the observed baryon asymmetry of the universe.

Chapter 5

Sterile neutrinos

Although most observations of neutrino oscillations fit the three-neutrino mixing framework, certain exceptions raise the possibility of a more expansive neutrino sector. An extended mixing model could include one or more sterile neutrinos, so called because they do not engage in normal weak interactions. The first section of this chapter outlines ways in which sterile neutrinos could modify the oscillation formalism developed in Ch. 3. The following section reviews experimental evidence for sterile neutrinos, along with observations, sometimes conflicting, that limit the parameter space they may occupy.

5.1 Phenomenology

Measurements of the invisible decay width of the Z boson constrain the number of light, active neutrino flavors to three [39]. In this context, light describes a mass of less than $M_Z/2$, and active means that the neutrino couples to the Z boson through the SM weak interaction. Other types of light neutrinos could exist if they are SM gauge singlets. Many SM extensions feature at least one such particle, a sterile neutrino. In most models, sterile neutrinos are lighter versions of the predominately right-handed neutrinos derived in Sec. 4.2.3. In other models, they may be supersymmetric partners of axions, mirror matter particles, or other constructs.

Sterile states can impact neutrino oscillations if they mix significantly with the active flavors. This effect can be parametrized by expanding the mixing matrix $\mathbf{U}$ from 3×3 to $(3+n_s)\times(3+n_s)$,

where n_s is the number of sterile states. For example, with $n_s = 1$, $\mathbf{U}$ becomes:

$$\mathbf{U} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} & U_{\mu 4} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} & U_{\tau 4} \\ U_{s1} & U_{s2} & U_{s3} & U_{s4} \end{bmatrix} \tag{5.1}$$

This $3+1$ model naturally includes new mixing angles, θ_{i4} for $i = 1, 2, 3$.

Following the logic of Ch. 3, various oscillation probabilities can be derived in $3+1$ models. In a hierarchical $3+1$ scheme, the fourth mass state is much heavier than the first three, so that $m_4 \sim \sqrt{\Delta m^2_{4i}}$, with $i = 1, 2, 3$. In this case, active-sterile mixing may be observable at much shorter baselines than standard oscillations, with $L \sim \langle E_\nu \rangle / \Delta m^2_{41}$. At these short baselines, approximate appearance and disappearance probabilities resemble the two-neutrino formulas of Sec. 3.2:

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta (\bar{\nu}_\alpha \to \bar{\nu}_\beta)} &= \sin^2 2\theta_{\alpha\beta} \sin^2 \left(\frac{\Delta m^2_{41} L}{4E_\nu} \right) \qquad (\alpha \neq \beta) \\ P_{\nu_\alpha \to \nu_\alpha (\bar{\nu}_\alpha \to \bar{\nu}_\alpha)} &= 1 - \sin^2 2\theta_{\alpha\alpha} \sin^2 \left(\frac{\Delta m^2_{41} L}{4E_\nu} \right) \end{aligned} \tag{5.2}$$

where the transition amplitudes are defined as:

$$\begin{aligned} \sin^2 2\theta_{\alpha\beta} &= 4|U_{\alpha 4}|^2 |U_{\beta 4}|^2 \qquad (\alpha \neq \beta) \\ \sin^2 2\theta_{\alpha\alpha} &= 4|U_{\alpha 4}|^2 \left(1 - |U_{\alpha 4}|^2\right) \end{aligned} \tag{5.3}$$

For instance, the short-baseline survival probability of electron antineutrinos can be approximated by:

$$P_{\bar{\nu}_e \to \bar{\nu}_e} = 1 - 4|U_{e4}|^2 \left(1 - |U_{e4}|^2\right) \sin^2 \left(\frac{\Delta m^2_{41} L}{4E_\nu} \right) \tag{5.4}$$

Using the parametrization of $\mathbf{U}$ given in [40], in which $|U_{e4}| = \sin\theta_{14}$, this probability can be written:

$$P_{\bar{\nu}_e \to \bar{\nu}_e} = 1 - \sin^2 2\theta_{14} \left(\frac{\Delta m^2_{41} L}{4E_\nu} \right) \tag{5.5}$$

Oscillation probabilities can also be derived for more general $3+1$ models in which m_4 is not necessarily much larger than the other masses. In a medium-baseline scenario ($L \lesssim 2$ km, with $E_\nu \gtrsim 2$ MeV), with the data-driven constraint of $|U_{e4}|^2 \lesssim 10^{-2}$, the electron antineutrino survival probability is approximately:

$$P_{\bar{\nu}_e \to \bar{\nu}_e} = 1 - \cos^4 \theta_{14} \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m^2_{31} L}{4E_\nu} \right) - \sin^2 2\theta_{14} \left(\frac{\Delta m^2_{41} L}{4E_\nu} \right) \tag{5.6}$$

Details of this approximation are presented in [41].

Models with $n > 1$ sterile neutrinos produce more complicated oscillation formulas. Beyond the additional mixing angles and mass splittings, these models can include extra CP-violating phases which induce differences between neutrino and antineutrino oscillations.

In addition to modifying oscillation probabilities, sterile neutrinos may affect certain astrophysical observables, such as the CMB, the power spectrum of large-scale matter distributions, and the dynamics of core collapse supernovae. A review of these possibilities is presented in [42]. The following section focuses more narrowly on potential signs of sterile neutrinos in oscillation measurements.

5.2 Evidence in oscillation experiments

5.2.1 Electron neutrino and antineutrino appearance

The first evidence for sterile neutrinos appeared in the Liquid Scintillator Antineutrino Detector (LSND) at Los Alamos National Laboratory. LSND was exposed to muon antineutrinos from pions decaying at rest, a process with $< 10^{-5}$ contamination from electron antineutrinos (mainly produced by muons which decay before being captured by a nucleus). The detector was located at $L \sim 30$ m from the source, and the neutrino energies covered approximately $E_\nu \sim 35$–50 MeV. LSND observed a large excess of electron antineutrino-like events, claiming a significance of 3.8σ [43]. Muon-to-electron flavor oscillations provide the simplest explanation, but the required $\Delta m^2 \sim 1$ eV^2 strongly disagrees with the three mass splittings known from other experiments. Figure 3.6 illustrates the incompatibility of the LSND observation with solar and atmospheric neutrino data. The hypothesis of a sterile neutrino associated with a fourth, heavier mass state relieves this conflict.

A larger liquid scintillator detector, MiniBooNE, was built at Fermilab to test the sterile neutrino interpretation of LSND. MiniBooNE was positioned $L \sim 550$ m from an accelerator-driven source of muon neutrinos or antineutrinos with $E_\nu \sim 300$–3000 MeV, allowing the same L/E_ν coverage as LSND. In neutrino mode, the MiniBooNE results solidly disfavored the LSND oscillation solutions [44]. The accelerator-based Karlsruhe Rutherford Medium Energy Neutrino (KARMEN) experiment, in the UK, and the Neutrino Oscillation Magnetic Detector (NOMAD), at CERN, also failed to observe anomalous electron antineutrino appearance. By contrast, in antineutrino mode, MiniBooNE observed an excess of electron-flavored events which agreed well with the LSND

parameter space [45].

5.2.2 Electron antineutrino disappearance

Another hint of sterile neutrinos occurs in short-baseline electron antineutrino measurements. Beginning in 2011, new calculations of electron antineutrino production in reactors suggested that the flux is 3% higher than previous estimates [46]. Most of this change came from re-evaluations of two nuclear effects in the beta decay rates of fission products. Meanwhile, a new measurement of the neutron lifetime decreased the expected cross section of the inverse beta decay reaction, the prime channel for detecting reactor antineutrinos (see Sec. 6.2 for this connection) [47]. Additionally, reactor flux predictions began to include the previously neglected antineutrino flux from long-lived fission products. In combination, these effects increased the expected rate of reactor-derived electron antineutrino interactions by 6% [46]. The new prediction clashes with more than a dozen short-baseline measurements of the reactor flux, all of which matched the original flux expectation. This discordance, called the reactor antineutrino anomaly, can be alleviated by a sterile neutrino associated with $\Delta m^2_{41} > 1$ eV^2.

More signs of sterile neutrinos may come from calibration campaigns in two solar neutrino experiments. Like the Homestake experiment, the Gallium Experiment (GALLEX) in Italy and Soviet-American Gallium Experiment (SAGE) in Russia used a radiochemical technique to detect electron neutrinos. Both experiments calibrated their detectors with high-activity ^{51}Cr and ^{37}Ar sources, which decay via electron capture and emit electron neutrinos with energies up to several hundred keV. Each experiment observed fewer neutrinos than expected from these sources, with a combined significance of 2.8σ [42]. These deficits could be explained by oscillations involving a sterile neutrino with $\Delta m^2_{41} > 1$ eV^2.

5.2.3 Additional limits and searches

While the LSND result, MiniBooNE antineutrino result, reactor anomaly, and radiochemical calibration observations provide relatively consistent evidence for a sterile neutrino associated with $\Delta m^2_{41} \gtrsim 1$ eV^2, other observations complicate this interpretation. The MiniBooNE neutrino result, along with KARMEN and NOMAD, present major conflicts in a $3+1$ model, although they could be accommodated by a $3+2$ model with CP violation. All models are challenged by the

lack of muon neutrino disappearance observations in accelerator-based experiments such as MINOS and the CERN–Dortmund–Heidelberg–Saclay–Warsaw (CDHSW) experiment at CERN. Figure 5.1 illustrates some of these tensions.

Several current and proposed experiments hope to clarify this situation. A program of short-baseline liquid argon detectors at Fermilab will provide unprecedented sensitivity to electron neutrino (and antineutrino) appearance and muon neutrino (and antineutrino) disappearance. Very short-baseline reactor experiments will search for L/E_ν–dependent signatures of electron antineutrino disappearance. Observations of high-intensity, radioactive neutrino sources may provide similar information. Many other experiments have also been suggested. Since Double Chooz and similar reactor experiments were optimized for $|\Delta m^2_{31}| \sim 2.4 \times 10^{-3}$ eV^2, they do not offer ideal sensitivity to oscillations at the scale of $|\Delta m^2| \gtrsim 1$ eV^2 favored by current sterile neutrino fits. Still, they may offer some useful limits, as recently demonstrated by Daya Bay [48]. Chapter 5 discusses prospects for sterile neutrino searches in the two-detector phase of Double Chooz.

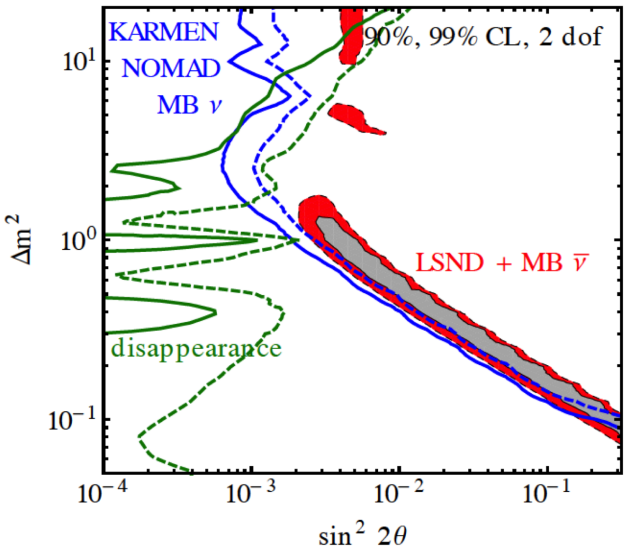

Figure 5.1: Allowed and excluded regions from various observations analyzed in a $3 + 1$ sterile neutrino model [42]. The filled curves indicate allowed regions in a combined LSND–MiniBooNE antineutrino analysis [43, 45]. The blue contours outline the region (to the upper right of the plot) excluded by a combined analysis of the KARMEN, NOMAD, and the MiniBooNE neutrino data [49, 50, 44]. The green contours outline the region (to the right of the plot) excluded by a combined analysis of disappearance channels in MINOS, CDHSW, and short-baseline reactor experiments [51, 52, 53, 54]. Note that the blue and green contours exclude most of the LSND–MiniBooNE allowed region in this model.

Chapter 6

Reactor antineutrinos as an experimental tool

Nuclear reactors have long been an experimentally important source of antineutrinos. The first observation of neutrinos of any kind occurred in 1956 at a reactor on the Savannah River Site in South Carolina. That experiment, orchestrated by Clyde Cowan and Frederick Reines, used the inverse beta decay (IBD) interaction to detect electron antineutrinos emanating from the reactor core [55]. The KamLAND, CHOOZ, and short-baseline reactor experiments mentioned in Chs. 3 and 5 have helped frame the three-neutrino mixing model; the latter may provide hints of physics beyond it. The reactor antineutrino experiments currently measuring θ_{13}, including Double Chooz, build on strategies developed in these earlier experiments. This chapter introduces the basic principles of antineutrino production, detection, and θ_{13}-driven oscillation analyses at reactors.

6.1 Reactor antineutrino generation

In a reactor core, the fissions that liberate nuclear energy also initiate a chain of antineutrino-producing reactions. When an actinide such as ^{235}U is fissioned by a thermal neutron, it usually yields two large fragments and a small number of free neutrons. The fragments are generally not alike in mass, instead following a bimodal mass distribution similar to the illustration in Fig. 6.1. Both fragments are excessively neutron-rich and undergo a series of beta decays, $n \to p + e^- + \bar{\nu}_e$, with Q values up to about 10 MeV. Two sample decay chains appear in Fig. 6.1. Given the

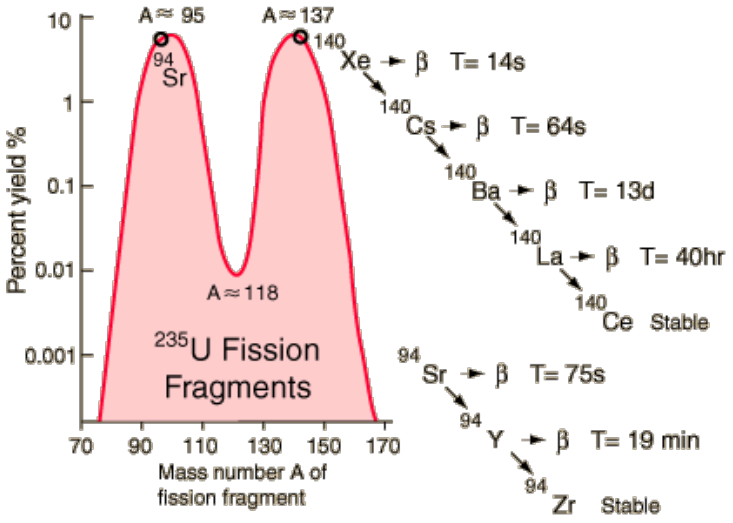

Figure 6.1: An illustration of the thermal neutron fission fragment yield of ^{235}U, along with decay chains for two of the most common fragments [56].

large variety of fission fragments in a typical reactor, many of which can decay through multiple branches, thousands of different beta decays are represented. Each fission leads to an average of six beta decays, and each of these produces one electron antineutrino. Most decays proceed rapidly, with half-lives on the order of seconds or minutes, but longer-lived products also affect the reactor dynamics and, to a lesser extent, the emitted neutrino flux.

Assuming that the beta decay rate is in equilibrium, the rate of antineutrino emission is proportional to the thermal power of the reactor. For a reactor operating at 4 $\mathrm{GW_{th}}$,[1] the antineutrino rate is on the order of:

$$R_{\bar{\nu}_e} \approx 4000\ \mathrm{MW_{th}} \cdot \frac{10^6\ \mathrm{J/s}}{\mathrm{MW}} \cdot \frac{1\ \mathrm{MeV}}{1.6 \cdot 10^{-13}\ \mathrm{J}} \cdot \frac{1\ \mathrm{fission}}{\sim 200\ \mathrm{MeV}} \cdot \frac{\sim 6\ \bar{\nu}_e}{\mathrm{fission}} \approx 10^{21}\ \bar{\nu}_e/\mathrm{s} \tag{6.1}$$

Fission fragments are emitted with no preferred direction in the reactor reference frame, and they quickly come to rest. Consequently, the antineutrino emission is isotropic.

The antineutrino flux is a sum over all beta decay branches of all fissioning isotopes in a reactor. In a commercial pressurized water reactor (PWR), such as those observed by modern antineutrino experiments, the composition of fissioning isotopes changes over time. PWRs are typically fueled

[1] Following common usage in the electrical power industry, $\mathrm{GW_{th}}$ denotes a GW of thermal power produced in the reactor, while $\mathrm{GW_e}$ denotes a GW of net electrical power.

with low-enriched uranium, with a ^{235}U content of up to 5%. Just after fueling, the majority of fissioning nuclei are ^{235}U. As the reactor cycle proceeds, ^{239}Pu is formed from neutron absorption by ^{238}U. An increasing fraction of fissions occur from ^{239}Pu, culminating near 40% at the end of a cycle. Fissions from ^{238}U and ^{241}Pu also contribute at the $\sim 10\%$ level, along with trace contributions from other isotopes. The isotope composition of a reactor core can be modeled fairly well, an important capability both academically and commercially.

Each fissioning isotope produces a different distribution of fragments, leading to a distinct rate and spectrum of emitted antineutrinos. The change in fission fractions over a reactor cycle lead to time-varying effects of order 10% in the antineutrino flux. As an aside, this phenomenon underlies one of the few practical applications so far explored in neutrino physics. By monitoring the antineutrino flux from a reactor, regulatory agents may be able to detect illicit removal of fissile material or other tampering in a civil reactor. The antineutrino detection principles outlined in the following sections may eventually find use in such nonproliferation efforts.

6.2 Inverse beta decay

In the energy range of reactor antineutrinos, $E_\nu \lesssim 10$ MeV, the most readily observable interaction is CC quasi-elastic scattering off free protons (^{1}H nuclei), often called inverse beta decay:

$$\bar{\nu}_e + p \rightarrow e^+ + n \tag{6.2}$$

Figure 6.2 shows the simplest Feynman diagram of the IBD interaction. The IBD cross section exceeds nearly all other SM cross sections for antineutrinos in this energy range. The only exception is coherent antineutrino-nucleus scattering, which produces nuclear recoils so low in energy that they have never been detected. The IBD interaction has the additional advantage of yielding two identifiable products. Searching for both the positron and neutron, in close coincidence, allows greater signal efficiency and background rejection than searching for a single signal.

In the laboratory frame, assuming that the proton is initially at rest and that the mass of the neutrino is negligible, kinematic constraints give the IBD interaction a threshold of:

$$E_\nu^{min} = \frac{(m_n + m_e)^2 - m_p^2}{2m_p} = 1.806 \text{ MeV} \tag{6.3}$$

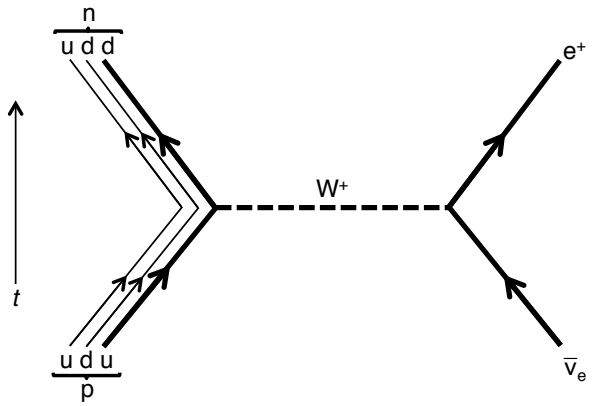

Figure 6.2: A tree-level Feynman diagram of the inverse beta decay (IBD) interaction, including the two spectator quarks in the nucleons. The upward arrow indicates time increasing along the vertical axis.

where m_n, m_e, and m_p are the mass of the neutron, electron, and proton, respectively. Above this threshold, the cross section can be expanded in powers of $1/M \approx 1/m_n \approx 1/m_p$, as in [57]. The zeroth order cross section, equivalent to the assumption of infinite nucleon mass, is a good approximation at energies up to ~ 10 MeV, but first-order corrections have also been determined. At zeroth order, the IBD cross section can be written:

$$\sigma_{IBD} = \frac{G_F^2 \cos^2\theta_C}{\pi} \left(1 + \delta_{inner}^R\right) \left(f^2 + 3g^2\right) E_e p_e \tag{6.4}$$

In this expression, G_F comes from the basic CC scattering amplitude; θ_C is the Cabibbo angle of the CKM matrix; $\delta_{inner}^R \approx 0.024$ covers inner radiative corrections; $f = 1$ and $g = 1.26$ are proton form factors; and E_e and p_e are the energy and momentum of the positron.

Since the same matrix element governs neutron decay, the IBD cross section can be expressed in terms of the neutron lifetime, τ_n:

$$\sigma_{IBD} = \frac{2\pi^2}{m_e^5} \frac{1}{f^R \tau_n} E_e p_e \tag{6.5}$$

where $f^R = 1.7152$ is a phase space factor including Coulomb, weak magnetism, recoil, and outer radiative corrections. Because of its commonality with neutron beta decay, the IBD cross section is known to high precision. Combining the prefactors in Eq. 6.5, and expressing E_e and p_e in MeV,

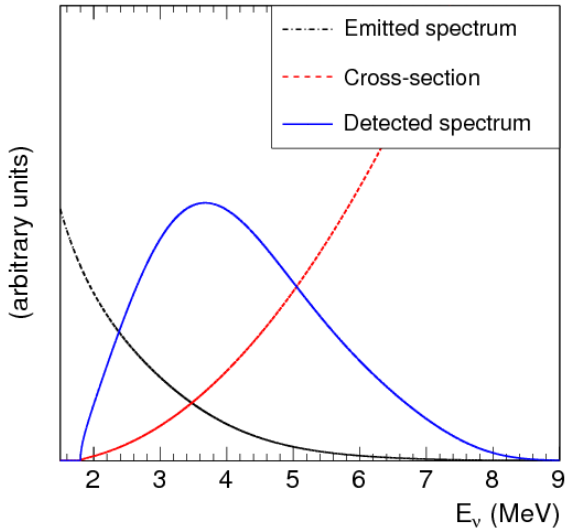

Figure 6.3: A schematic of the antineutrino flux emitted from a reactor (black), convolved with the IBD cross section (red) to yield an observable spectrum (blue) [58].

the cross section can be written as:

$$\sigma_{IBD} = \kappa E_e p_e / \mathrm{MeV}^2 \tag{6.6}$$

A recent measurement of the neutron lifetime in the MAMBO-II experiment yields $\kappa = 0.961 \times 10^{-43}$ cm^2 [47]. The convolution of the reactor antineutrino spectrum and the IBD cross section produces a characteristic spectrum of neutrino energies observable through this channel, as pictured in Fig. 6.3.

To first order in $1/M$, and neglecting small effects which occur near the threshold, the positron and antineutrino energy are related by:

$$E_e = (E_v - \Delta)\left(1 - \frac{E_v}{M}(1 - \cos\phi)\right) - \frac{\Delta^2 - m_e^2}{2M} \tag{6.7}$$

where ϕ is the angle between the antineutrino and positron trajectories in the laboratory frame and $\Delta = m_n - m_p = 1.293$ MeV. At zeroth order, Eq. 6.7 reduces to $E_e = E_\nu - \Delta$. An IBD detection technique that records the positron energy can therefore provide the incident neutrino energy, necessary for spectrum shape-based oscillation analyses.

6.3 IBD detection in liquid scintillator

Many reactor antineutrino experiments, including Double Chooz, Daya Bay, and RENO, use liquid scintillator detectors to identify IBD events. Organic scintillators offer plentiful free protons as interaction targets. Following an IBD interaction, the positron travels for up to a few centimeters, ionizing the scintillator along this path. The positron then annihilates with an electron, creating two gamma rays with total energy $2m_e = 1.022$ MeV. The total energy released through ionization and annihilation, often called the visible energy of the positron, is at zeroth order equal to $E_{vis} = E_\nu - \Delta - m_e = E_\nu - 0.78$ MeV. This energy is deposited within picoseconds of the IBD interaction.

The neutron thermalizes through elastic scattering off protons in the scintillator, traveling a net distance on the order of centimeters or tens of centimeters. Eventually, the neutron is captured by a nucleus. Capture on a hydrogen nucleus occurs in a mean time of a 200 µs and releases a single 2.22 MeV gamma ray. Double Chooz, Daya Bay, and RENO have doped their scintillators with Gd to enhance the neutron capture signal. Natural Gd is a mixture of seven isotopes, with five at abundances of 15–25%. All of these isotopes have unusually high cross sections for thermal neutron absorption; one of them, ^{157}Gd, has the highest of any stable nucleus [59]. In a typical liquid scintillator doped with Gd at $\sim 0.1\%$, over 80% of neutron captures occur on Gd, with a mean capture time of about 30 µs. These captures generate cascades of gamma rays with an average multiplicity of four and total energy of ~ 8 MeV. This signal is significantly higher in energy than most ambient radioactivity. The short time between the positron signal and neutron capture limits the probability of uncorrelated signals coincidentally mimicking an IBD event.

The photons produced by positron annihilation and neutron capture transfer energy to the scintillator through the usual processes of pair production, Compton scattering, and the photoelectric effect. Ultimately, the IBD signal begins with an energy deposition with E_{vis} ranging from $(E_\nu^{min} - 0.78) \approx 1.02$ MeV to $(E_\nu^{max} - 0.78)$ MeV, where $E_\nu^{max} \sim 10$ MeV is the upper energy bound of the reactor antineutrino spectrum. This prompt signal is followed, within a period of tens or hundreds of microseconds, by a delayed signal from the neutron capture on Gd, H, or, rarely, another nucleus such as carbon. While current reactor experiments were designed around the Gd capture channel, all have now expanded their IBD searches to include H captures.

6.4 Measuring θ_{13}-driven oscillations

Figure 6.4 illustrates how liquid scintillator detectors can be used to search for θ_{13}-driven antineutrino oscillations near a reactor. A far detector observes the rate and/or energy spectrum of IBD events at a distance of $L_{far} \sim 1$ km from the reactor. The spectral density of the IBD signal, ρ_{far} (in events per fission per unit energy), is related to the spectrum of electron antineutrinos produced in the reactor, $S(E\nu)$ (in antineutrinos per fission per unit energy), by:

$$\rho_{far}(E_\nu) = \frac{1}{4\pi L_{far}^2} P_{\bar{\nu}_e \to \bar{\nu}_e}(L_{far}, E_\nu, \sin^2 2\theta_{13}, \Delta m^2_{31}) \cdot S(E_\nu) \cdot N_p \, \sigma_{IBD}(E_\nu) \, \epsilon_{det}(E_\nu) \qquad (6.8)$$

where $P_{\bar{\nu}_e \to \bar{\nu}_e}$ is the electron antineutrino oscillation probability given by Eq. 3.10, N_p is the number of free protons in the far detector, and ϵ_{det} includes all efficiency factors from the detection process. Thus, if the antineutrino production rate (R_ν), baseline (L_{far}), mass splitting (Δm^2_{31}), proton number (N_p), and detector efficiencies (ϵ_{det}) are well known, the rate of IBD events in the far detector provides a measurement of $\sin^2 2\theta_{13}$. The spectrum of prompt visible energy from IBD events provides additional information, allowing the L/E_ν–dependent oscillation pattern to be mapped.

In practice, several elements make these strategies less straightforward. Cosmic ray muons create a variety of backgrounds to the IBD search, an effect which is mitigated but not eliminated by placing the detector underground. Accidental coincidences of radioactivity and other signals contribute additional backgrounds. Special techniques must be developed to measure the proton number and detector efficiencies to high precision. The strongest limitation for a single-detector $\sin^2 2\theta_{13}$ measurement comes from imperfect knowledge of the antineutrino production rate. Reactor operators can provide information about thermal power and fuel composition, but its precision is not absolute. Moreover, the rate and spectrum of antineutrinos produced by the main fissioning isotopes are known only from inexact, semi-empirical calculations based on a small number of datasets. In total, reactor-related uncertainties contribute roughly 2.5% uncertainty to the antineutrino flux normalization. Given $\sin^2 2\theta_{13} \sim 0.1$, reactor uncertainty limits $\sin^2 2\theta_{13}$ precision to about 25% in a single-detector, rate-based experiment.

A near detector allows experiments to reach beyond that limit. This detector is positioned close to the reactor, at an L where the electron antineutrino survival probability is still close to unity. By observing the IBD rate at this baseline, the near detector provides a measurement of the unoscillated

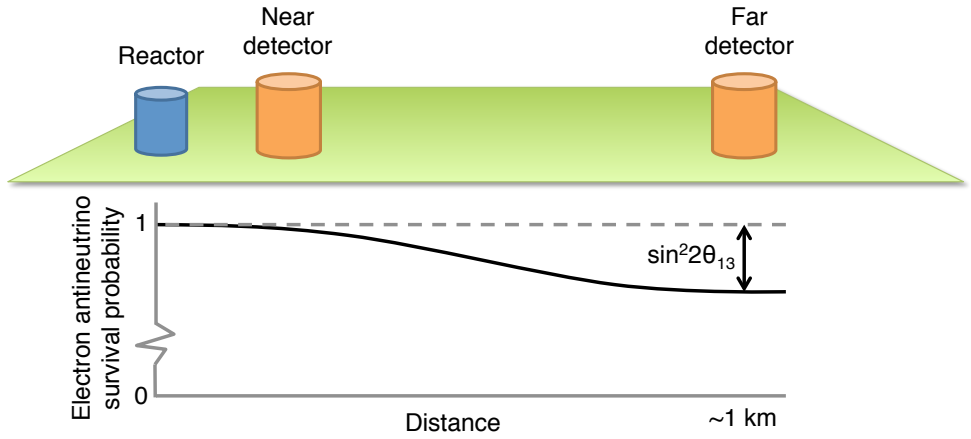

Figure 6.4: Schematic of a reactor antineutrino experiment designed to measure $\sin^2 2\theta_{13}$. The top of the diagram indicates the approximate positions of near and far detectors, with respect to a reactor. The bottom of the diagram shows the survival probability for an electron antineutrino of a typical reactor energy at these locations.

antineutrino flux. In a near-plus-far analysis, this high-statistics, empirical measurement supersedes the reactor flux simulation. If the detectors share nearly identical designs, a combined analysis also eliminates most detector-related uncertainties. Through this two-detector approach, $\sin^2 2\theta_{13}$ precision can reach the few-percent level.

Double Chooz, Daya Bay, and RENO have all pursued a near-plus-far detector strategy for measuring $\sin^2 2\theta_{13}$. Different numbers and locations of reactors at each experiment site have motivated different numbers and locations of detectors. The Daya Bay site, in the Guangdong province of southeastern China, includes six reactors, divided into two clusters separated by ~ 1 km. Each group is monitored by a pair of near detectors. A far detector hall, housing four detectors, is located with baselines of $L \sim 1.5$–2.0 km to the reactors. At the RENO site, in the Jeollanam-do province of South Korea, six reactors are arrayed in a line ~ 1.5 km long. Two detectors are located along the perpendicular bisector of that line: a near detector at a distance of ~ 150 m from the reactor line, and a far detector at a distance of ~ 1.5 km. Figure 6.5 presents diagrams of the Daya Bay and RENO sites. With six reactors each, these experiments benefit from a large

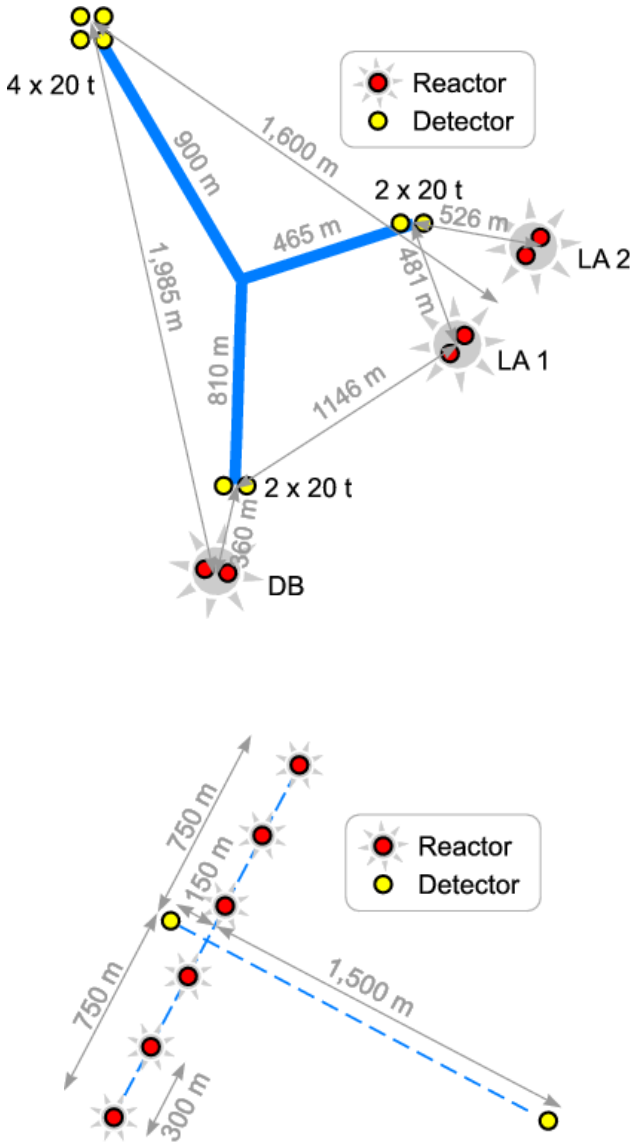

Figure 6.5: Top: Diagram of the Daya Bay experiment site, indicating the relative locations of reactors. Bottom: Diagram of the RENO site [60].

flux of antineutrinos. However, the multiplicity of reactor-detector baselines complicates oscillation analysis. Chapter 7 describes the simpler Double Chooz site.

Part II

Design of the Double Chooz experiment

Chapter 7

Experiment site

The Double Chooz experiment is located at the Chooz Nuclear Power Plant, next to the town of Chooz, in the Ardennes department (Champagne-Ardenne region) of France. The town and plant lie inside a meander of the Meuse River, a few kilometers from the Belgian border. This chapter describes the background of the site and its layout.

7.1 Site selection

The Chooz Nuclear Power Plant has several features which distinguish it from the 18 other nuclear power stations in France. The original reactor at this site, known as Chooz A, was the first commercial PWR to be constructed in Europe. The unusual design of Chooz A placed the reactor and many auxiliary components in caverns dug into the bedrock of a 200 m–high hill. Chooz A supplied power to the electrical grid, at around 300 MW, from 1967 to 1991 [61]. Decommissioning of the plant began in 2007 and continues through the present.

Construction of two more advanced PWRs, called Chooz B1 and B2, begin nearby in 1984. The design of these reactors would make them two of the most powerful cores in the world. By the early 1990s, the site became an attractive site for the CHOOZ experiment. An out-of-service access tunnel for the Chooz A reactor provided an underground location for the detector, and the potential for taking data both before and after the reactors began operation offered a direct opportunity to measure backgrounds. Units B1 and B2 began nominal-power operation in April and August of 1997, respectively [26].

Three features made the Chooz Nuclear Power Plant an especially appealing site to host a two-detector reactor neutrino experiment: the ability to reuse the CHOOZ detector hall; the high power of the reactors; and the history of a successful relationship between the plant operators and a physics collaboration.

7.2 Site layout

Figure 7.1 identifies the main features of the Double Chooz site, and Fig. 7.2 indicates the distances between each reactor-detector pair. The location of the far detector, at approximately 1050 m from the barycenter of the reactors, was fixed by the preexisting CHOOZ detector hall. It is closer to the reactors than the ideal position for measuring $\sin^2 2\theta_{13}$, given the currently known value of $|\Delta m^2_{32}| \sim 2.4 \times 10^{-3}$ eV^2 and the average energy of reactor antineutrinos observed in IBD interactions, $\langle E_\nu \rangle \sim 4.2$ MeV. The optimal far detector baseline would be closer to the first oscillation maximum at $L_{max} = \frac{\pi}{2}\langle E_\nu \rangle / (1.267|\Delta m^2_{32}|) \approx 2100$ m, perhaps with some consideration for the reduction in signal statistics caused by the decreased solid angle. This displacement strongly limits the Double Chooz ability to measure $|\Delta m^2_{32}|$ in a spectrum shape-based analysis, since at $L = 1050$ m, the oscillation maximum occurs just barely above the IBD threshold. However, it makes Double Chooz a valuable addition to a multi-experiment, purely rate-based measurement of $|\Delta m^2_{32}|$, as demonstrated by some Double Chooz collaborators [62].

Ideally, the near detector would be as close as practical to the reactors to maximize signal statistics and minimize the impact of oscillation. A large overburden would also be an advantage. A further asset would be an isoflux position with respect to the far detector. Detectors at isoflux positions observe the same ratio of flux from the two reactors:

$$\text{Isoflux condition:} \qquad \frac{\Phi^{ND}_{B1}}{\Phi^{ND}_{B2}} = \frac{\Phi^{FD}_{B1}}{\Phi^{FD}_{B2}} \tag{7.1}$$

where Φ^d_r refers to the flux from the r^{th} reactor observed in the d^{th} detector. The isotropy of the reactor flux implies $\Phi^d_r \propto 1/(L^d_r)^2$, where L^d_r is the baseline between the r^{th} reactor and the d_{th} detector. Therefore, in the isoflux condition of Eq. 7.1, the four reactor-detector baselines would satisfy:

$$\frac{L^{ND}_{B1}}{L^{ND}_{B2}} = \frac{L^{FD}_{B1}}{L^{FD}_{B2}} \tag{7.2}$$

Figure 7.1: The site of the Double Chooz experiment, highlighting the two operating reactors, B1 and B2; buildings used for decommissioning of the Chooz A reactor; and the locations of the two Double Chooz detectors (both underground, as indicated in the labels). The town of Chooz is also identified.

Because B1 and B2 are operated with somewhat different power profiles and fuel histories, they do not emit identical rates and spectra of antineutrinos. Consequently, only a near detector at an isoflux location can act as a perfect flux monitor for the far detector. This condition allows maximal suppression of reactor-related uncertainties, leading to maximal precision in two-detector oscillation analyses. With increasing deviation from the isoflux arrangement, a two-detector analysis suffers from an increased amount of irreducible reactor uncertainty. The analysis advantages of a near-isoflux geometry are discussed further in Chs. 18–19.

The choice of the near detector site was a compromise of isoflux considerations, overburden, and technical feasibility. As displayed in Fig. 7.2, the location does not exactly coincide with the isoflux curve of the far detector. Still, its proximity to that curve will benefit two-detector measurements.

7.3 Baseline measurements

Precise assessments of the reactor-detector differences are important for oscillation analyses. Multiple geodetic surveys have been conducted at the Double Chooz site in order to locate the reactors and the far detector in a common reference frame. Surveys of the near detector location are ongoing. Table 7.1 gives the measured baselines for the far detector and preliminary estimates for the near detector.

Table 7.1: Distances between each reactor-detector pair, as determined from geodetic surveys [63]. Near detector baselines are preliminary estimates.

Detector	Distance to Reactor B1 (m)	Distance to Reactor B2 (m)
Far	1114.656 ± 0.015	997.839 ± 0.015
Near	465	351

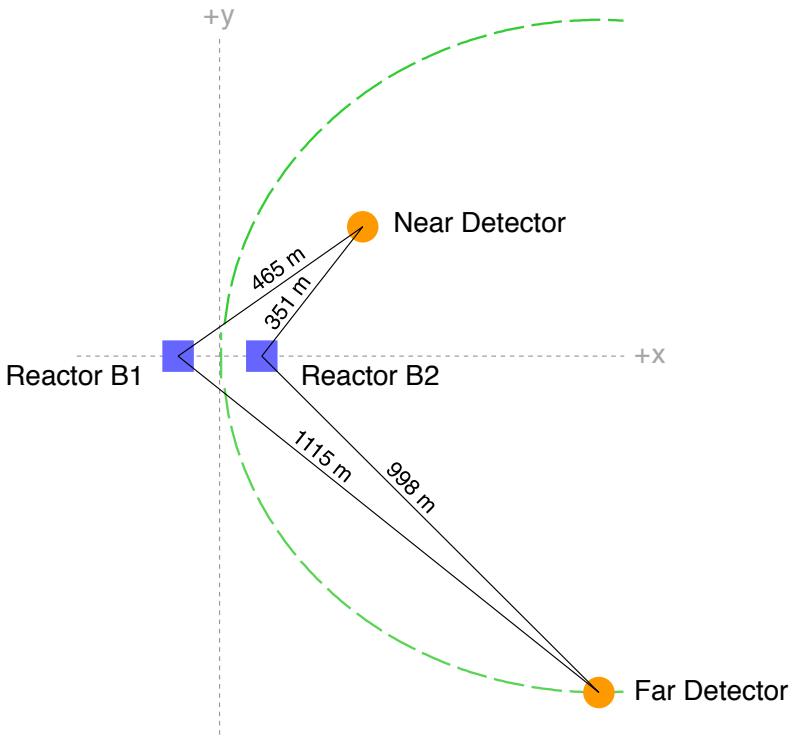

Figure 7.2: A schematic of the Double Chooz site, showing the positions of the reactors (blue squares) and detectors (orange circles) in a two-dimensional coordinate system centered between the reactors. In this system, reactor B1 is located at $(x = -d, y = 0)$, and reactor B2 is located at $(x = d, y = 0)$, where $d = 82.5$ m. The green dashed curve indicates all points at an isoflux location with respect to the far detector, assuming full-power operation of both reactors and assuming that the near detector is located at a depth of $z^{ND} = 40$ m. The curve is described by the equation $y^2 = \left((x+d)^2 - k^2(x-d)^2 + (1-k^2)(z^{ND})^2\right)/(k^2-1)$, where $k = L_{B1}^{FD}/L_{B2}^{FD} = 1.12$.

Chapter 8

Reactors

This chapter gives a brief overview of the design and operation of the two reactors currently running at the Chooz Nuclear Power Plant, focusing on features which are relevant for antineutrino production.

8.1 Design and operation

The Chooz B1 and B2 reactors were designed and built by Framatome (now Areva). They were the first reactors constructed in the N4 design, named for the four steam generators in each unit. This design affords a higher nominal power and more load-following capability than previous PWRs constructed in France, along with more advanced safety features. The reactors are operated by Électricité de France (EDF). Originally, the nominal power of the reactors was 1455 MW_e, but a 2003 upgrade increased the output to 1500 MW_e [64]. This electrical output corresponds to a thermal power of 4.27 GW_{th} [65]. Only two other reactors in France and a small number elsewhere in the world have such a high thermal power, which is directly proportional to the antineutrino flux from the core (as demonstrated in Eq. 6.1).

The Chooz reactors are fueled with low enriched uranium dioxide, encapsulated in Zircaloy-clad pellets. These pellets are stacked inside rods, which are then grouped into fuel assemblies. Four types of assemblies are used in the Chooz reactors, with initial ^{235}U enrichment of either 1.8%, 3.4%, or 4%. More than 99.7% of fissions occur in the two dominant uranium isotopes, ^{235}U and ^{238}U, and two plutonium isotopes bred during operation, ^{239}Pu and ^{241}Pu.

Each reactor core consists of 205 assemblies, each containing 264 fuel rods, arranged in a cylindrical bundle with radius 3.47 m and height 4.27 m [65, 66]. Figure 8.1 shows the structure of the core, fuel assemblies, and fuel rods. A special feature of the N4 core is a set of "gray" control rods, so called because they absorb fewer neutrons than standard "black" control rods. In addition to smoothing thermal power output across the core, these rods allow the total power to be varied between maximal and lower levels relatively quickly [66]. The ability to respond to electrical demand is valuable in the heavily nuclear-dominated French system.

In each Chooz reactor, as in a typical PWR, the core is enclosed within a pressure vessel filled with light water, which acts as a neutron moderator and primary coolant. Fission reactions heat the water to approximately 300°, but it remains liquid due to the high pressures, on the order of 15 MPa, maintained in the vessel [26]. The primary coolant water contains a small amount of boron, a potent thermal neutron absorber, which is monitored for information about core conditions. Heat from the primary circuit is exchanged with a secondary circuit inside the steam generators. These components, along with the core, are enclosed within a concrete containment building. The secondary circuit flows outside the containment building to steam turbines which drive electrical generators. Steam is then condensed by thermal contact with a tertiary circuit, which draws cooling water from the Meuse River. Much of the water in the tertiary circuit is recycled through a cooling tower. Each reactor at Chooz has its own containment building, water circuits, and cooling tower, and the two reactor systems are controlled independently.

8.2 Refueling

Each reactor is refueled once every 12–14 months. The process takes about one month, during which the reactor is shut down. Refueling is typically time-staggered so that one unit can continue to operate while the other is being serviced. In a refueling event, about a third of the assemblies are removed and replaced with a fresh set. The old and new assemblies are rearranged in specific patterns to maintain a homogeneous neutron flux across the core. Spent assemblies are stored in pools on the Chooz site. Decays of long-lived isotopes in the spent fuel produce an antineutrino flux concentrated at low energies, but the rate is small enough to neglect in current analyses.

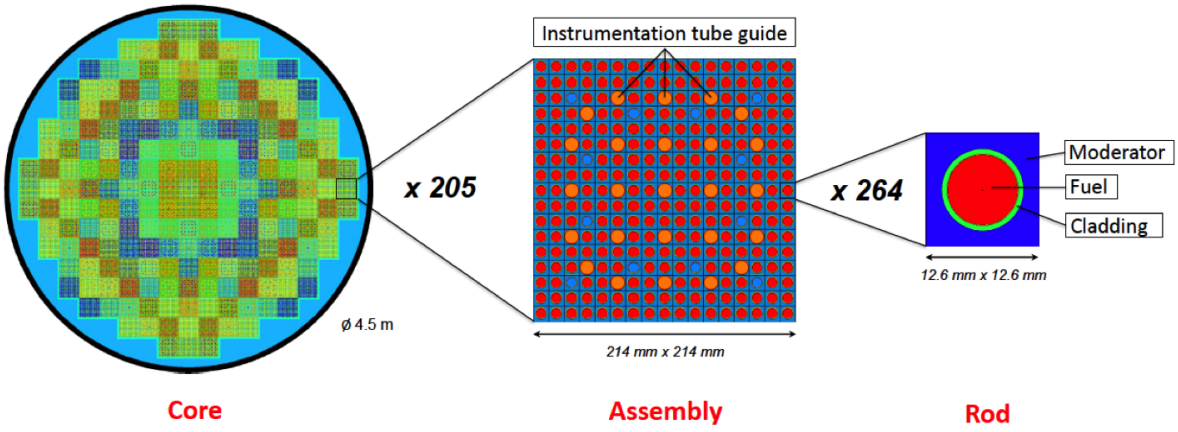

Figure 8.1: Diagram of the Chooz reactor core, with colored squares indicating different types of fresh and used fuel assemblies (left); a single fuel assembly (center); and a single fuel rod (right) [67].

Chapter 9

Detector design

The Double Chooz detectors were designed to efficiently observe IBD interactions in the manner described in Section 6.3, with good energy reconstruction and minimal contamination from non-IBD events. The designs of the near and far detectors are identical in most respects to maximize cancellation of detector-related uncertainties in a two-detector analysis. After presenting an overview of the detector geometry, this chapter describes the design and function of each subsystem.

9.1 Overview

The main portion of each detector consists of a set of concentric cylinders, as shown in the cross-sectional view of Fig. 9.2. The cylindrical form is the most symmetrical geometry that was also technically achievable, as fabrication of three nested spherical vessels was considered too difficult. The innermost volume is the neutrino target (NT), the fiducial volume for the Gd-based analysis. Surrounding it, and optically connected, are the gamma catcher (GC) and buffer. The GC was designed to enhance signal collection from the NT, and it also serves as an extended fiducial volume for the H-based analysis. The buffer provides some insulation from backgrounds. Together, the NT, GC, and buffer constitute the inner detector (ID). Installed on the buffer walls are photomultiplier tubes (PMTs) which collect light from the ID.

Surrounding the buffer, but optically separated, is the first of two cosmic ray veto systems. This system is called the inner veto (IV), since it is located within the detector shielding. The ID and IV volumes are encased in non-instrumented material to limit the infiltration of external

gamma rays and neutrons. For the far detector, a 15-cm layer of demagnetized steel was chosen for shielding. The larger near detector hall enabled the more economical choice of a 1 m-thick water shield. Extending up from the neutrino target is the chimney, used for filling the liquid volumes, and a glovebox. The second cosmic ray veto system, the outer veto (OV), consists of flat panels installed above the main detector components. The OV array is centered around the chimney and covers an area well beyond the detector pit, particularly in the far detector hall.

During normal operations, data acquisition systems continuously read out signals from the ID, IV, and OV. Dedicated calibration systems are also built into each detector.

9.2 Neutrino target

The NT was designed to isolate the primary class of signal events, IBD interactions followed by neutron capture on Gd. The vessel consists of an 8 mm-thick, cylindrical acrylic vessel with a shallowly conical top and bottom, 2458 mm in height and 1150 mm in radius. It is filled with 10.3 m^3 of Gd-doped, organic liquid scintillator. The scintillator formula was developed specifically for the Double Chooz experiment, with a composition tuned to optimize optical properties, radiopurity, and chemical stability. Particular emphasis was placed on the last feature, since gradual oxidation of the CHOOZ scintillator caused reduced sensitivity and ultimately the premature end of that experiment.

The bulk of the NT liquid is a mixture of 1,2-dimethyl-4-(1-phenylethyl)-benzene, more commonly called ortho-phenylxylylethane (o-PXE), and n-dodecane, in a 1:4 volume ratio. With its aromatic structure, o-PXE is the major scintillating agent, while n-dodecane is added to enhance material compatibility with the acrylic vessel. Advantages of this combination include its well-defined chemical composition, versus, for example, most mineral oils; relatively high flash point, a safety improvement over pseudocumene-based scintillators; and the ability to tune density and light yield by adjusting the component ratio. Both components have previously been used or explored for use in multi-ton liquid scintillator detectors [68, 69].

The o-PXE/n-dodecane mixture emits ultraviolet (UV) scintillation light but is most transparent to light in the visible range. That range is also the regime in which the ID PMTs have the highest quantum efficiency. To transfer the scintillation light to longer wavelengths, two wavelength-

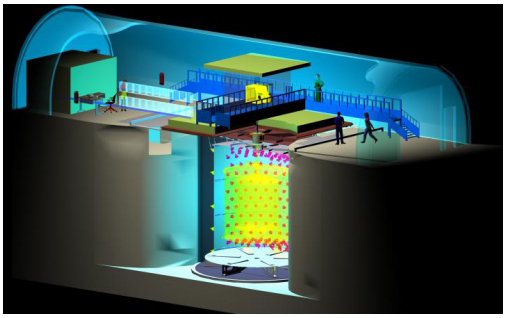

Figure 9.1: A cutaway view of the Double Chooz far detector laboratory, showing the neutrino detector inside the detector pit, with the gamma catcher and neutrino target highlighted in yellow-green.

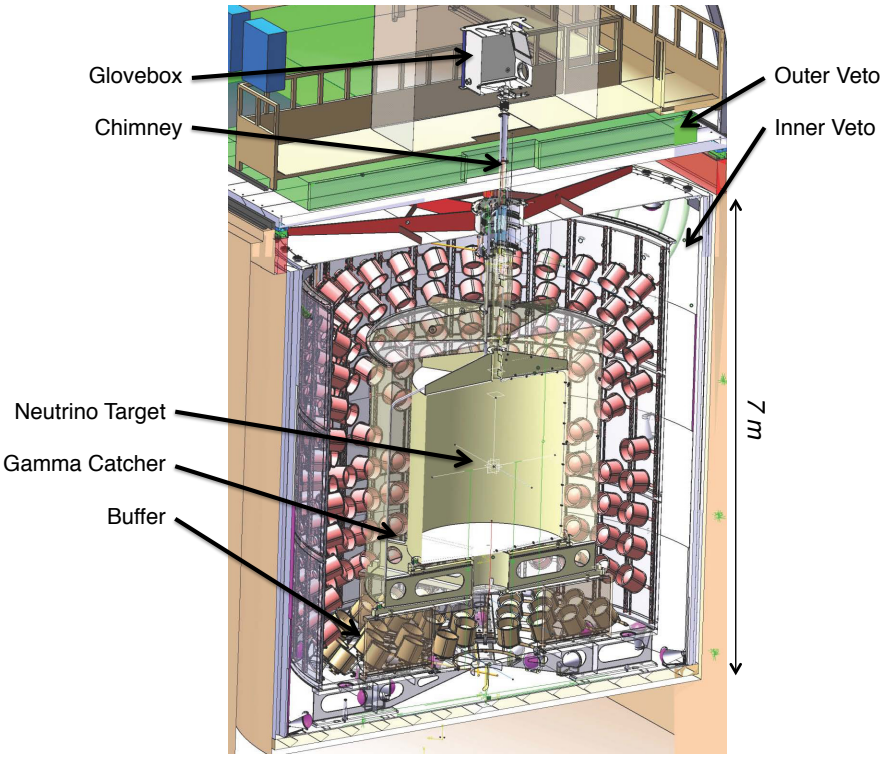

Figure 9.2: A cross-sectional view of a Double Chooz detector. This diagram does not show the true extent of the outer veto: in the far detector, it is significantly larger than the size shown here; in the near detector, it is smaller.

shifting solutes were added to the NT liquid. The primary fluor, 2,5-diphenyloxazole (PPO), participates mainly in non-radiative energy transfer from o-PXE, which occurs through Coulombic interactions. PPO then emits photons which are shifted by the secondary wavelength shifter, 4-bis-(2-methylstyryl)benzene (bis-MSB). The final scintillation light, after non-radiative transfer and wavelength shifting, occurs in the 350–500 nm range [70].

The rare earth metal Gd does not dissolve readily enough in organic solvents to produce the concentration required for IBD tagging. To increase solubility, Gd was incorporated into a metalorganic complex. The chosen complex is a metal-β-diketone, Gd(III)-tris-(2,2,6,6-tetramethyl-heptane-3,5-dionate) ($Gd(thd)_3$). The stability and high vapor pressure of this complex allowed it to be purified by sublimation, removing most radioimpurities. The $Gd(thd)_3$ complex is dissolved in oxolane, commonly known as tetrahydrofuran (THF), before addition to the NT solution. The final Gd concentration in the NT is 0.123% by weight, or approximately 1 g/l.

The composition of the NT liquid is summarized in Tab. 9.1. Careful selection of suppliers, preparation techniques, and purification methods minimized chemical contaminants, particularly water, and radioimpurities. The scintillator for the near and far detectors was prepared at the same time to ensure similarity. After preparation, a variety of tests were performed to assess the chemical and optical properties of the liquid. The attenuation length at 430 nm was measured at 7.8 ± 0.5 m, safely longer than the linear dimensions of the ID [70]. Radiopurity standards were met. A determination of the number of free protons in the NT volume, a crucial quantity for oscillation analyses, is discussed in Sec. 13.2.2.

9.3 Gamma catcher

The NT is surrounded by a 55 cm-thick region called the gamma catcher. The GC consists of a 12 mm-thick, cylindrical acrylic vessel containing 22.5 m^3 of liquid scintillator with no Gd dopant. The acrylics separating the GC from adjacent volumes are transparent in the visible to UV range. Originally, the purpose of the GC was to allow gammas which escape the NT to deposit their energy in scintillator. With the development of the H capture-based IBD selection, the GC also serves an extended neutrino target.

The GC scintillator was designed to match the density and light yield of the NT. Density

Table 9.1: Composition of the neutrino target and gamma catcher scintillators [70].

Neutrino target

Component	Amount
n-dodecane	80% by volume
1,2-dimethyl-4-(1-phenylethyl)-benzene (ortho-phenylxylylethane, o-PXE)	20% by volume
Gd(III)-tris-(2,2,6,6-tetramethyl-heptane-3,5-dionate) ($Gd(thd)_3$)	4.5 g/l
oxolane (tetrahydrofuran, THF)	0.5% by weight
2,5-diphenyloxazole (PPO)	7 g/l
4-bis-(2-methylstyryl)benzene (bis-MSB)	20 mg/l

Gamma catcher

Component	Amount
Shell Ondina 909 mineral oil	66% by volume
n-dodecane	30% by volume
1,2-dimethyl-4-(1-phenylethyl)-benzene (ortho-phenylxylylethane, o-PXE)	4% by volume
2,5-diphenyloxazole (PPO)	2 g/l
4-bis-(2-methylstyryl)benzene (bis-MSB)	20 mg/l

matching to at least 1% was required to prevent mechanical stress on the thin NT vessel, while light yield equivalence creates a nearly-uniform detector response across the NT-GC volume. Since the $Gd(thd)_3$ complex lowers the light yield of the NT, the GC formula could not use the same o-PXE concentration. Mineral oil was introduced as an additional, non-scintillating solvent. Ondina 909, a mixture of branched and unbranched alkanes manufactured by Shell Chemicals, was chosen for its transparency and low aromaticity. Relative to the NT, the concentration of PPO was reduced to produce a difference in the scintillation emission profiles of the two volumes. Scintillation light is emitted more slowly with lower PPO concentrations, so in principle, events in the GC and NT can be distinguished by their PMT pulse profiles. This technique has not yet been employed in Double Chooz analyses, but in the future, it may be used to study the spill-in and spill-out currents discussed in Section 13.2.4.

Table 9.1 summarizes the composition of the GC scintillator. As with the NT, the GC solution was prepared with attention to chemical and radiological purity.

9.4 Buffer and inner detector PMTs

Outside the GC is the 105 cm-thick buffer region. The buffer consists of a stainless steel cylindrical vessel filled with 110 m^3 of non-scintillating oil. On the top, bottom, and side walls of the vessel, a total of 390 PMTs are installed to view the ID volumes.

The buffer volume exists to shield the ID scintillator from gamma radiation originating in the material outside the detector and in the PMT glass. It is a significant design improvement over the CHOOZ experiment, in which PMTs were located in the GC volume and radioactivity was a major background [26]. The buffer liquid is blend of Shell Ondina 917 mineral oil, at 53% by volume, and a more refined mixture of n-alkanes, at 43% by volume. This formula was chosen to maximize transparency, minimize scintillation, and match the density of the NT and GC liquids [71].

The PMTs installed on the buffer walls provide 13% photocathode coverage, a level chosen to approximately match the 15% coverage successfully used in the CHOOZ detector. The photocathodes are oriented to face the center of the NT, allowing fairly uniform light collection across the detector. The PMT model is the 10-inch-diameter, semi-hemispherical R7081, manufactured by Hamamatsu Photonics K. K. [72]. This model was chosen to suit the Double Chooz detector dimensions and

because its glass can be produced with very low levels of the radioimpurities typically found in borosilicates. Each PMT has a 10-stage dynode chain which provides a gain of 10^7 when operated at a voltage between approximately 1200 and 1600 V [73]. The bialkali cathode is sensitive in the 300–600 nm range, with a maximum response around 400 nm [74].

To protect the circuitry from the buffer liquid, each PMT base is encased in transparent, oil-resistant epoxy coating. In the far detector, some PMT base circuits were observed to spontaneously emit light flashes, which can propagate through the epoxy, trigger the data acquisition system, and become background for IBD searches. For some time, 14 of the most problematic emitters were switched off, but the development of new analysis-level cuts allowed them to be restored. To reduce PMT light emission in the near detector, the base of each PMT was covered with a black polyester material. In both detectors, the entire PMT is surrounded by a cylinder of high-permeability mu-metal, which shields the dynodes from the magnetic field of the earth and, in the far detector, the residual magnetic field of the steel shielding [75].

Before the 780 near and far detector PMTs were installed, a number of their characteristics were assessed. The average photon detection efficiency, defined as the product of quantum efficiency and photon collection efficiency at the first dynode, was 23% in the multi-photo electron regime [73]. The PMT response was observed to be linear for signals up to 300 photo electrons (PEs). Measurements were also made of the ratio of the single PE peak to the pedestal, spread in PE transit time, after-pulse probability, and dark count rate [73, 74].

9.5 Inner veto

Surrounding the buffer, and optically separated by the steel buffer wall, is a 50 cm-thick region of scintillator designed to detect cosmic ray muons. ID data taken directly after a muon detection in the IV can be excluded from analyses, reducing cosmogenic backgrounds in oscillation searches. The IV consists of a cylindrical stainless steel vessel filled with 90 m^3 of a third formulation of scintillator. A total of 78 PMTs are installed on its walls. In addition to its role as an active cosmic veto, the IV acts as an extra layer of passive shielding from gamma rays and neutrons. In recent analyses, it has also provided tagging information for certain backgrounds.

The IV liquid is a 1:1 mixture, by volume, of n-alkanes and linear alkylbenzene (LAB), a

scintillator with molecules consisting of a benzene ring and linear, saturated hydrocarbon chain of varying length. PPO is included as a primary fluor, at 2 g/l, and bis-MSB is included as a secondary wavelength shifter, at 20 mg/l [71].

The IV PMTs, of the 8-inch Hamamatsu R1408 model, were recycled from the SuperKamiokande experiment and, originally, the Irvine-Michigan-Brookhaven (IMB) experiment. They are located strategically within the IV to maximize muon detection efficiency, with 24 PMTs on the top wall, 42 on the bottom, and 12 circling the side wall at its midpoint. To increase light collection, all inner surfaces of the IV are covered with reflective sheets or white paint.

9.6 Outer veto

The OV is a large array of solid scintillator modules installed above the main detector components. It was designed to identify cosmic ray muons which pass near or clip a small portion of the liquid scintillator volumes. Such muons can produce backgrounds for IBD searches, either by stopping and decaying in the NT or GC or by spalling neutrons from the surrounding materials.

Each OV module consists of 64 strips of polystyrene containing two wavelength shifters: PPO, at 1%, and 1,4-bis(5-phenyloxazole-2-yl)benzene (POPOP), at 0.03%. The strips are extruded with a cross-sectional area of 5 cm $\times$ 1 cm, with a hole running through the center. A 1.5 mm–diameter wavelength shifting fiber, with a polystyrene core and multi-component, wavelength-shifting cladding, is strung through each hole. The strips are coated with a 0.25 mm–thick, reflective layer of titanium dioxide. Within each module, the strips are arranged in two layers, with a 2.5-cm offset to eliminate dead area at the strip interfaces. This arrangement results in a module width of 1.6 m; two module lengths, 3.2 m and 3.6 m, were chosen to fit the detector cavern dimensions. One end of each fiber is fitted with a mirror. At the other end, the fibers are coupled to a 64-channel, Hamamatsu H8804 multi-anode PMT (M64). The modules are wrapped in thin aluminum sheets, with all junctures sealed against external light.

The layout of the OV differs significantly between the near and far detectors. In the far detector, the OV includes 44 modules, arranged in two bilayers. With each bilayer, the modules are oriented in orthogonal directions to allow two-dimensional localization of through-going muons. Within each layer, adjacent modules overlap to maintain continuous two-strip coverage. The lower bilayer, called

the lower OV, rests on the steel shielding above the IV and covers an active area of 13.1 m × 7.2 m. These dimensions were chosen to maximize muon detection area within the limited width of the far detector cavern. The upper OV, approximately 3.2 m × 6.4 m, is suspended about 2 m above the lower OV. It covers the gap in the lower OV where the chimney emerges and allows hodoscopic tracking of muons which also penetrate the lower OV.

A similar but larger double-bilayer design was originally planned for the near detector. This design included up to 70 modules, divided between a 13.2 m × 11 m lower OV and a 6.4 m × 6.4 m upper OV. However, in the three years between the installation of the far detector and near detector, multiple techniques (described in Sec. 13.3) were developed to powerfully reduce stopping muon and fast neutron backgrounds without the use of the OV. A study was performed to determine whether a full OV was necessary for two-detector θ_{13} analyses. Results of this study suggested that a full near detector OV would have negligible impact on such analyses.

Ultimately, an alternative and much smaller OV design was chosen for the near detector. This design was developed to preserve some muon-tagging capability, which is useful for backgrounds studies, while avoiding the expense and complicated installation required for a near-hermetic muon veto. The current near detector OV includes a single layer of eight modules, arranged without overlap along the center line of the detector, plus two modules in a second layer, flanking the chimney. Eventually, some modules may be moved or added to enhance muon veto efficiency over the chimney.

9.7 Calibration systems

Oscillation analyses require precise knowledge of the energy response and neutron detection efficiency of the detector. Much of this information is collected from dedicated calibration systems built into or attachable to the detector. These include an LED-based light injection system and three systems for radioactive source deployments. Additional calibration information comes from cosmogenic spallation neutrons.

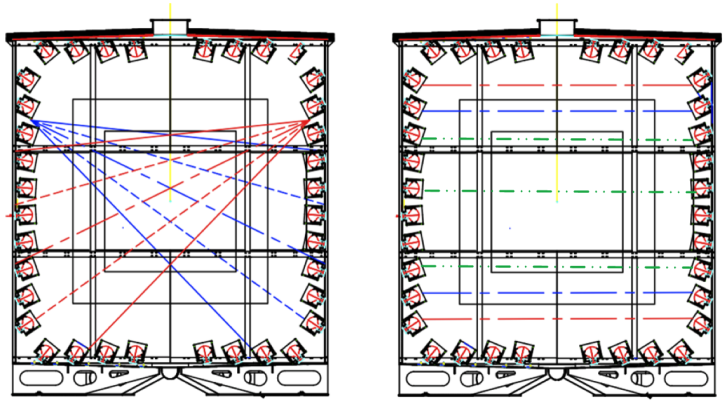

Figure 9.3: Illustrations of the inner detector light injection system [76]. The left image includes two examples of diffuse beams, each of which illuminates many PMTs. The right image includes seven pencil beams: two traversing only the buffer (red), two traversing the buffer and gamma catcher (blue), and three traversing all three ID volumes.

9.7.1 Light injection system

Both the ID and IV are equipped with multi-wavelength light injection systems. The systems are activated at regular intervals to monitor readout channel gains, time offsets, and scintillator properties.

In the inner detector, the light injection system begins with three light emitting diodes (LEDs), each producing a different wavelength. The light is channeled to 46 injection points interspersed between the PMTs. For 32 of these points, the light is routed through polymer optical fibers and, at their ends, through diffuser plates which spread the beam to an opening angle of about 22°. For the remaining 14 points, the light is routed through quartz optical fibers which produce narrower pencil beams of about 7°. The LEDs can be operated with variable flash rates, light intensities, and pulse time profiles. Light emitted by the 385-nm LED is mostly absorbed in the GC, while light emitted by the 425-nm LED is partially absorbed in the GC and NT, and light emitted by the 470-nm LED is not appreciably attenuated. Different combinations of wavelengths, LED settings, and beam profiles provide complementary probes of the two scintillator volumes and multiple PMT gain regimes. Figure 9.3 illustrates diffuse and pencil beam injections.

A similar light injection system is installed in the IV. The IV system uses only two LEDs, emitting light at wavelengths of 365 nm and 475 nm. Uses of the light injection system in data

reconstruction and MC tuning are discussed in Chs. 11–12.

9.7.2 *z*-axis system

Radioactive sources, or laser or LED light sources, can be deployed along the z-axis (central vertical axis) of the chimney and NT using a fishing line–style system. This system consists of a thin, Teflon-coated, stainless steel cable mounted on a pulley-and-weight device controlled by a stepper motor. Sources are introduced through a glovebox mounted at the top of the chimney. Each source can be positioned at any height between the chimney and 1 cm above the bottom of the NT, with its location known to within 1 mm.

Three gamma ray sources, each with an activity of approximately 50 Bq, have been deployed along the z-axis in millimeter-scale capsules. Decays of the ^{137}Cs source usually produce a single 0.662 MeV gamma. The ^{68}Ge source decays by electron capture to ^{68}Ga, which in turn beta-plus decays to ^{68}Zn, leading to the emission of two 0.511 MeV gammas from the subsequent positron annihilation. The ^{60}Co source decays with either a 1.173 MeV or 1.333 MeV gamma. In addition to the gamma sources, tagged and untagged ^{252}Cf sources have been deployed to study neutron capture processes. The majority of ^{252}Cf decays proceed through alpha emission, but 3% occur through spontaneous fission, releasing an average of 3.73 neutrons with average energy of about 2 MeV. In the tagged source, the ^{252}Cf is enclosed within a gas ionization chamber capable of detecting fission products and thus identifying neutron-emitting events.

A laser system has also been deployed along the z-axis. In this system, 380-nm or 470-nm light is channeled through an optical fiber to a silicone diffuser ball.

9.7.3 Guide tube system

Radioactive sources can also be deployed in the GC inside a 5 mm–diameter, Teflon-lined, stainless steel tube. The guide tube traces a loop between the inner and outer GC walls, as shown in Fig. 9.4. Inside it, sources are pulled on a motor-driven cable, with positions known to a precision of 1 cm. The same gamma- and neutron-emitting isotopes have deployed along the z-axis have been used in the guide tube. A separate guide tube is installed in the buffer, but this tube has not been instrumented.

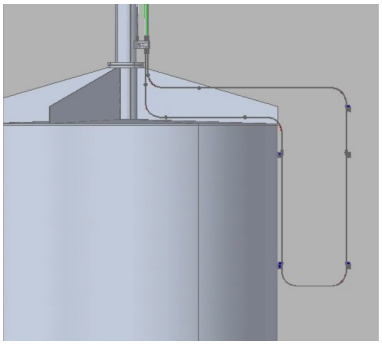

Figure 9.4: A diagram of the gamma catcher guide tube. The cylindrical vessel is the neutrino target. The outer GC wall, not shown here, is just beyond the outer vertical piece of the guide tube.

9.7.4 Articulated arm

A third source deployment system, called the articulated arm, would allow radioisotopes to be moved with full three-dimensional freedom through the NT volume. This coverage would provide valuable information about the detector response across the entire Gd-channel signal region. The articulated arm consists of a telescoping support rod, pivot, and variable-length, source-bearing segment, along with a drive assembly. When installed on top of the glovebox, the articulated arm could be extended into the NT, and the source-bearing segment could be raised or lowered by a cable. The system has not yet been successfully deployed in either detector.

9.8 Data acquisition systems

A single data acquisition system (DAQ), called the neutrino DAQ (ν-DAQ) collects signals from the ID and IV PMTs. An independent system, the OV-DAQ, reads out the OV.

9.8.1 Neutrino DAQ

Figure 9.5 illustrates the flow of signals in the ν-DAQ. To streamline cabling and reduce ground loop risk, each PMT in the ID and IV is connected to a single coaxial cable (type RG-303/U) which transmits both high voltage (HV) and output signals. The HV, supplied by CAEN-A1535P modules [77], distributes $\sim$ 1.3 kV across each PMT channel. The PMT signals carry $\sim$ 5 mV per PE. The PMT cables are routed to a custom high voltage splitter, consisting of passive filter circuits for each

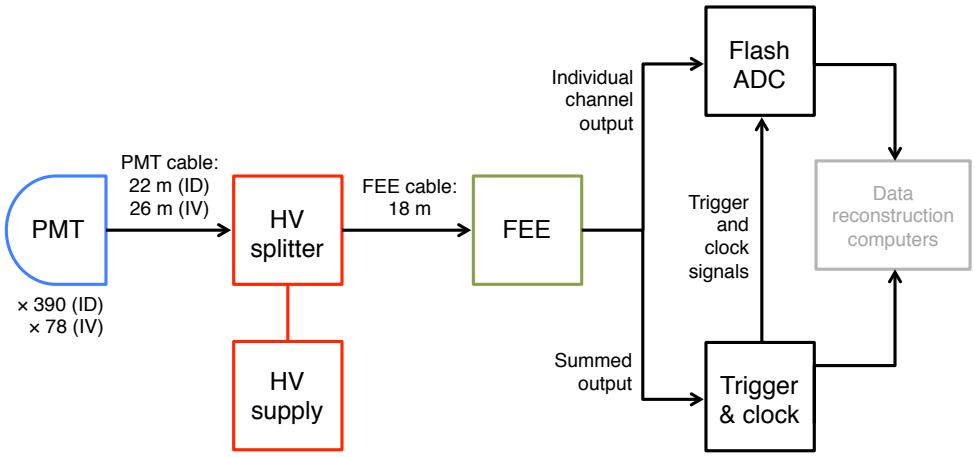

Figure 9.5: A flow chart of the neutrino DAQ, with information transfer indicated by arrows. Adapted from [78].

channel. These circuits decouple the HV and PMT signals while damping HV noise.

From the splitters, the PMT signals are transmitted through RG-58 cables to custom front-end electronics (FEE) modules [79]. The FEE circuits first filter noise from the PMT channels. The resulting signal is processed through two paths: the neutrino path, which receives an overall gain of 7.8, and the muon path, which receives a gain of 0.55. Each path corrects for the drift in PMT baselines that often follows large energy depositions. The neutrino path produces two outputs: the processed signal from each channel, which is delivered to flash analog-to-digital converters (ADCs), and the sums of groups of eight channels, which are delivered to the trigger system. The output of the muon path is conveyed to a dedicated muon electronics system.

The trigger system consists of a trigger master board, two ID trigger boards, and an IV trigger board. Each ID trigger board receives the summed analog signals from half of the eight-PMT groups. Groups are assigned in order to give the trigger boards equivalent coverage of the detector volume. The IV channels are bundled into smaller groups of three to six, all read by the IV trigger board. Any one of the three boards can generate a trigger signal which, via the master board, causes simultaneous readout of all ID and IV channels. The trigger logic depends on the total charge collected on each board and multiplicity among the groups. A variety of trigger conditions can be implemented. In the standard configuration, at least one of the ID boards must receive a total

charge equivalent to about 0.35 MeV, or the IV board must receive charge equivalent to 10 MeV, corresponding to an 8-cm minimum-ionizing muon track. The trigger is designed to be virtually deadtime free, and trigger efficiency is 100.0% above all analysis thresholds. During standard data-taking, the trigger rate is roughly 300 Hz. Data is also collected with fixed-rate, $\sim$ 1 Hz triggers to monitor various detector properties and effects of analysis cuts. The trigger master board distributes a 62.5 MHz clock signal which is used to synchronize the DAQ subsystems.

The channel-by-channel signals from the neutrino path are digitized by a set of 64 CAEN-Vx1721 (VME64x) cards, developed specifically for this application [80, 81]. Each card contains 8 channels, and each channel is equipped with an 8-bit flash ADC (FADC). The FADCs process 500 million samples per second, resulting in digitized waveforms with 2-ns time bins. The dynamic range of the FADCs extends to approximately 100 MeV. The FADC baselines are stable to within about 1 ADC count, but for very low-energy signals, post-power-cycle baseline shifts can cause a bias in charge reconstruction. This bias is removed at the analysis level, as explained in Ch. 11.

Each FADC channel has memory consisting of 1024 circular buffers, with each buffer sized to store a 4 µs–long digitized waveform. Data is continuously written to one buffer until a trigger signal is received. Following a trigger, the current buffer is frozen and writing resumes on a new buffer. Data from the frozen buffers is recorded for a 256-ns window, a duration chosen to include at least 90% of scintillation light from a single energy deposition. The digitized waveforms from all channels are passed to the ν-DAQ computers for further reconstruction.

9.8.2 Outer veto DAQ

The OV is served by its own data collection and trigger systems. The M64 PMT in each OV module is attached to a custom readout board containing a MAROC2 (Multi-Anode ReadOut Chip) ASIC [82] and an Altera Cyclone FPGA [83]. The MAROC2 applies individual gains to compensate for variations across the 64 PMT channels and compares signals to an analog threshold. Signals exceeding the threshold are passed through a fast shaper to the FPGA, which executes a geometrically-based trigger logic. Generally, this logic requires at least two overlapping scintillator strips to receive hits, a condition met by nearly all muons but only a small fraction of background radioactivity. Meanwhile, analog signals are processed with a slower, more accurate shaper. If the trigger condition is achieved, signals are digitized by FADCs and stored in FPGA memory.

Digitized signal packets are collected through a USB daisy chain. Each chain strings together up to 10 OV modules and a USB board. Signal packets travel unidirectionally around the chain to the USB board, where they are conveyed to the OV DAQ computers. Test and diagnostic signals can also be sent along the USB chain. A vulnerability of the daisy chain design is that one malfunctioning board can prevent readout from all downstream modules. Such a problem occurred in the summer of 2012, when a board in the lowest layer of the far detector OV stopped sending signals. A rerouting of USB connections restored some downstream modules, but one eighth of the far detector lower OV remains nonoperational. Since this portion is covered by another layer of functioning modules, the muon veto efficiency of the OV is minimally affected.

The OV readout boards received clock signals and periodic synchronization pulses from the ν-DAQ trigger master board. This synchronization enables offline merging of the two data streams.

Part III

Far detector oscillation analyses

Chapter 10

Overview of far detector analyses

The Double Chooz far detector began taking data in April 2011. The first oscillation results, based on Gd capture signals in six month of far detector data, were announced in December 2011 and published three months later [28]. This first oscillation fit used both rate and spectral shape information, providing better sensitivity than a completely rate-based fit. For well over a year, Double Chooz remained the only reactor experiment to use the Rate+Shape technique in a $\sin^2 2\theta_{13}$ measurement.

In July 2012, the Double Chooz collaboration published an updated Gd-based analysis, derived from approximately double the livetime of the previous analysis [78]. A $\sin^2 2\theta_{13}$ measurement based on H capture signals, the first from any reactor experiment, was published in 2013 [84]. In addition to Rate+Shape fits, the new Reactor Rate Modulation (RRM) fit was performed on these datasets [85]. Combined fits of the Gd-channel and H-channel data were prepared for conferences in 2013.

The following chapters describe the final set of $\sin^2 2\theta_{13}$ measurements made using data from only the far detector. These analyses include roughly twice the livetime of the previous analyses, along with powerful improvements in signal selection and reductions in systematic uncertainties. Chapters 11 and 12 present common elements of data processing and MC production in the Gd- and H-channel analyses. Chapters 13 and 14 discuss the selection and makeup of the IBD candidates in the Gd and H analyses, respectively. Chapter 15 discusses the Rate+Shape fits in detail, and Ch. 16 briefly presents the RRM fits.

The most recent Gd analysis was the subject of a 2014 publication [86]; methods, figures, and results presented here expand upon information provided in that publication. A paper describing

the latest H analysis was under preparation when this thesis was completed.

Chapter 11

Data preparation

Data collected in the Double Chooz detectors undergoes a series of reconstruction steps before entering oscillation analyses. Signal processing begins with real-time amplification and digitization in the ν-DAQ and OV-DAQ, described in Sec. 9.8. From the DAQs, digitized signals are passed through reconstruction algorithms which attempt to identify the charge, time, and other properties of the original charge deposition. Partway through reconstruction, signals from the ν-DAQ and OV-DAQ are merged into a single data stream. Data that passes basic quality checks proceeds to the analysis level. This chapter describes each step in the path between DAQ and analysis.

11.1 Inner detector and inner veto data reconstruction

11.1.1 Charge and time reconstruction

The first step in signal reconstruction is finding the charge deposited in each ID and IV PMT, along with the time of that deposition. These quantities are determined from the digitized waveforms produced in the FADCs. As explained in Sec. 9.8.1, each waveform covers 256 ns in time bins of width 2 ns. The charge in the i^{th} channel, q_i, is defined as the sum of ADC counts in a 112-ns sub-window of the 256 ns readout window, minus the channel baseline. The baseline, B_i, and its RMS spread, σ_{B_i}, are computed from 256-ns samples collected with a dedicated 1 Hz trigger. The length of the signal integration window is based on the width of a single PE pulse. Since most channels see approximately one PE for signals up to a few MeV, this length optimizes the resolution and efficiency of the charge integration. The position of the integration window is dynamically

determined for each waveform. It is chosen to maximize q_i, subject to two constraints: at least one of the 2-ns time bins must have at least two ADC counts, corresponding to about a third of a PE, and q_i must exceed $\sigma_{B_i}/\sqrt{N_b}$, where $N_b = 56$ is the number of time bins in the window. If these conditions cannot be met, the waveform is discarded as noise.

The time of the signal, t_i, is defined as the time at which the waveform reaches 20% of its maximum value. A correction is applied for the relative PMT time offsets, as measured with the ID and IV light injection systems.

11.1.2 Vertex reconstruction

For each triggered readout, charges and times from all channels in the ID are combined to estimate the location and time of the source event. This event is assumed to be an instantaneous, point-like emission of light that can be fully characterized by $\mathbf{X} = (x, y, z, T, \Phi)$, where ($x$, y, z) is the spatial vertex, T is the time of emission, and Φ is the light intensity in photons per steradian. From this assumption, the charge in PE deposited in the i^{th} PMT (q_i') and the time of deposition (t_i') can be predicted using:

$$\begin{aligned} q_i' &= \Phi \epsilon_i \Omega_i A_i \\ t_i' &= T + \frac{r_i}{c_n} \end{aligned} \tag{11.1}$$

where ϵ_i is the quantum efficiency of the PMT; Ω_i is the solid angle subtended by the PMT from the point of view of the source event; r_i is the distance from the source event to the PMT; A_i is the light transmission amplitude in the intervening medium, including the scintillator and PMT glass; and c_n is the speed of light in that medium. The last three parameters are extracted from radioisotope deployments and bench-top tests of the scintillators.

For each event, a maximum likelihood algorithm is used to find $\hat{\mathbf{X}}$, the best estimator for $\mathbf{X}$ (see Sec. 15.1.1 for background to this approach). A likelihood function can be constructed from solely charge or time information, but including both parameter sets maximizes the accuracy and stability of the technique. The likelihood function is defined as:

$$\mathcal{L}(\mathbf{X}) = \prod_{i:\ q_i=0} f_q(0; q_i') \prod_{i:\ q_i>0} f_q(q_i; q_i') f_t(t_i; t_i', q_i') \tag{11.2}$$

The first product includes the PMT channels that recorded no charge in this event, and the second product includes the rest of the PMTs. The probability of recording a PMT charge of

q_i, given a predicted PMT charge of q'_i, is expressed with $f_q(q_i; q'_i)$. Similarly, the probability of recording a PMT time of t_i given a predicted charge of q'_i and time of t'_i, is expressed with $f_t(t_i; t'_i, q'_i)$. Both of these functions are developed from laser calibrations performed with the z-axis system. The best estimator of the vertex, $\hat{\mathbf{X}}$, is found by minimizing $-\ln \mathcal{L}(\mathbf{X})$.

The performance of this reconstruction algorithm has been evaluated with radioisotope deployments in the z-axis and guide tube systems. When sources are positioned in the center of the NT, typical spatial resolutions are 20–30 cm.

11.1.3 Linearized PE reconstruction

The next stage of reconstruction converts raw charge in each channel to an equivalent number of PE, correcting for nonlinearity that arises from imperfect baseline subtraction. The number of PE in each event is defined[1] as:

$$N_{PE}^{data} = \sum_i \frac{q_i}{G_i(q_i, T)} \tag{11.3}$$

where i runs over all PMT channels, occasionally excluding a small number which have been tagged as problematic by online checks of the waveforms. In this expression, G_i is a function describing the gain of of the i^{th} channel. The gain function is charge-dependent because the quantized nature of the baseline sampling leads to nonlinearity for low-energy signals. The time dependence comes mainly from baseline shifts following power cycles.

Gain functions are determined using the light injection system, run with multiple light intensities and injection points. For each light injection trial, the mean μ_i and standard deviation σ_i are found for the charge distribution in the i^{th} PMT. Since Poissonian photon statistics imply $\mu_i \propto g_i N_i$ and $\sigma_i \propto g_i \sqrt{N_i}$, where N is the number of observed photons, these quantities can be combined to find:

$$g_i = \alpha \sigma_i^2 / \mu_i \tag{11.4}$$

The parameter α corrects for spread in the charge distribution due to single PE width and electronic noise. Its value is fixed by an analysis of the 2.22 MeV signal from spallation neutron captures on H. Figure 11.1 shows one typical set of g_i. As shown in the figure, the set of g_i is fit with a piecewise linear function, G_i. New gain functions are determined following each detector power cycle.

[1] The superscript *data* distinguishes quantities from the MC equivalents introduced in Ch. 12.

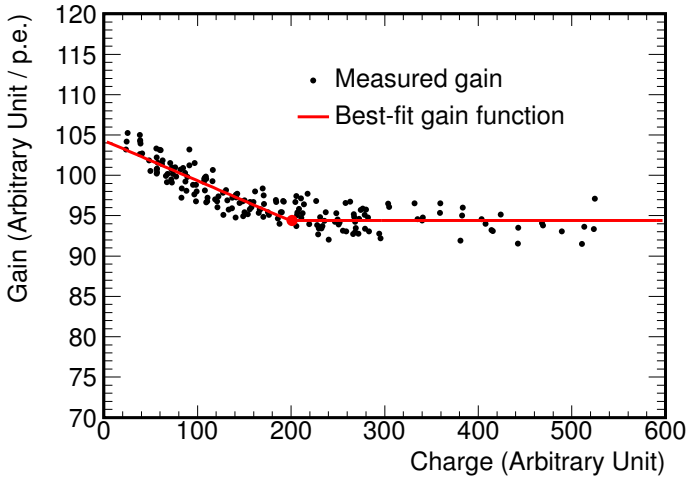

Figure 11.1: Observations of gain, $g_i = \alpha\sigma_i^2/\mu_i$, versus charge, q_i, in one PMT over multiple light injection trials (black points). The observations are fit with a piecewise linear function, G_i (red lines, with intersection indicated by a red point). Figure from [78].

11.1.4 Uniformity correction

Although the detectors were designed to provide a nearly uniform response across the entire active volume, geometric effects and and differences between the NT and GC liquids cause up to $\sim 10\%$ variations in N_{PE}^{data} reconstructed for equivalent events in different locations. These variations would be immaterial at the analysis level if they were exactly reproduced in the detector Monte Carlo (MC). However, MC tuning can produce only approximate agreement, so it is advantageous to calibrate out as much volume-dependence as possible from both data and MC.

Volume dependence of N_{PE}^{data} is corrected by the function $f_u^{data}(\rho, z)$. This function is extracted from spallation neutron captures on H, which occur in large numbers throughout the NT and GC. The first step is dividing the NT and GC volumes into a total of 108 ring-shaped sub-volumes, based on nine divisions in ρ and 12 divisions in z. Next, in each sub-volume, the mean N_{PE}^{data} of H captures is determined, and the ratio of this value to N_{PE}^{data} at the center of the detector is computed. This ratio is smoothly interpolated between sub-volumes using the Delauney triangles algorithm. Figure 11.2 shows the resulting "map" of $f_u^{data}(\rho, z) = N_{PE}^{data}(\rho, z)/N_{PE}^{data}(0, 0)$.

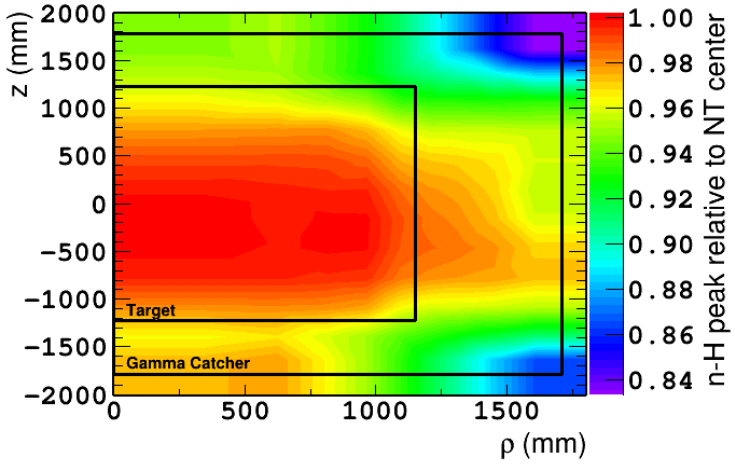

Figure 11.2: The volume-dependent energy correction function $f_u^{data}(\rho, z)$ for data. Black boxes mark the outlines of the NT (Target) and GC. Figure from [78].

11.1.5 Absolute energy scale correction

The absolute energy scale is defined by the 2.223 MeV gamma ray resulting from neutron capture on H, as measured from a ^{252}Cf source deployed at the center of the detector. This signal yields a charge-to-energy conversion factor of $1/f_{abs}^{data} = 186.2$ PE MeV^{-1}. The same definition is used in MC, anchoring the energy scales of both data and MC in the region most sensitive to θ_{13}-driven oscillations.

11.1.6 Stability correction

The detector response varies over time due to discrete changes in electronic gains, usually following power cycles, and more gradual trends such as scintillator aging. These time variations are not simulated in MC. Therefore, if uncorrected, time variations in data would broaden the overall energy resolution with respect to the MC. In analyses with multiple time bins, they could also cause discrepancies between the data and MC energy scales. To minimize these data-MC differences, time variations are corrected as much as possible in data. The correction function, $f_s(E_{vis}^{data}, T)$, is constructed from a time series of data samples, including spallation neutron captures on Gd and H and alpha particles from ^{212}Po decays. The latter are strongly quenched in the scintillator, resulting in a visible energy of about 1 MeV.

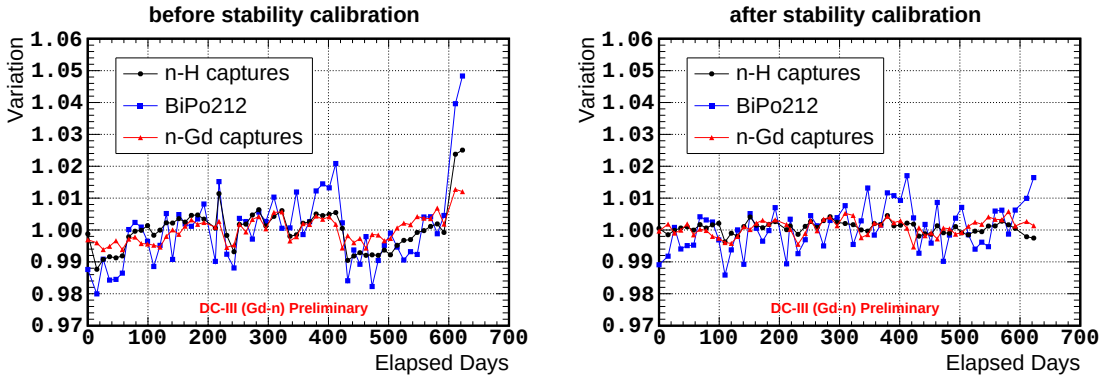

Figure 11.3: A comparison of the time variation in signals from neutron capture on H (black), alpha decays of ^{212}Po (blue), and neutron captures on Gd (red) before application of the stability correction function f_s (left) and after its application (right). Figure adapted from [86].

The function f_s is constructed as:

$$f_s = 1 + (\delta f_s)^{gain} + (\delta f_s)^{scint} \tag{11.5}$$

The term $(\delta f_s)^{gain}$ corrects for changes in the mean gain of all channels and can be viewed as an adjustment to α in Eq. 11.4. Its functional form is:

$$(\delta f_s)^{gain} = \delta\alpha(T)(p_0 + p_1 E_{vis}^{data}) \tag{11.6}$$

Here, $\delta\alpha(T)$ reflects the change in the H capture peak position, which is measured at regular time intervals. Just as the basic gain functions G_i are known to be nonlinear, time-dependent changes in gain are expected to affect low-energy signals more than high-energy signals. The second factor in Eq. 11.6 accommodates this effect. Values of $p_0 = 0.7866$ and $p_1 = 0.07101/\text{MeV}$ are determined from fits including ^{212}Po decays and the Gd capture peak.

The term $(\delta f_s)^{scint}$ attempts to correct for remaining time variation, including changes in scintillator response. It has the energy-independent form $(\delta f_s)^{scint} = mT$, where T is measured in days since the reference point of May 15, 2012. A fit to the H capture energy as a function of time yields $m = -8.24 \times 10^{-6}/\text{day}$. Figure 11.3 shows the improvement in detector response stability following the application of f_s.

11.1.7 Summary of energy scale in data

Including the corrections described in Secs. 11.1.3–11.1.6, the relationship between observed charge in PE and visible energy in MeV is:

$$E_{vis}^{data} = N_{PE}^{data} \cdot f_u^{data}(\rho, z) \cdot f_{abs}^{data} \cdot f_s(E_{vis}^{data}, T) \tag{11.7}$$

11.2 Outer veto data reconstruction and merging

OV data is reconstructed from the digitized signals collected along the OV-DAQ daisy chains. As noted in Sec. 9.8.2, normal trigger settings require signals to include hits in at least two overlapping scintillator strips. The first step of reconstruction subtracts the baseline from each channel, using pedestal measurements made at the beginning of each OV run. Next, signals are compared to an offline threshold that requires at least 73 ADC counts, equivalent to ~ 1.5 PE, in each channel of an overlapping-strip pair. Signals are then time-ordered, correcting the small amount of misordering that arises within daisy chains. At this stage, signals from all daisy chains are merged and events are defined from signals recorded within 48 ns of one another.

11.3 Muon track reconstruction

Analyses of cosmogenic backgrounds, especially the long-lived beta-neutron emitters ^{9}Li and ^{8}He, rely on reconstructions of muon tracks that traverse the detectors. Two reconstruction algorithms have been developed to deduce these tracks from the spatial and temporal distribution of PMT signals. One algorithm uses information exclusively from the ID. Tracks reconstructed by this method are used in the estimation of ^{9}Li and ^{8}He background rates. The other algorithm, described fully in [87], includes information from the IV and OV as well as the ID. This algorithm finds applications in a likelihood used to reject ^{9}Li and ^{8}He events and in the measurement of the ^{9}Li and ^{8}He energy spectrum.

11.4 Dataset selection

The Gd and H analyses include data taken since the inauguration of the Far Detector on April 13, 2011. The last day of data included in these analyses is January 15, 2013. Most data was taken

in hour-long runs, although some runs were cut short by DAQ problems or other interruptions. A small number of runs in this period are excluded from analysis because they are extremely short (less than 300 s in duration), experienced specific failures during reconstruction, or were taken under atypical detector conditions. The final dataset contains 12304 runs, covering a runtime of 489.51 days.

This runtime includes 7.53 days in which both Chooz reactors were shut down. One 22.5-hr period of reactor-off time occurred on October 22, 2011, during a refueling period for B1, when B2 was shut down to be serviced. The second period lasted from May 27 to June 4, 2012, when B2 was off for a lengthy repair and B1 was temporarily shut down to address a separate issue. Because backgrounds in Double Chooz are entirely independent of the reactors, these periods allow direct measurement of the background rate and spectrum.

Chapter 12

Signal simulation

In the single-detector phase of Double Chooz, simulation of the IBD signal plays an essential role in oscillation analyses. This simulation is divided into two parts: modeling of antineutrino production in the reactors and detector response to IBD events. This chapter begins by describing the reactor model, which combines information from the reactor operator, detailed simulations, and results of previous reactor-based experiments. The product of this stage is the predicted rate and prompt energy spectrum of IBD interactions in the far detector. This prediction is passed through a full detector response MC, including light propagation, readout electronics, and event reconstruction. Later sections of this chapter details these steps. The final product of the signal simulation is an ensemble of IBD events with an extensive set of reconstructed and MC truth information.

12.1 Reactor model

12.1.1 Information from reactor operator

The reactor operator, EDF, has agreed to provide the Double Chooz collaboration with information about the fuel loading and operational dynamics of each reactor. For each fuel cycle, EDF furnishes a map of the various types of fuel assemblies loaded into the core. For previously used assemblies, EDF reports burnup, a measure of the amount of energy extracted from fissionable material. EDF also grants controlled access to a database of more than 3000 parameters monitored during reactor operation. Among these measurements are the position of fuel rods, boron concentration in the primary coolant, and thermal power, P_{th}.

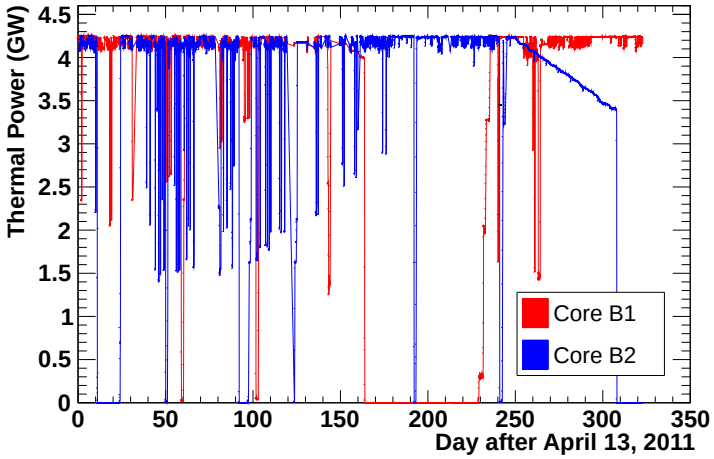

Figure 12.1: Thermal power of reactors B1 and B2 as a function of time during the first $\sim$ 300 days of far detector data-taking. Reactor B1 is shut down for a refueling period covering days $\sim$ 160–240. Beginning around day 240, the power of B2 was ramped down and then brought to zero for a refueling period.

Assessments of the thermal power are made at time intervals of one minute or less. They are based on multiple measurements taken inside the reactor core, including temperature of the primary coolant and the neutron flux. The in-core instrumentation is checked weekly with a second type of thermal power assay. This method analyzes the enthalpy balance between the primary and secondary coolants in the steam generators. It is more precise than the in-core assessments, and if results of the in-core and enthalpy balance measurements disagree by more than the uncertainty on the latter, the in-core system is re-calibrated. Because the upper bound of safe operating power is limited by the precision of thermal power measurements, EDF has intensively studied and optimized these systems [88]. The final thermal power uncertainty is approximately 0.5%. Uncertainties are slightly larger on lower power measurements, as the enthalpy balance test is performed only at full reactor power. Figure 12.1 shows examples of the thermal power profiles of both reactors.

12.1.2 Reactor core simulations

Several factors needed to predict antineutrino production do not appear directly in the EDF database. One set of important parameters, called α_k, describes the fraction of fissions occurring in

the four main fissioning isotope. Here, and in the rest of this chapter, k runs over ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu. Two reactor modeling packages have been used to estimate the α_k and other reactor parameters. The MCNP Utility for Reactor Evolution (MURE) package simulates three-dimensional movement of neutrons within the core using the Monte Carlo N-Particle (MCNP) code [89]. The DRAGON package employs faster, deterministic algorithms that solve neutron transport equations on a two-dimensional lattice [90]. Nominal predictions for the present analyses come from MURE, while DRAGON confirms these predictions and aids in systematic uncertainty assessment. Both codes have been benchmarked against a destructive assay of a PWR core [91].

The first simulation task is determining start-of-cycle fuel compositions, as EDF reports burnup in used fuel assemblies but not isotope fractions. Simulations using MURE and DRAGON are performed for cores loaded entirely with fresh fuel, and the resulting fuel assembly compositions are recorded. To validate these simulations, results for burnup and other parameters are compared to a similar run of APOLLO2-F, the proprietary code used by EDF. The difference between compositions calculated by MURE and DRAGON is treated as a systematic uncertainty on the α_k.

Next, MURE is configured to model the actual fuel cycles coincident with Double Chooz data-taking. Inputs include initial fuel composition, thermal power, control rod positions, and the temperature, density, and boron concentration of the primary coolant. Select information from nuclear databases is also included. The core evolution is simulated in time steps of 6 to 48 hours, depending on reactor conditions. The simulated values of α_k are written to a database for use in reactor antineutrino predictions. These values can be combined with $\langle E_f \rangle_k$, the mean energy released per fission of the k^{th} isotope, to determine the overall mean energy released per fission, $\langle E_f \rangle$:

$$\langle E_f \rangle = \sum_k \alpha_k \langle E_f \rangle_k \tag{12.1}$$

Table 12.1 lists simulated α_k for a representative portion of Double Chooz data-taking, along with $\langle E_f \rangle_k$ computed from the fractional yield and atomic mass excesses of the fission fragments [92].

Systematic errors on the α_k and $\langle E_f \rangle$, and correlations between the α_k uncertainties, are estimated by modeling 1σ excursions on each of the simulation inputs. Uncertainties from the various inputs are assumed to be uncorrelated.

Table 12.1: Characteristics of the four main fissioning isotopes (indexed by k) in the Chooz reactors: the mean energy per fission, $\langle E_f \rangle_k$ [92]; the mean fission fraction in a representative portion (approximately the first year) of Double Chooz data-taking time, $\langle \alpha_k \rangle$; and the mean fission fraction in the Bugey4 experiment, $\langle \alpha_k \rangle^{Bugey}$ [54]. Uncertainties on the $\langle \alpha_k \rangle^{Bugey}$ are not given in [54] and do not greatly affect reactor flux uncertainty in Double Chooz.

Isotope, k	$\langle E_f \rangle_k$ (MeV)	$\langle \alpha_k \rangle$	$\langle \alpha_k \rangle^{Bugey}$
^{235}U	201.92 ± 0.46	0.496 ± 0.016	0.538
^{239}Pu	209.99 ± 0.60	0.351 ± 0.013	0.328
^{238}U	205.52 ± 0.96	0.087 ± 0.006	0.078
^{241}Pu	213.60 ± 0.65	0.066 ± 0.007	0.056

12.1.3 Antineutrino spectra from fission

The P_{th} measurements from EDF, combined with the α_k and $\langle E_f \rangle$ predictions from the reactor simulation, yield the instantaneous rate of fissions, R_k, from the k^{th} isotope:

$$R_k = \frac{P_{th}}{\langle E_f \rangle} \alpha_k \tag{12.2}$$

To predict the resulting antineutrino flux, the R_k are combined with the total antineutrino spectra, $S_k(E_\nu)$, produced following fissions of the k^{th} isotope. These spectra are complicated because they sum over thousands of beta decay branches, as noted in Sec. 6.1. Two techniques have been developed to estimate the S_k. The summation method constructs S_k from individual beta branches described in nuclear databases. This construction relies on branching ratios and spectrum shapes which are not completely known. In particular, about 30% of the antineutrino flux comes from forbidden decays, whose spectrum shapes suffer from large nuclear uncertainties. The second method of S_k construction, the conversion method, begins by measuring the total electron spectrum following fissions of the k^{th} isotope. The electron spectrum is translated into the total antineutrino spectrum through virtual beta branches which do not correspond to physical decays.

Although previous Double Chooz analyses used the summation method for ^{238}U [58], the present analyses rely solely on the conversion method. The S_k for $k =$ ^{235}U, ^{239}Pu, and ^{241}Pu are derived from measurements made in the 1980s at the Institute Laue Langevin (ILL) research reactor [93, 94, 95]. In the ILL studies, thin foils containing each isotope were exposed to the thermal neutron

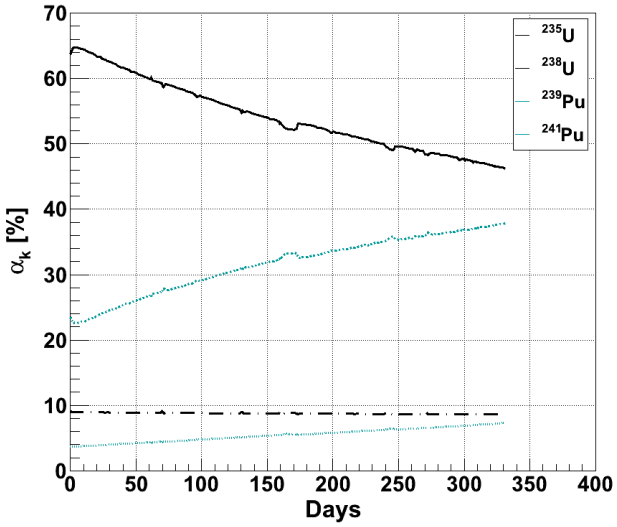

Figure 12.2: Fission fractions, α_k, of the four dominant isotopes in reactor B2 over the first $\sim$ 300 days of far detector data-taking. This evolution is typical of a single PWR cycle.

flux, and the electrons from subsequent beta decays were analyzed with a sensitive spectrometer. These challenging measurements required calorimetric accuracy over more than four decades of energy. The ILL electron spectra were converted to antineutrino spectra using the technique of [96], including corrections for long-lived fission products described by [46]. The S_k for $k =$ ^{238}U must be derived from a different type of measurement, because ^{238}U is fissioned by fast neutrons, with energies of at least $\sim$ 5 MeV, rather than thermal neutrons. Recently, the ^{238}U antineutrino spectrum between 3 MeV and 7.5 MeV was derived from a fast neutron experiment performed at the Heinz Maier-Leibnitz Neutron (FRM-II) research reactor [97]. For the Double Chooz analyses, an exponential-polynomial fit was used to extend this model beyond 3–7.5 MeV. The S_k used in the current analyses are binned at a resolution of 0.25 MeV.

Uncertainties on the S_k are of order 3%, with significant energy dependence and bin-to-bin correlations. The higher-energy portions of the spectra carry the largest uncertainties. On their own, the S_k uncertainties would be the most severe limitation on the Double Chooz signal prediction. A technique described in Sec. 12.1.5 helps to reduce this limitation.

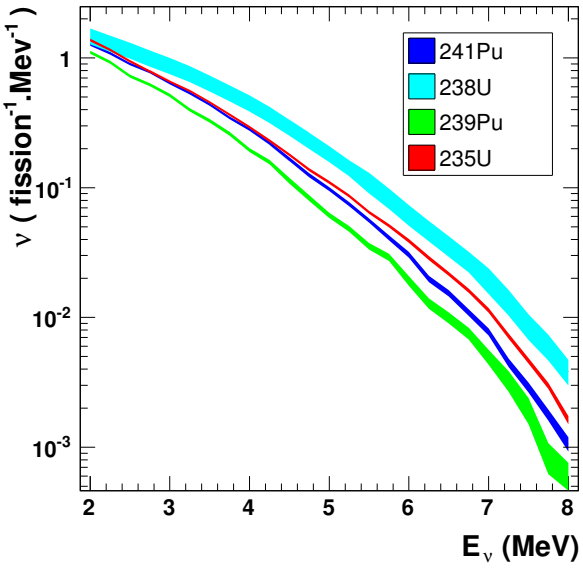

Figure 12.3: Antineutrino spectra generated by the fission products of the four dominant isotopes. The width of each curve indicates systematic uncertainty. The ^{235}U and Pu spectra are calculated by the procedure of [96], with corrections for long-live fission products as in [65]. The ^{238}U spectrum shown here, an *ab initio* calculation from [65], is not used in the current analyses; the empirically derived spectrum of [97] is used instead.

12.1.4 Average cross section per fission

The spectrum of detected antineutrinos is a convolution of the S_k and the IBD cross section. Double Chooz uses the cross section approximation of Eq. 6.6, with the value of κ quoted in Sec. 6.2. It is convenient to define quantity $\langle\sigma_f\rangle_k$, the flux-averaged IBD cross section per fission of the k^{th} isotope:

$$\langle\sigma_f\rangle_k = \int_0^\infty \mathrm{d}E_\nu S_k(E_\nu)\sigma_{IBD}(E_\nu) \tag{12.3}$$

The total flux-averaged cross section per fission is then:

$$\langle\sigma_f\rangle = \sum_k \alpha_k \langle\sigma_f\rangle_k \tag{12.4}$$

12.1.5 Bugey4 anchor

In the two-detector phase, the impact of the S_k uncertainties will be strongly suppressed by near detector measurements. In the single-detector phase, Double Chooz obtains some of the same benefit by treating the Bugey4 experiment as a virtual near detector, a technique first used in CHOOZ [26]. In the early 1990s, Bugey4 measured the rate of IBD interactions at distance of 15 m from a 2800 $\mathrm{MW_{th}}$ PWR near Lyon, France. From this data, the Bugey4 collaboration inferred a flux-averaged cross section of $\langle\sigma_f\rangle^{Bugey} = 5.75 \times 10^{-43}$ cm^2/fission, with 1.4% uncertainty [54]. The collaboration also reported the mean fission fractions, $\langle\alpha_k\rangle^{Bugey}$, in their observation period.

In the Double Chooz antineutrino prediction, $\langle\sigma_f\rangle^{Bugey}$ is used to renormalize the $\langle\sigma_f\rangle$ calculated through Eq. 12.3. A correction is made to account for the small difference between the Double Chooz α_k and $\langle\alpha_k\rangle^{Bugey}$. The Bugey4 renormalization factor, f^{Bugey}, is then:

$$f^{Bugey} = \frac{\langle\sigma_f\rangle^{Bugey} + \sum_k \left(\alpha_k - \langle\alpha_k\rangle^{Bugey}\right)\langle\sigma_f\rangle_k}{\sum_k \alpha_k \langle\sigma_f\rangle_k} \tag{12.5}$$

Using this rate anchor roughly halves the effective normalization uncertainty on the S_k. It also renders the far detector insensitive to the reactor antineutrino anomaly introduced in Sec. 5.2.2. As a result, Double Chooz can measure θ_{13}-driven oscillations without ambiguity from possible sterile neutrino signals. If sterile neutrino sensitivity is desired, as in the measurement suggested in Sec. 19.1, the antineutrino rate may be computed without f^{Bugey}.

12.1.6 IBD event generation

Following the form of Eq. 6.8, and incorporating Eqs. 12.2, 12.4, and 12.5, the instantaneous spectral density of antineutrinos from the r^{th} reactor (r = B1, B2) detected in the far detector (in events per unit time and energy) is predicted to be:

$$\rho_r(E_\nu) = \epsilon N_p \frac{1}{4\pi L_r^2} \frac{(P_{th})_r}{\langle E_f \rangle_r} f^{Bugey} \sum_k (\alpha_k)_r \left(\langle \sigma_f(E_\nu) \rangle_k\right)_r \tag{12.6}$$

In this expression, $(P_{th})_r$, $\langle E_f \rangle_r$, f^{Bugey}, $(\alpha_k)_r$, and $\left(\langle \sigma_f(E_\nu) \rangle_k\right)_r$ are time-dependent. The factor ϵ represents signal detection efficiency, which is modeled through the detector response simulation described in Sec. 12.2. The baselines L_r are measured as described in Sec. 7.3.

For each run in the dataset, the number of antineutrinos in each prompt E_{vis}^{MC} bin is predicted by integrating Eq. 12.6 over the corresponding time period and energy range. If both reactors are operating, the predicted spectrum sums over both contributions. With both reactors operating at full power, roughly two IBD interactions followed by Gd capture are expected to occur each hour. About twice as many IBD events will be followed by H capture in the NT or GC. To minimize the impact of statistical fluctuations in the MC sample, IBD events are generated at 100 times the predicted rate and later scaled appropriately.

For each event, E_ν is drawn uniformly from the relevant $\langle \sigma_f \rangle_r$ bin. Each antineutrino is assigned a creation point randomly drawn from the fission profile in the reactor core. Uncertainty on the barycenter of fission is insignificant for baseline assignment. An IBD interaction point is assigned according to the proton density in the ID, including the volumes of the NT, GC, Buffer, and acrylic vessels. Center-of-mass kinematics are calculated from the incident antineutrino energy and a random choice of the outgoing positron direction. The positron and neutron momenta are then boosted to the laboratory frame.

12.2 Detector model

12.2.1 Scintillator and PMT simulation

For each predicted IBD event, the interaction vertex, positron momentum, and neutron momentum are passed to a custom Geant4 (version 9.2.p02) simulation of the detector response [98]. This package models positron, neutron, and photon interactions in the detector and light collection at

the PMTs. The hadron interaction treatment is similar to the Geant4-supplied quark-gluon string precompound (QGSP) model, including the Bertini cascade package for intranuclear processes and the NeutronHP package for neutron transport [99]. Treatment of neutron thermalization has been customized, following the model of [100], and the neutron radiative capture model has also been improved. The custom optical model includes detailed light waveforms, relevant emission and re-emission spectra, Birks' law quenching [101], and a precise photocathode treatment. Parameters in the optical model, including light yields of the scintillators, attenuation lengths, indices of refraction, and material reflectivities, are derived from a variety of bench-top measurements and *in situ* calibrations [102, 70, 103]. The detector geometry is modeled in intricate detail, reflecting measurements made during construction. The dimensions of various detector components were verified during installation, and a photographic survey determined the PMT locations and orientations to sub-millimeter accuracy.

12.2.2 Readout system simulation

The next step in signal simulation accounts for the readout electronics, including the PMT circuits, FADCs, trigger, FEE, and DAQ. The readout model uses probability density functions (PDFs) which characterize the response of each ID channel to a single PE signal. The PDFs were developed with data taken in a PMT test stand. Simulated PMT signals are processed into digital waveforms, analogous to those created in the FADCs. Details of the real readout system, including electronic noise and channel-by-channel gain and baseline variations, are incorporated into the simulation.

12.2.3 Energy scale in MC

At this stage, simulated FADC waveforms are processed into calibrated visible energy. Raw charge and timing information in each channel is reconstructed by the same procedure used for data, as described in Sec. 11.1.1. The vertex reconstruction algorithm of Sec. 11.1.2 is applied next. A uniformity correction function, like the one described in 11.1.4, is generated from a sample of H captures following simulated IBD events. The absolute energy scale in MC is defined, as in data, from the H capture peak position, yielding $1/f_{abs}^{MC} = 186.6$ PE/MeV. These factors produce a first-order energy scale, defined as:

$$E_{vis}^{MC} = N_{PE} \cdot f_u^{MC}(\rho, z) \cdot f_{abs}^{MC} \tag{12.7}$$

Finally, the MC receives corrections to minimize nonlinearity relative to data. Nonlinearity results from inexact modeling of two processes: charge reconstruction and light production and propagation. The first source affects all signals, regardless of particle type. A compensating factor, f_{QNL}, is derived from a z-axis calibration run with a ^{252}Cf source at the center of the detector. Most neutrons from this source capture on Gd, with a smaller fraction capturing on H. Although the total visible energy of each Gd capture is about 8 MeV, the average energy of individual photons in the signal is approximately 2.2 MeV, very similar to the photon released upon a neutron capture on H. A difference between the data and MC energy scales for these signals must therefore originate in the readout electronics or charge integration algorithm, rather than in the light model. From a comparison of the ^{252}Cf calibration in data and MC, the following MC correction is derived:

$$f_{QNL} = b_{QNL} + c_{QNL} E_{vis}^{MC} \tag{12.8}$$

where $b_{QNL} = 0.9949 \pm 0.0036$ and $c_{QNL} = 0.0023 \pm 0.0006/\text{MeV}$ with a correlation of $\rho_{b_{QNL},c_{QNL}} = -0.60$. Figure 12.4 shows the data-to-MC energy ratio after f_{QNL} is applied.

The remaining data-MC discrepancy is believed to result from imperfect scintillator modeling, such as in the expected ratio of scintillation to Cherenkov light, L, or the Birks quenching factor, k_B. These factors should have somewhat particle-dependent effects, with distinct impacts for electrons or positrons and strongly quenched signals such as alpha particles. They are also expected to differ between the NT and GC scintillators. Fig. 12.4 shows the discrepancy in the NT. A nearly opposite pattern has been observed in the GC, with the data-to-MC ratio rising to approximately 1.015 at low energies. Consequently, separate light nonlinearity treatments are used for the data in the Gd analysis, where signal is confined to the NT, and in the H analysis, which includes signal in both the NT and GC.

In the Gd analysis, a correction factor optimized for positrons, and also appropriate for electrons and photons, is applied to the prompt E_{vis}^{MC} spectrum of IBD candidates. To produce this function, called f_{LNL}, variations on NT calibration MC samples are produced with L and k_B changed within their uncertainties. Combinations of L and k_B which yield reasonable agreement with calibration data are used to generate positron MC samples. Comparing these samples to the nominal positron MC motivates the correction factor:

$$f_{LNL} = \frac{a_{LNL}}{E_{vis}^{MC}} + b_{LNL} \tag{12.9}$$

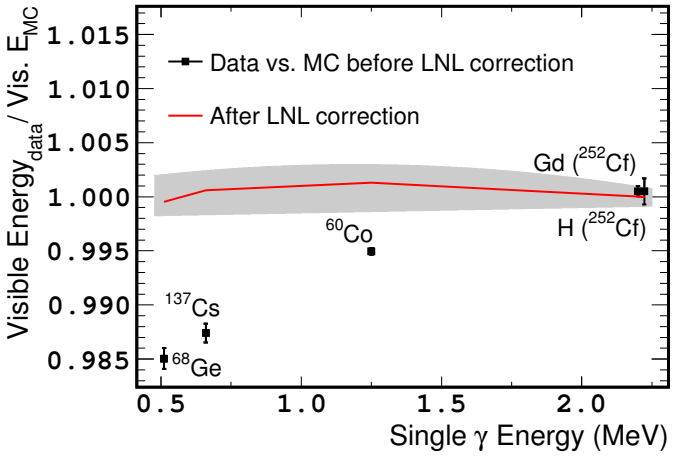

Figure 12.4: The ratio of visible energy in data to visible energy in MC for five calibration points, shown as a function of single-gamma energy, before the light nonlinearity calibration is applied (black points with statistical error bars). The average gamma multiplicity of Gd cascades is taken from MC. The red line connects the same five points after the light nonlinearity correction from the Gd analysis is applied. The gray band indicates systematic uncertainty on the correction factor. Figure from [86].

where $a^{Gd}_{LNL} = -0.027 \pm 0.0062$ MeV and $b^{Gd}_{LNL} = 1.008 \pm 0.0026$, with a correlation $\rho^{Gd}_{a_{LNL},b_{LNL}} = -0.81$. Application of f_{LNL} brings the data and MC energy scales for the Gd analysis into close agreement, as displayed in Fig. 12.4.

Producing a similar correction factor would be more difficult in the H analysis, since its signal population spans two scintillator volumes, and because calibration data does not exist for many locations in the GC. Moreover, as issues discussed in Ch. 17 came to light, a precisely calibrated energy scale became less essential to the H analysis strategy. As a result, no light nonlinearity correction is applied to H data, or equivalently, $a^{H}_{LNL} = 0$ and $b^{H}_{LNL} = 1$ are used in Eq. 12.9. The effective light nonlinearity in the H analysis is most likely an average of the pure NT and GC curves, and conservative uncertainties are assigned to a^{H}_{LNL} and b^{H}_{LNL} to cover both of those extremes. In MC-based tests, values of $\sigma^{H}_{a} = 2(a^{Gd} + \sigma^{Gd}_{a}) = 0.0664$ and $\sigma^{H}_{b_{LNL}} = 2(b^{Gd}_{LNL} + \sigma^{Gd}_{b_{LNL}}) = 0.0212$ have been shown to provide sufficient freedom to accommodate NT- and GC-like light nonlinearity curves.

Including all correction factors, the nominal MC energy scale for oscillation analyses is given by:

$$E^{MC,corr}_{vis} = E^{MC,corr}_{vis} \cdot f_{QNL}(E^{MC}_{vis}) \cdot \begin{cases} f_{LNL}(E^{MC}_{vis}) & \text{if a positron} \\ 1 & \text{otherwise} \end{cases} \tag{12.10}$$

The final definitions of visible energy in data and MC produce excellent agreement, both in absolute energy scale and in resolution, as portrayed in Fig. 12.5. The energy resolutions, σ, of data and MC have been fit with the function $\sigma/E_{vis} = \sqrt{s_0^2/E_{vis} + s_1^2 + s_2^2/E_{vis}^2}$. In this description, s_0 accounts for fluctuations in photoelectron statistics, s_1 accounts for electronic noise, and s_2 accounts for other instrumental effects which are not energy-dependent. For data (MC), the fit yields $s_0 = 0.077 \pm 0.002$ MeV$^{1/2}$ ($s_0 = 0.077 \pm 0.002$ MeV$^{1/2}$), $s_1 = 0.018 \pm 0.001$ ($s_1 = 0.018 \pm 0.001$), and $s_2 = 0.017 \pm 0.011$ MeV ($s_2 = 0.024 \pm 0.006$).

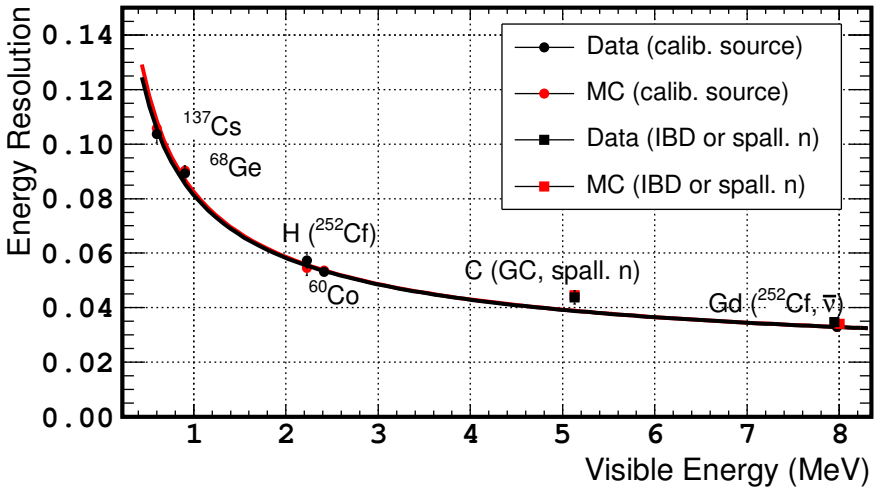

Figure 12.5: The energy resolution of calibration points in data and MC, from Gaussian fits, versus visible energy. Black (red) circles represent z-axis source deployments in the center of the detector in data (MC). The black (red) square represents spallation neutron captures on carbon in data (delayed signals from IBD events with neutron captures on carbon in MC). Because these captures occur throughout the entire NT and GC volume, the overall resolution is slightly worse than the resolution of sources at the detector center. The black (red) curve is a fit to the data (MC). Figure from [86].

Chapter 13

Gd-channel candidates

The present Gd-based analysis involves the highest signal purity and lowest systematic uncertainties of any single-channel analysis yet performed in Double Chooz. The signal selection algorithm is optimized to collect more than 98% of IBD events followed by Gd capture during the analysis livetime. Backgrounds are reduced through targeted veto strategies. Multiple, complementary techniques have been developed to assess the signal detection efficiency with precision better than 1%. This chapter outlines the selection of Gd-channel IBD candidates, estimates of residual backgrounds, and assessment of efficiency-related uncertainties.

13.1 Signal selection

As noted in Secs. 9.8.1 and 12.1.6, signals in the ID trigger hundreds of readouts each second, while just one or two Gd-channel IBD events are expected each hour. Highly selective cuts are needed to cull true IBD signals from the large background population. The basic strategy is to choose events which share the most distinctive features of the IBD signal: the close time correlation of the prompt positron and delayed neutron signals ($\tau \sim 30\,\mu s$) and the high energy of the neutron capture on Gd ($E_{vis} \sim 8$ MeV). Coincidences of unrelated signals, such as background radioactivity and spallation neutrons, sometimes pass the time correlation requirement. However, most uncorrelated energy depositions occur below the Gd capture energy.

The more problematic backgrounds involve correlated signals, especially if the later one is naturally high in energy. These backgrounds include fast neutrons which cause proton recoils before

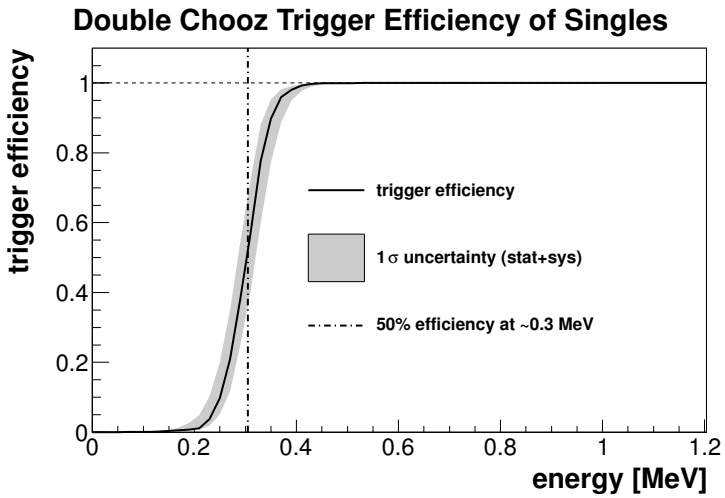

Figure 13.1: ID trigger efficiency as a function of visible energy, extracted from a sample of normal data runs and fit with an error function. Figure from [104].

capturing, muons which stop and decay in the detector, and long-lived cosmogenic isotopes which undergo beta-neutron decays. Each of these backgrounds can be reduced with specific cuts. At the same time, limiting the number of cuts, and placing them in regions that are relatively insensitive to the details of the MC, minimizes systematic uncertainties on the detection efficiency. The Gd selection strategy is a balance of these two objectives.

13.1.1 Pre-selection

Before the IBD selection is performed, signals are subjected to a set of basic quality cuts. Requiring $E_{vis} > 0.4$ MeV restricts the analysis to a regime in which the ID trigger is nearly 100% efficient, with small uncertainty, as illustrated in Fig. 13.1. Signals with $E_{vis} > 20$ MeV or an inner veto energy of $E_{IV} > 16$ MeV are identified as cosmic ray muons and rejected. Signals occurring within the first 1 ms after a muon are also rejected, because they are very likely to be spallation products or fast-decaying cosmogenic isotopes.

Given the muon rate of ~ 40 Hz, the 1-ms veto creates a modest amount of analysis deadtime. Accounting for deadtime due to muon vetoes of prompt IBD signals yields a Gd analysis livetime of 460.67 days in the reactor(s)-on period and 7.24 days in the reactors-off period. The normalization of the MC is corrected to account for this deadtime, along with the small probability of a muon

occurring between prompt and delayed signals. The MC correction factor is $C^{\mu-veto}_{data/MC} = 0.9551$, with negligible uncertainty.

Light noise, the spontaneous light emission from PMT bases described in Sec. 9.4, is removed using low-level charge and time cuts. These criteria target the highly localized, peripherally-located topology of light noise events. The light noise cuts also reject 0.01% of IBD events, and the MC is corrected accordingly with $C^{LN}_{data/MC} = 0.9999$. All pre-selection cuts are summarized in Tab. 13.1.

13.1.2 IBD selection

Signals that pass the pre-selection are scanned for IBD-like events. IBD candidates must include two signals with close time correlation, loose space correlation, and specific prompt and delayed energies. In the Gd analysis, the time coincidence requirement is $0.5\,\mu s < \Delta T < 150\,\mu s$, where $\Delta T = T_d - T_p$ and T_p (T_d) is the time of the prompt (delayed) signal. The upper bound for ΔT is chosen to maximize signal efficiency while minimizing systematic uncertainties from the spill in/out current (to be introduced in Sec. 13.2.4).

A multiplicity cut prevents ambiguity in event definitions and removes some muon-correlated backgrounds. This cut demands that no signals occur in the 200 µs preceding the prompt signal and that no signals other than the delayed signal occur in the 600 µs after the prompt signal. This requirement introduces a small signal inefficiency, equal to the rate of uncorrelated signals, 13.22 Hz, times the length of the veto window, 800 µs.

An additional cut requires a loose correlation of the reconstructed prompt and delayed vertices, (x_p, y_p, z_p) and (x_d, y_d, z_d). Most IBD events have $\Delta R = \sqrt{(x_d - x_p)^2 + (y_d - y_p)^2 + (z_d - z_p)^2}$ of a few tens of centimeters or less, so the requirement of $\Delta R < 100$ cm accepts nearly all signal while rejecting accidental coincidences.

The energies of Gd captures are centered around 8 MeV, but the capture peak is broadened by the energy resolution of the detector and the multiple isotopes of Gd, each of which produces a somewhat different gamma cascade. The Gd analysis requires the delayed signal to satisfy 4 MeV $< E_{vis} <$ 10 MeV. Keeping these bounds loose reduces systematic uncertainties from inexact modeling of the neutron capture process.

Ideal prompt signals have an energy between 1.02 MeV and about 10 MeV, as demonstrated in Sec. 6.2, but detector resolution can shift signals beyond those bounds, especially on the low energy

Table 13.1: Cuts applied to all signals before the Gd-based IBD selection is performed. Symbols are defined in the text unless otherwise noted.

Cut name	Condition for passage	Note
Trigger threshold	$E_{vis} > 0.4$	
Muon rejection	$E_{vis} \leq 20$ MeV	
	$E_{IV} \leq 16$ MeV	
After-muon rejection	$\Delta T_{\mu} \geq 1$ ms	ΔT_{μ} is the time since the most recent muon, defined by above criteria.
Light noise rejection	$q_{max}/q_{tot} \leq 0.12$	q_{max} (q_{tot}) is the maximum (total) PMT charge.
	$\sigma_t \leq 36$ ns	σ_t is the standard deviation of the PMT hit times.
	$\sigma_q \leq (464 - 8\ \text{ns}^{-1}\sigma_t)$ CU	σ_q is the standard deviation of the PMT charges. CU is the unit of charge defined in the charge integration algorithm.
	$1/N \sum_i^N (q_{max} - q_i)^2/q_i$ $\leq 3 \times 10^4$ CU	N is the number of PMTs within a 1-m radius sphere centered at the PMT with q_{max}.

side. In addition, collecting events with prompt energy below 1 MeV and above 10 MeV provides more information about background populations. Constraining backgrounds through their prompt spectra improves $\sin^2 2\theta_{13}$ precision, in the manner described in Ch. 15. The prompt energy range is chosen as 0.5 MeV $< E_{vis} <$ 20 MeV, with the lower bound set to match 100% trigger efficiency, with negligible uncertainty, and the upper bound complementary to the muon definition.

13.1.3 Background vetoes

A final set of cuts targets fast neutrons, stopping muons, and decays of the cosmogenic isotopes ^{9}Li and ^{8}He. Many events of these types are rejected by the simple muon identification described above. In other cases, the progenitor muon may escape detection by barely grazing the active detector volumes, traveling only a short distance before decaying, or passing entirely outside the detector but generating fast neutrons which travel inside. These muons are frequently identified by the OV. Demanding that prompt signals have no coincident OV activity eliminates about 60% of fast neutron and stopping muon (FN+SM) backgrounds. The OV-based veto has varying efficiency across the Gd dataset because the OV was installed in stages. The lower OV became operational about two months after data-taking began, and the upper OV was installed about one year later.

The IV is also used to reject FN+SM events. Most muons deposit energy in the IV before entering the ID, and some fast neutrons cause proton recoils in the IV. The criteria for the IV-based cut include the number of IV PMTs hit; the total IV charge, Q_{IV}; the distance between the IV and ID vertices, ΔR_{IV-ID}; and the amount of time by which the IV event preceded the ID event, ΔT_{IV-ID}.

The OV and IV cuts are unable to tag muons which enter the detector through the chimney. If such muons stop before depositing $E_{vis} > 20$ MeV, their decays can remain as backgrounds. These events typically have poor delayed vertex reconstructions, because the tracks of Michel electrons and positrons in the chimney, and associated reflections, do not match the reconstruction model of a point-like light source inside the main detector volume. A cut based on $\mathcal{L}_{vtx}$, defined in Eq. 11.2, eliminates most of these events. The $\mathcal{L}_{vtx}$ cut also removes light noise.

Finally, about 55% of beta-neutron decays of ^{9}Li and ^{8}He are rejected with a likelihood-based cut. The likelihood function, $\mathcal{L}_{Li}$, is constructed from two parameters: the distance between the event and a preceding muon track, and the number of neutron captures which occur close in time. These

parameters effectively identify cosmogenic isotopes, which tend to be created near muon tracks and are often accompanied by multiple spallation neutrons. The full definition and motivation for $\mathcal{L}_{Li}$ is given in [105]. These four background vetoes have very small signal inefficiencies, as calculated from data-based studies and listed in Tab. 13.3.

Table 13.2 summarizes the IBD selection cuts for the Gd channel.

13.1.4 Observed IBD candidates

Applying all selection cuts to the Gd dataset yields 17351 IBD candidates in the reactor(s)-on period and seven candidates in the reactors-off period. Figure 13.2 plots the rate of candidates as a function of time, superimposed upon the predicted IBD event rate. In this plot, the MC prediction does not include backgrounds or oscillation, both up to $\mathcal{O}(10\%)$ effects (in opposite directions). Still, observations clearly track changes in event rate predicted by the MC. Three distinct rate levels correspond to different reactor operating conditions, with the highest set occurring when both reactors are operating near full power, the middle set occurring when only one reactor is operating, and the small set near zero occurring during periods when both reactors are shut down.

13.2 Signal detection efficiency

Understanding the IBD detection efficiency in data, relative to MC, is essential for a single-detector θ_{13} measurement. Uncertainty on the data/MC efficiency ratio contributes directly to signal normalization uncertainty, in the same manner as reactor flux uncertainty. That effect motivates detailed studies of each contribution to the total detection efficiency: the IBD selection efficiency; IBD inefficiencies of the background vetoes; the number of free protons N_p in the NT; the fraction of IBD events in the NT followed by Gd capture; and the effect of events with either a prompt or delayed signal outside the NT, known as the spill-in/out current. The following sections outline the estimation of $C_{data/MC}$, the ratio of detection efficiency in data to the same quantity in MC, for each of these factors. Methods for quantifying the uncertainty on $C_{data/MC}$ are also highlighted.

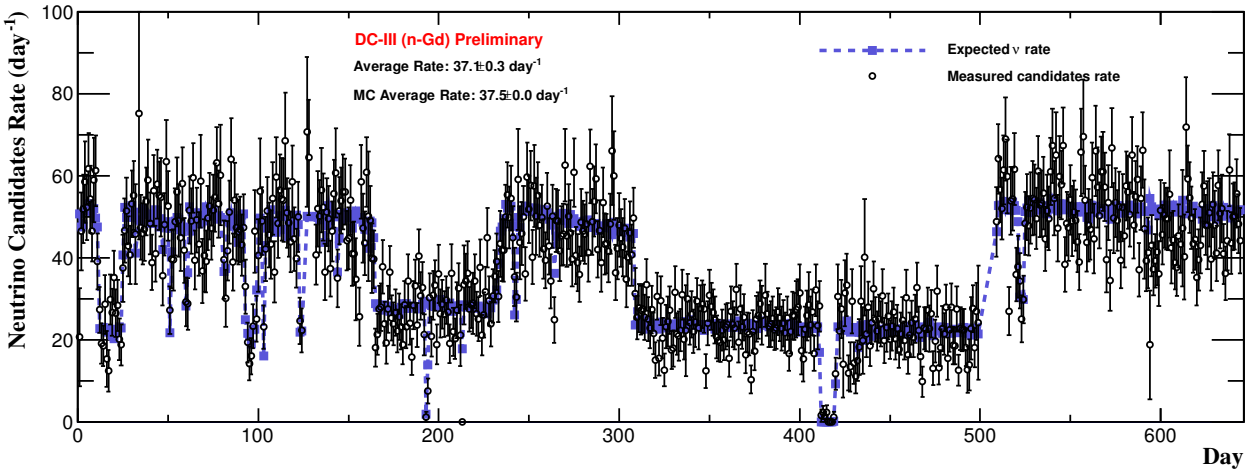

Figure 13.2: Daily rate of IBD candidates from the Gd selection (black open circles with statistical error bars) superimposed upon the expected IBD event rate (blue filled squares), as a function of days since the start of far detector data-taking in April 2011. The three distinct rate levels correspond to both reactors operating, one reactor operating, and no reactors operating. In this plot, the MC prediction includes neither oscillation effects nor backgrounds. Figure from [106].

Table 13.2: Selection cuts for IBD candidates in the Gd channel, including the energy, time correlation, and vertex correlation cuts that define a basic IBD candidate, plus four targeted background vetoes. Symbols are defined in the text.

Cut name	Condition for passage
IBD selection	Prompt signal: 0.5 MeV $< E_{vis} <$ 20 MeV
	Delayed signal: 4 MeV $< E_{vis} <$ 10 MeV
	0.5 µs $< \Delta T <$ 150 µs
	$\Delta R <$ 100 cm
Multiplicity cut	No signals in 200 µs before prompt; no signals other than delayed signal in 0.5–600 µs after prompt
OV-based veto	Prompt signal not coincident within 224 ns with OV activity
IV-based veto	Prompt signal has *at least one* of the following properties:
	IV PMT multiplicity < 2
	$Q_{IV} \leq$ 400 CU
	$\Delta R_{IV-ID} \geq$ 3.7 m
	$\Delta T_{IV-ID} <$ 10 ns, or $\Delta T_{IV-ID} >$ 100 ns
Vertex quality veto	Delayed signal: $E_{vis}/\text{MeV} > 0.068\ \text{e}^{-\ln \mathcal{L}_{vtx}}$
^{9}Li/^{8}He likelihood veto	$\mathcal{L}_{Li} < 0.4$

13.2.1 IBD selection efficiency

The joint efficiency of the IBD selection cuts (prompt E_{vis}, delayed E_{vis}, ΔT, and ΔR) is evaluated in two ways. The first method uses calibration data from ^{252}Cf deployments along the z-axis system and corresponding MC simulations. The efficiency of the prompt E_{vis} cut is found to be 100% with negligible uncertainty. From the remaining three cuts, the IBD selection efficiency, ϵ_{IBD}, is defined as:

$$\epsilon_{IBD} = \frac{\text{Events passing delayed } E_{vis}, \Delta T, \text{ and } \Delta R \text{ cuts}}{\text{Events passing loose cuts}} \tag{13.1}$$

where the loose cuts are: 3.5 MeV $< E_{vis} <$ 10 MeV for delayed energy, 0.25 µs $< \Delta T <$ 1000 µs, and no constraint on ΔR. While the selection efficiency could instead be defined as the product of delayed E_{vis}, ΔT, and ΔR efficiencies, Eq. 13.1 has the advantage of accounting for correlations. Evaluating this expression for the deployment at the center of the detector produces ϵ^0_{IBD}. The ϵ_{IBD} for all other locations in the NT is then computed as:

$$\epsilon_{IBD}(\rho, z) = f_\rho(\rho)\ f_z(z)\ \epsilon^0_{IBD} \tag{13.2}$$

where f_ρ and f_z characterize the spatial dependence of the efficiency, which tends to decrease towards the NT perimeter. These functions are extracted from calibrations along the z-axis, with the MC-supported assumption that efficiency depends simply on distance to an NT wall, regardless of whether it is the top, bottom, or side. Data yields $\langle\epsilon_{IBD}\rangle = 98.29 \pm 0.23\%$, while MC yields $\langle\epsilon_{IBD}\rangle = 98.26 \pm 0.22\%$, resulting in $C^{sel}_{data/MC} = 1.0003 \pm 0.0032$.

The ^{252}Cf result is confirmed by the second method, which uses IBD candidates. In this case, ϵ_{IBD} is determined from IBD events spread across the entire NT volume. The resulting $C^{sel}_{data/MC}$ is fully consistent with the ^{252}Cf result.

13.2.2 Proton number

Equation 12.6 shows that the predicted IBD rate is directly proportional to N_p. It is therefore convenient to group data-MC uncertainties for N_p with efficiency uncertainties. For a Gd-based analysis, N_p is the number of free protons (^{1}H nuclei) in the NT scintillator. The central value of N_p is determined from the weight and chemical composition of this liquid. A sensitive weighing tank system was deployed to measure the NT scintillator during filling. Measurements were corrected for temperature differences between the weighing setup and the actual detector environment. The

uncertainty on the weight measurement is 0.04%. During data-taking, temperature variations in the NT do not exceed 1°. These changes have a negligible impact on the liquid density and thus on the mass of liquid inside the NT, versus in the attached expansion tank.

From the well-known composition of the NT liquid, the ^{1}H fraction is determined to be approximately 13.60% by weight. Uncertainty on the chemical composition is dominated by imprecise knowledge of the n-dodecane/o-PXE ratio, amounting to about 10 kg out of the 18000 kg of scintillator prepared for the NT of both the near and far detectors. To assess systematic uncertainties on the ^{1}H fraction, calculations were performed with the relative n-dodecane/o-PXE amounts shifted by 10 kg in both directions. Impurities in the scintillator chemicals were also considered. The overall uncertainty on the ^{1}H fraction is taken to be 0.3% [70]. The weight uncertainty is negligible compared to this number, so the total uncertainty on N_p in the NT is 0.30%.

13.2.3 Gd capture fraction

The high neutron cross section of Gd ensures that most IBD interactions in the NT are followed by a Gd capture. The remaining neutrons capture mainly on H, with a very small fraction capturing on C, Fe, or another nucleus. To accurately predict the number of signal events in the Gd channel, the MC must correctly model these capture probabilities. The Gd capture fraction is measured with ^{252}Cf deployments at the center of the detector. Gamma rays released in the spontaneous fissions of ^{252}Cf serve as a prompt tag of neutron emission. The fraction of neutron captures on Gd is defined as:

$$f_{Gd} = \frac{\text{Number of events with 3.5 MeV} < \text{delayed } E_{vis} < \text{10 MeV}}{\text{Number of events with 0.5 MeV} < \text{delayed } E_{vis} < \text{10 MeV}} \tag{13.3}$$

The threshold of 3.5 MeV removes nearly all H captures but permits the acceptance of C captures. To be counted in this fraction, events must pass a basic set of background-removal cuts. Backgrounds from accidental coincidences are measured with off-time windows, as in Sec. 13.3.1, and subtracted. Variations on the energy windows are considered to determine the systematic uncertainty on f_{Gd}. Analysis of data produces $f_{Gd} = 85.30 \pm 0.08\%$, while analysis of MC produces $f_{Gd} = 87.49 \pm 0.04\%$. The ratio of these fractions gives $C^{Gd}_{data/MC} = 0.9750 \pm 0.0011(\text{stat}) \pm 0.0041(\text{syst})$.

13.2.4 Spill-in and spill-out

In an ideal Gd-channel IBD event, both the prompt and delayed signals occur inside the NT volume. Most selected events match this description, but a small number straddle the NT-GC boundary. Cases where the positron signal appears in the NT, while the neutron captures on H in the GC, constitute a spill-out current. Cases in which the IBD interaction and positron signal occur in the GC, while the neutron enters the NT and captures on Gd, constitute a spill-in current. Naively, these currents might be expected to cancel one another in signal rate computations. However, the detector geometry and longer neutron transit length in the GC lead to an excess of the spill-in current. This net spill current increases the expected number of signal events by a few percent and must be modeled in the MC.

In the nominal MC, the spill current represents 2.08% of the Gd-channel IBD events. This percentage is sensitive to details of the neutron thermalization model described in Sec. 12.2.1. Because the spill current cannot be measured readily in data, $C^{spill}_{data/MC}$ is taken as 1.0, and a conservative uncertainty is assigned. To estimate this uncertainty, the nominal spill current is compared to the spill current calculated with alternative neutron thermalization model. The second model, TRIPOLI-4, is more empirically based and a somewhat better match to neutron thermalization observables in Double Chooz [107]. The TRIPOLI-4 model predicts a spill current equal to 2.36% of the total Gd-channel IBD events, leading to a $C^{spill}_{data/MC}$ uncertainty of 0.27%.

13.2.5 Overall MC efficiency correction factor

The overall data-to-MC efficiency ratio is the product of all $C_{data/MC}$ factors listed in Secs. 13.2.1–13.2.4. Table 13.3 summarizes this $C^{tot}_{data/MC}$ calculation. In oscillation fits, $C^{tot}_{data/MC}$ is applied to signal MC as a kind of second-order normalization tuning. This step brings the MC into closer agreement with the data, correcting for small discrepancies which are difficult to tune in the original MC production. Uncertainties on the individual $C_{data/MC}$ are uncorrelated, so the total uncertainty is easily calculable. In the Gd analysis, the uncertainty on $C^{tot}_{data/MC}$ is 0.63%. The factor $C^{tot}_{data/MC}$ and its uncertainty have no prompt energy dependence, reflecting the nearly complete independence of the positron energy and neutron behavior following IBD interactions.

Table 13.3: Data-to-MC ratios of efficiencies and other normalization factors in the Gd analysis, denoted $C_{data/MC}$, and their relative uncertainties. The overall factor $C^{tot}_{data/MC}$, the product of all individual factors, is listed in the final row. This factor is used to correct the MC normalization in oscillation fits. The total uncertainty is the quadratic sum of all contributions. Uncertainties listed as $< 0.1\%$ have been determined to negligibly affect the overall MC normalization uncertainty.

Source	$C_{data/MC}$	Relative uncertainty on $C_{data/MC}$
Electronics, trigger, and DAQ inefficiency	1.0000	$< 0.1\%$
Veto for 1 ms after muon	0.9551	$< 0.1\%$
Light noise rejection	0.9999	$< 0.1\%$
OV-based veto	0.9994	$< 0.1\%$
IV-based veto	0.9996	0.01%
Vertex quality veto	0.9994	0.11%
^{9}Li/^{8}He likelihood veto	0.9946	0.02%
IBD selection efficiency	1.0000	0.19%
Multiplicity cut	0.9894	$< 0.1\%$
NT proton number	1.0000	0.30%
Gd fraction	0.9750	0.43%
Spill current	1.0000	0.27%
Overall	0.9149	0.63%

13.3 Background estimates

After all Gd-channel selection cuts are applied, backgrounds comprise about 4% of candidates. Multiple techniques are used to estimate the rate and energy spectra of the residual background populations. Strong constraints on these quantities increase the precision of the $\sin^2 2\theta_{13}$ fits. Although previous Double Chooz analyses used MC-based methods to predict some backgrounds, the current analysis relies solely on data-driven studies.

13.3.1 Accidental coincidences

As noted in the previous section, accidental coincidences are very effectively suppressed by the requirements of high delayed E_{vis} and short ΔT. In the remaining accidental coincidences, prompt signal are mainly gamma rays from the decays of ^{40}K (1.46 MeV), ^{208}Tl (2.62 MeV, from the ^{232}Th series), or other primordial radionuclides in or just outside the detector. Delayed signals are a mixture of spallation neutron captures on Gd, electrons from beta decays of cosmogenic ^{12}B, and neutron captures on C, Fe, and other nuclei.

The accidental rate and spectrum can be precisely measured from data because the random prompt and delayed signals are, by definition, uncorrelated. Up to small corrections, the same rate and spectrum will occur in any pair sufficiently separated and muon-isolated windows covering 0.5 MeV $< E_{vis} <$ 20 MeV and 4 MeV $< E_{vis} <$ 10 MeV. In the main accidentals analysis, prompt signals are selected according to the ordinary IBD criteria. Then, delayed signals are selected from off-time windows, intervals of length 149.5 µs opened at least 1 s after the prompt signal. Many consecutive off-time windows are opened to increase the sample statistics. This selection produces accidental rate and spectrum measurements with very small statistical uncertainties. The muon veto, $\mathcal{L}_{Li}$ veto, and multiplicity cut have slightly different effects on the rate of off-time candidates than the rate of normal candidates. These differences are corrected, a process that adds a small systematic uncertainty to the accidental rate. The final accidental rate is 0.0701 ± 0.0003 (stat) ± 0.0026 (syst) events per day. Figure 13.3 shows the prompt spectrum of accidental events.

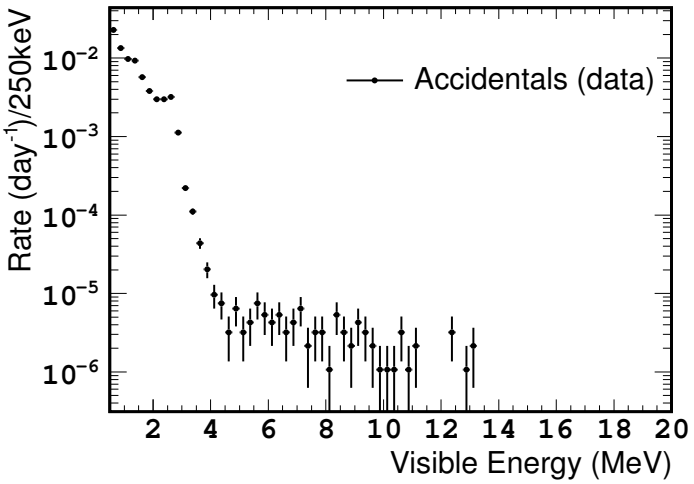

Figure 13.3: The prompt spectrum of Gd-channel accidental coincidences, as measured in an off-time selection. Figure from [86].

13.3.2 Fast neutrons and stopping muons

Two classes of cosmogenic events, fast neutrons and stopping muons, are studied together because they share similar prompt spectra and OV/IV activity. Fast neutrons, with energies of a few keV or more, are created by muon spallation on nuclei outside the detector. If these neutrons enter the detector and scatter off protons, the proton recoils can satisfy the criteria for a prompt IBD signal. The spectrum of proton recoils is expected to be flat in the range of 0.5 MeV $< E_{vis} <$ 20 MeV, with the possibility of slight energy dependence from detector geometry or quenching. Fast neutrons can satisfy the delayed signal criteria by thermalizing and capturing on Gd. Prompt and delayed signals may come from different neutrons, usually generated in the same spallation event. Because these signals are separated by the neutron thermalization process, the fast neutron ΔT distribution is not distinguishable from that of real IBD events.

As previously noted, the ionization tracks of chimney-entering muons can meet the prompt signal requirement for IBD events. Within the prompt E_{vis} window, these signals are approximately uniformly distributed. Aside from a small fraction of μ^- which capture on nuclei, most stopped muons decay. The extremely dominant decay mode is $\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu$ or its charge conjugate, $\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$. Since the Michel spectrum extends to about 50 MeV, a subset of the decay

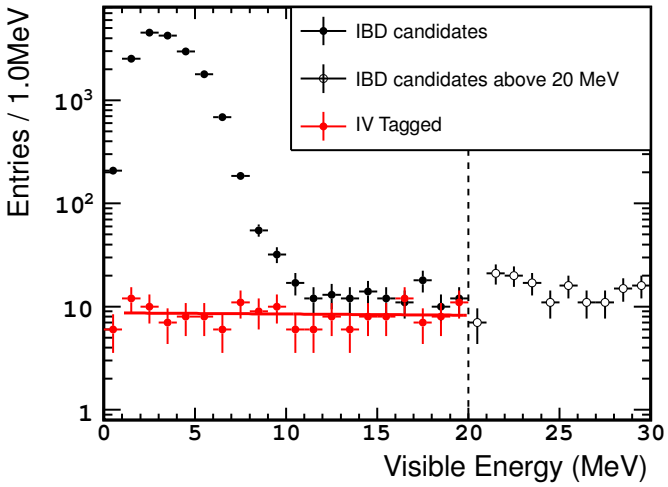

Figure 13.4: The spectrum of IV-tagged candidates, expected to be FN+SM (filled red circles); IBD candidates (filled black circles); and candidates from an IBD selection with the prompt energy extended to 30 MeV (open black circles). The red line is a zeroth order polynomial fit to the IV-tagged candidates. The rate of IV-tagged candidates is somewhat lower than the rate of IBD candidates above ~ 10 MeV because not all FN+SM are tagged with the IV; the important aspect of the fit is its shape, which is consistent between the IV-tagged and IBD samples. Figure from [86].

electrons and positrons satisfy the delayed signal criteria. The characteristic ΔT of stopping muon events is the muon lifetime, $\tau_\mu = 2.20\,\mu s$. This short time constant allows stopping muons to be isolated in some studies. The same principle was applied in the previous Gd-based selection, where the ΔT threshold was $2\,\mu s$. That cut occurred in a ΔT interval which is very sensitive to the neutron thermalization model of the MC, increasing uncertainty on the signal normalization. In the current analysis, where stopping muons are largely removed by the cuts of Sec. 13.1.3, discarding slightly more background is less valuable than controlling signal uncertainties.

The IV-based veto was designed to remove FN+SM events from the IBD selection. Inverting that cut, while maintaining the other IBD cuts, produces an IV-tagged sample of FN+SM. The spectrum of these backgrounds is deduced from the IV-tagged sample. Figure 13.4 shows the energy spectrum of IV-tagged candidates, fit with a horizontal line. An alternative fit that allows a slope, m, to take on nonzero values yields $m = -0.02 \pm 0.11$ events/MeV^2, in good agreement with zero.

Figure 13.4 also shows candidates chosen by an extended version of the IBD selection, in which the prompt energy window ends at 30 MeV and the muon definition is modified accordingly. The candidates between 20 and 30 MeV are assumed to be FN+SM, as no other background spectra extend that high in energy. With the assumption of a flat shape, the 20–30 MeV sample is used to estimate the rate of FN+SM among normal IBD candidates. For the portion of the dataset in which at least one reactor was operating, this rate comes to 0.604 ± 0.051 events/day. This rate averages over various OV coverage levels available during that period. For the reactor-off period, which had full coverage from the lower OV, the FN+SM rate is expected to be 0.529 ± 089. The rates and spectrum shape have been verified with an independent set of OV-tagged candidates.

13.3.3 Cosmogenic ^{9}Li and ^{8}He decays

The largest background among Gd-channel candidates comes from the decays of ^{9}Li and ^{8}He, both created by cosmic muons or their products interacting with C nuclei in the detector. Several other unstable isotopes are produced in the same type of interactions, but the majority undergo ordinary beta decays and make only a minor contribution to the accidental background. Decays of ^{9}Li (^{8}He) are unique because approximately 50% (16%) involve emission of a neutron as well as an electron [108]. Figure 13.5 shows the decay branches for the dominant isotope, ^{9}Li.

The Q value of beta-neutron decays of ^{9}Li (^{8}He) is 11.9 MeV (8.6 MeV), so the decay electrons can easily meet the prompt signal requirements [109, 110]. Gd captures of the decay neutrons can create delayed signals. The mean lifetimes of these isotopes, 257.2 ms for ^{9}Li and 171.8 ms for ^{8}He [108], make them impractical to target with a simple after-muon veto. Given the through-going muon rate of about 40 Hz, a veto length of even 25 ms would leave no detector livetime. The similarity of the ^{9}Li and ^{8}He endpoints and lifetimes leads to their treatment as a single background class, Li+He. Observations in the KamLAND [111] and Borexino detectors [112] imply that ^{9}Li decays constitute the majority of this class, with ^{8}He contributing up to $\mathcal{O}(10\%)$.

The spectrum of Li+He events is measured by inverting the likelihood-based cut introduced in Sec. 13.1.3 and subtracting the accidental background. To enhance statistics, candidates from a basic H-channel selection are added to these $\mathcal{L}_{Li}$-tagged events. The H selection replaces the delayed E_{vis} and ΔT requirements of the Gd selection with $1.8 < E_{vis} < 2.6$ MeV and $0.5\,\mu s < \Delta T < 600\,\mu s$. As in IBD events, prompt Li+He signals are independent of the delayed neutron behavior, so the

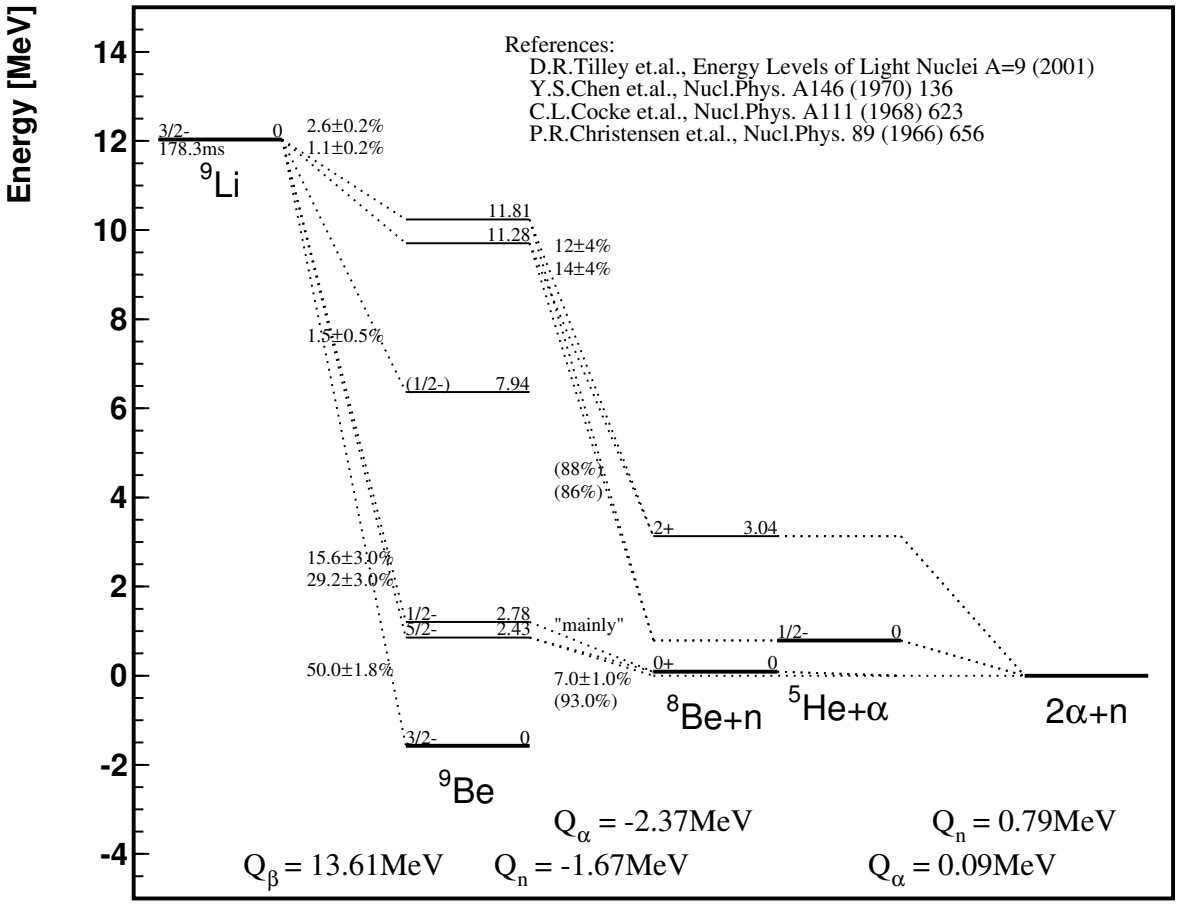

Figure 13.5: The decay schemes of ^{9}Li. Horizontal line segments show energy levels. Each level is labeled with its total angular momentum and parity, in the format J^P, along with its energy relative to the ground state, in MeV. Branching ratios and Q values are also shown. Numbers in parentheses have significant uncertainties. Figure from [113] (used with permission).

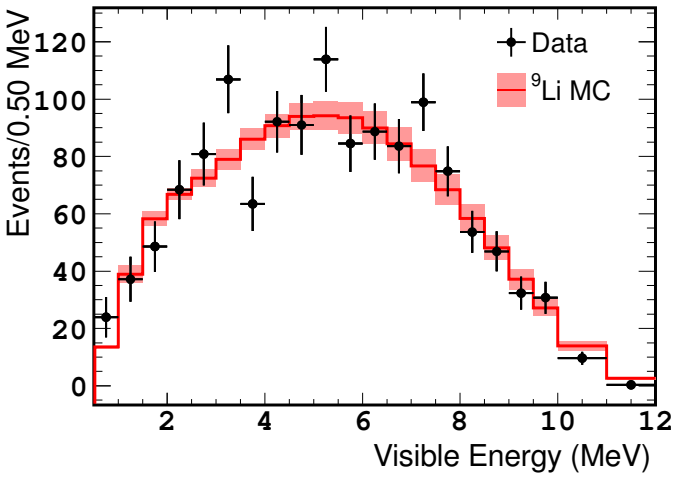

Figure 13.6: The prompt energy spectrum of ^{9}Li and ^{8}He candidates selected from the Gd-channel IBD cuts, except with the $\mathcal{L}_{Li}$ cut inverted, plus a similar H-channel selection (black points with statistical error bars). The red histogram shows an MC prediction for the Li+He spectrum, with the shaded area indicating systematic uncertainty in each bin; these uncertainties are highly correlated between bins. For this plot, the normalization of the MC was determined by a χ^2 minimization with respect to the data. Figure from [86].

Gd and H selections sample the same prompt spectrum, up to small differences between the NT and GC energy scales. These differences are not significant compared to the statistical uncertainty on the combined spectrum. Figure 13.6 shows this data-derived spectrum, which agrees well with the MC simulation shown in the same figure. Many of the ^{9}Li and ^{8}He branching ratios are not well known, leading to significant systematic uncertainty on the MC spectrum. Sensitivity studies indicate that the data-derived spectrum and MC spectrum, with their respective uncertainties, allow comparable $\sin^2 2\theta_{13}$ precision. The data-derived spectrum is chosen to limit model dependence of the oscillation analyses.

The Li+He rate is first derived for an IBD selection without the $\mathcal{L}_{Li}$ cut. The efficiency of that cut is then calculated and applied. The non-vetoed rate is derived from the distribution of $\Delta T_{\mu-p} = (T_{prompt} - T_{\mu})$ for the IBD candidates. Here, T_{μ} is the time of every preceding muon. In principle, an exponential fit using the isotope lifetimes could reveal the level of Li+He contamination. Realistically, the signal from correlated muon–Li/He pairs is overpowered by the much larger set

of uncorrelated pairs. A fit to $\Delta T_{\mu\text{–}p}$ is only informative if some extra cut is applied to enhance Li+He purity, such as a requirement that the progenitor muon deposit a large amount of energy in the detector. Alternatively, a cut can be made on $\Delta R_{\mu\text{–}p}$, the distance between the prompt event and the track of a preceding muon.

Both of these approaches are used to estimate the Li+He rate. One group of IBD events is selected based on the high energy deposited by a preceding muon. Other groups are selected by requiring $\Delta R_{\mu\text{–}p} < 75$ cm. For the latter groups, spatial cut efficiencies are estimated from the measured Li+He production profile and detector acceptance. In all groups, the Li+He rate is estimated from a fit to $\Delta T_{\mu\text{–}p}$. Results are combined to yield an overall rate estimate.

A more stringent lower bound can be derived from the $\Delta T_{\mu\text{–}p}$ profile of an especially high-purity Li+He sample. The main selection criterion for this sample is a muon energy deposition of several hundred MeV. Combining these two approaches yields a Li+He rate estimate of $2.06^{+0.37}_{-0.14}$ per day. Applying the $\mathcal{L}_{Li}$ veto efficiency, derived from a $\Delta T_{\mu\text{–}p}$ fit to vetoed candidates, returns $0.97^{+0.41}_{-0.16}$ events per day.

13.4 Summary of candidates

Table 13.4 lists the number of observed and expected IBD candidates in the Gd-channel analysis. In the reactor(s)-on period, a deficit of approximately 5% is observed with respect to the no-oscillation prediction. Chapters 15 and 16 describe how this deficit is analyzed in the context of $\sin^2 2\theta_{13}$-driven oscillation. A deficit is also observed in the reactors-off period. This difference is too large to be fully explained by oscillation of the antineutrino signal and most likely reflects a statistical fluctuation and/or an overestimate of the backgrounds. Considering statistical and systematic uncertainties, the probability of compatibility between the reactor-off data and expectation is roughly 10%, corresponding to agreement within 2σ.

Table 13.4: The number of observed and expected events in the reactor(s)-on and reactors-off periods of the Gd analysis. The correction factor $C^{tot}_{data/MC}$ and the nominal energy scale corrections have been applied to the signal MC. No oscillation has been applied. Systematic uncertainties, and treatments of statistical uncertainties, are described in the text.

Event type	Events in reactor(s)-on period	Events in reactor-off period
No-oscillation expectation		
Antineutrino signal	17534	1.6
Accidentals	32	0.5
FN+SM	278	3.8
Li+He	447	7.0
Total	18291	12.9
Observation		
Total	17351	7

Chapter 14

H-channel candidates

Selection of H-channel IBD candidates involves several techniques developed for the Gd-based analysis. It also includes a new, more advanced approach to signal identification and a new background veto. Through these advances, the H channel yields about twice the signal of the Gd channel, with only modestly worse background contamination and systematic uncertainties. This chapter outlines the signal selection and background analyses for the H channel, focusing on features unique to this H-based analysis.

14.1 Signal selection

14.1.1 Pre-selection

Nearly the same pre-selection cuts are used in the H analysis as the Gd analysis. The exception is the after-muon veto length, which is increased to 1.25 ms to control the higher cosmogenic background rate in this channel. Table 14.1 summarizes the H-channel pre-selection cuts.

14.1.2 IBD selection

The previous H-based analysis in Double Chooz relied on one-dimensional cuts in prompt and delayed E_{vis}, ΔT, and ΔR, much like the current Gd-based analysis. This selection admitted a huge amount of accidental backgrounds, resulting a signal-to-background ratio close to 1:1. Sensitivity studies showed that reducing the accidental background rate would greatly improve $\sin^2 2\theta_{13}$ precision. The delayed E_{vis}, ΔT, and ΔR distributions of signal MC and the off-time accidental

selection are markedly different, suggesting that much better separation should be possible with more sophisticated cuts. Correlations between pairs of parameters are also different for the two event types.

The current H analysis exploits these differences through a multivariate selection algorithm. The chosen strategy is an artificial neural network (ANN), implemented through the multi-layer perceptron (MLP) algorithm with backpropagation. The TMVA package in ROOT [114] provides the computational structure. As in a typical ANN, information about each IBD candidate is passed through layers of neurons, or nodes, which are connected by adaptive weights.

In the H-channel ANN, four input nodes (delayed E_{vis}, ΔT, ΔR, and one bias node), one nine-node hidden layer, and one output node provide optimal discrimination power and efficiency. For the input and output nodes, the activation functions are linear; for the hidden layer, they are hyperbolic tangents. The output node produces a continuous variable ranging from -1.2 (highly accidental-like) to 1.2 (highly signal-like). Specially collected samples of signal MC and off-time accidentals are used to train the network. Loose cuts on delayed E_{vis}, ΔT, and ΔR, along with a multiplicity cut, are applied before the ANN algorithm. Requiring the ANN output to exceed -0.23 yields a good compromise between signal efficiency and background rejection, as illustrated in Fig. 14.1. In the ANN–based selection, the signal-to-background ratio is more than $10\times$ that of the previous H analysis.

The prompt E_{vis} threshold of the H analysis is set at 1 MeV, rather than 0.5 MeV as in the Gd analysis. This higher threshold eliminates a small class of IBD interactions occurring in the buffer. In these events, the prompt signal consists of a single 0.511 MeV gamma ray which enters the GC. Some discrepancy has been observed between the data and MC model for these events, so they are simply removed from this analysis.

All IBD selection cuts for the H channel are summarized in Tab. 14.2

14.1.3 Background vetoes

The Li+He likelihood veto and OV-based veto are applied in the H selection exactly as in the Gd selection. An IV-based veto, similar to the Gd-channel cut, is applied to both prompt and delayed signals. A vertex quality cut is applied to the delayed signal, with parameters optimized for the H selection.

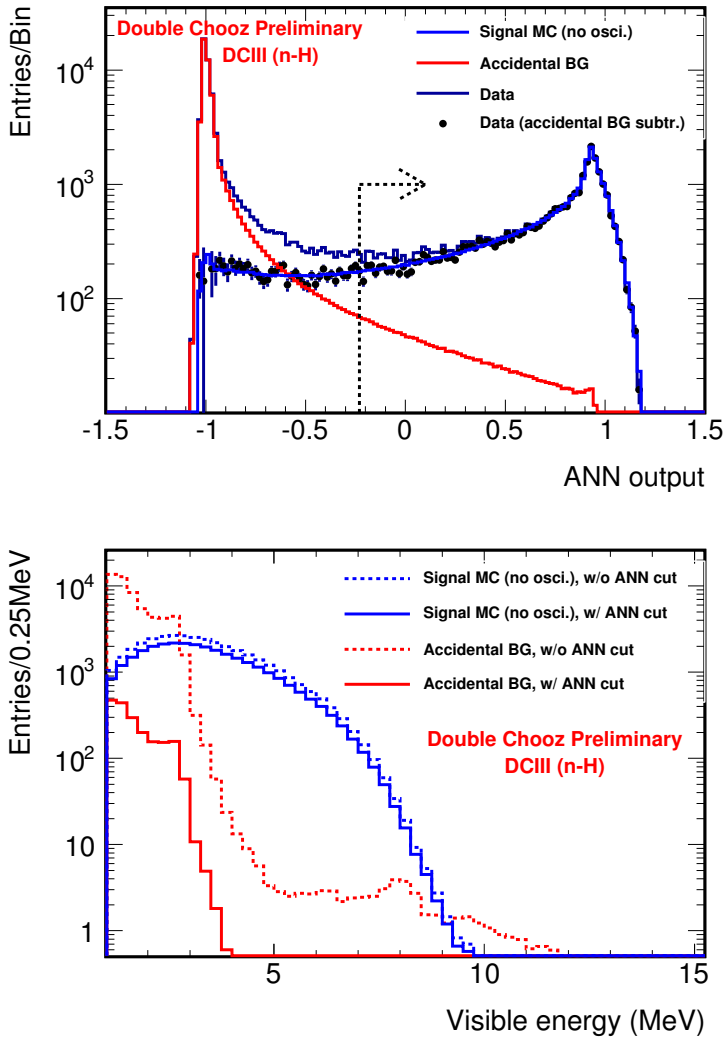

Figure 14.1: Top: The distribution of ANN outputs for an off-time accidental sample (red), signal MC (bright blue), IBD candidates in data (dark blue), and IBD candidates with accidentals subtracted (black points with statistical error bars). Bottom: The prompt energy spectrum of IBD candidates in data, without the ANN cut (dashed blue) and with the ANN cut (solid blue), and the off-time accidentals sample without the ANN cut (dashed red) and with the ANN cut (solid red). Figures from the forthcoming H publication.

Table 14.1: Cuts applied to all signals before the H-based IBD selection is performed. Pre-selection cuts for the Gd analysis are listed in Tab. 13.1.

Cut name	Condition for passage	Note
Trigger threshold	Same as Gd pre-selection	
Muon rejection	Same as Gd pre-selection	
After-muon rejection	$\Delta T_\mu \geq 1.25$ ms	ΔT_μ is the time since the most recent muon.
Light noise rejection	Same as Gd pre-selection	

One new background veto is included in the H analysis. This cut uses PMT waveform information to target fast neutrons. For each prompt signal, the start time of each PMT pulse is determined and corrected for the reconstructed time-of-flight of the signal. A histogram of these times represents the pulse shape of the event. In fast neutron events, the prompt pulse shape often includes multiple, lower-energy proton recoils before the main signal. As a result, the peak of the pulse shape may not occur at the beginning of the PMT pulse time distribution, as it typically does for positron events. Pulse peak positions are defined through a Gaussian fit around the highest bin. Generally, if more than 5 ns occurs between the first PMT pulse and the point located 1.8σ below the peak, the event is rejected as a probable fast neutron.

An extra condition must be included for events with prompt $E_{vis} < 3$ MeV. At these low energies, decays of the positron-electron bound state ortho-positronium (o-Ps) can have pulse shapes similar to fast neutrons. An o-Ps tag, also based on pulse shape information, has been developed to identify such events [115]. IBD candidates which receive that tag and have prompt E_{vis} between 1.2 and 3 MeV are not subjected to the pulse shape–based veto. Candidates with E_{vis} below 1.2 MeV, where o-Ps decay always leads to a pulse shape shift, are also exempted.

14.1.4 Observed IBD candidates

The H-channel selection yields 31835 candidates in the reactor(s)-on period and 63 in the reactors-off period.

Table 14.2: Selection cuts for IBD candidates in the H channel, including the energy, time correlation, and vertex correlation cuts that define a basic IBD candidate, plus five targeted background vetoes. The symbols μ_{PS} and σ_{PS} denote the mean and standard deviation of the highest peak in the PMT pulse shape. Other symbols are defined in the text.

Cut name	Condition for passage
IBD selection	Prompt signal: 1.0 MeV $< E_{vis} <$ 20 MeV
	Delayed signal: 1.3 MeV $< E_{vis} <$ 3 MeV
	0.5 µs $< \Delta T <$ 800 µs
	$\Delta R <$ 120 cm
	ANN output > -0.23
Multiplicity cut	No signals in 800 µs before prompt; no signals except delayed signal in 0.5–900 µs after prompt
OV-based veto	Prompt signal not coincident within 224 ns with OV activity
IV-based veto	Both signals have *at least one* of the following properties:
	IV PMT multiplicity < 2
	$Q_{IV} \leq$ 400 CU
	$\Delta R_{IV-ID} \geq$ 4 m
	$\Delta T_{IV-ID} <$ 20 ns, or $\Delta T_{IV-ID} >$ 100 ns
Vertex quality veto	Delayed signal: $E_{vis}/\text{MeV} > 0.2755\ \text{e}^{-\ln \mathcal{L}_{vtx}/2.0125}$
^{9}Li/^{8}He likelihood veto	$\mathcal{L}_{Li} < 0.4$
Pulse shape veto	Prompt signals with $E_{vis} \geq$ 3 MeV: $\mu_{PS} - 1.8\sigma_{PS} > 5$ ns
	Prompt signals with 1.2 MeV $< E_{vis} <$ 3 MeV, no o-Ps tag:
	$\mu_{PS} - 1.8\sigma_{PS} > 5$

14.2 Signal detection efficiency

As in the Gd analysis, MC correction factors $C_{data/MC}$ and associated uncertainties are evaluated. In the H analysis, some correction factors differ depending on the volume in which prompt and delayed signals occur. These differences are noted in the text below and applied to individual MC events accordingly. Table 14.3 lists the population-averaged corrections.

14.2.1 IBD selection efficiency

Methods analogous to those in Sec. 13.2.1 are used to find $C^{sel}_{data/MC} = 1.0000 \pm 0.0022$. In the H analysis, this number requires some extra consideration. In principle, the ANN cut could introduce prompt energy dependence to the selection efficiency. If that dependence differed between data and MC, the overall normalization adjustment performed by $C^{sel}_{data/MC}$ would be inadequate. The effect of the ANN cut was studied in both data and MC, and negligible prompt energy dependence was found.

14.2.2 Proton number

In the NT, $C^{N_p}_{data/MC}$ and its uncertainty are the same as those in the Gd analysis. Since the GC was not originally intended to be part of the fiducial volume, its scintillator was not weighed as precisely as the NT liquid. The N_p uncertainty in the GC is 1.04%, compared to 0.30% in the NT. Beyond the NT and GC liquids, a small number of IBD interactions occur inside the acrylic vessels. These events constitute less than 1% of candidates, and they receive no special treatment in the Gd analysis. In the H analysis, a special $C^{sel,TV}_{data/MC} = 1.390$ is applied to MC events with interaction vertices in the NT vessel and associated acrylic supports. This factor corrects for incomplete modeling of the support structures in the detector MC model.

14.2.3 H capture fraction

In the NT, the fraction of IBD events followed by neutron captures on hydrogen, f_H, can be taken as the complement of the Gd fraction described in Sec. 13.2.3. Comparing f_H in data and MC yields the correction factor $C^{H,NT}_{data/MC} = 1.1750 \pm 0.0078$ (stat) ± 0.0265 (syst). In the GC, f_H is derived from ^{252}Cf deployments in the portions of the guide tube farthest from the NT walls. A fraction

analogous to the definition in Sec. 13.2.3 gives f_H for data and MC, leading to the correction factor $C^{H,GC}_{data/MC} = 1.0020 \pm 0.0007$ (stat) ± 0.0003 (syst). An analysis involving IBD captures throughout the entire GC volume produces a consistent result.

14.2.4 Spill-in and spill-out

Spill currents in the H analysis include crossover between the GC and buffer volumes as well as between the GC and NT. As in the Gd analysis, the systematic uncertainty on MC spill modeling is evaluated through a TRIPOLI-4 comparison.

14.2.5 Overall MC efficiency correction factor

Table 14.3 summarizes the population-averaged MC correction factors for the H analysis.

14.3 Background estimates

14.3.1 Accidental coincidences

As in the Gd analysis, the accidental background rate and spectrum are measured through data from an off-time selection. Figure 14.2 shows the prompt spectrum of the off-time selection. The rate is 4.334 ± 0.007 (stat) ± 0.008 (syst) events per day in the reactor(s)-on period and 4.319 ± 0.056 (stat) ± 0.008 (syst) events per day in the reactors-off period.

14.3.2 Fast neutrons and stopping muons

The FN+SM rate and spectrum are measured through the same IV-tagging technique used in the Gd analysis. In this case, the extended IBD selection includes prompt E_{vis} up to 60 MeV. Detector acceptance effects give the H-channel FN+SM spectrum a different shape than the Gd-channel spectrum, as shown in Fig. 14.3. This shape is well described by the exponential form $y = F_0 e^{F_1 x} + F_2$, with $F_0 = 12.52 \pm 1.36$, $F_1 = 0.042 \pm 0.015$, and $F_2 = 0.79 \pm 1.39$. The estimated FN+SM rate is 1.55 ± 0.15 events per day in the reactor(s)-on period and 1.45 ± 0.20 events per day in the reactors-off period.

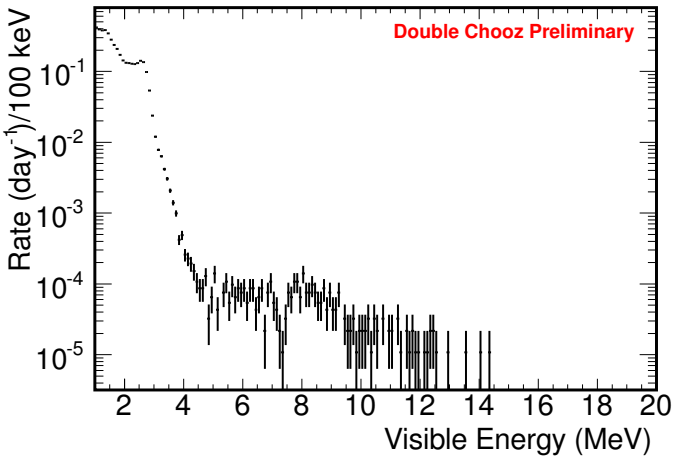

Figure 14.2: The prompt spectrum of H-channel accidental coincidences, as measured in an off-time window selection. Figure from the forthcoming H publication.

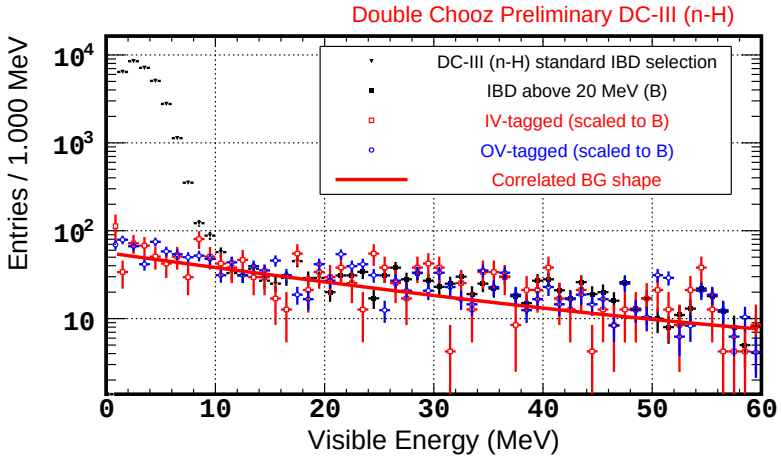

Figure 14.3: The spectrum of IV-tagged candidates (red squares) and OV-tagged candidates (blue squares), expected to be FN+SM; IBD candidates (black triangles); and candidates from an IBD selection with the prompt energy extended to 60 MeV (black squares). In this plot, the IV- and OV-tagged samples are normalized to match the IBD candidates above 20 MeV. The red line is a fit to the IV-tagged candidates of the form $y = F_0 e^{F_1 x} + F_2$. Figure from the forthcoming H publication.

Table 14.3: Data-to-MC ratios of efficiencies and other normalization factors in the H analysis, denoted $C_{data/MC}$, and their relative uncertainties. These factors are averages over events occurring in all detector volumes. The overall factor $C^{tot}_{data/MC}$, the product of all individual factors, is listed in the final row. This factor is used to correct the MC normalization in oscillation fits. The total uncertainty is the quadratic sum of all contributions. Uncertainties listed as $< 0.1\%$ have been determined to negligibly affect the overall MC normalization uncertainty.

Source	$C_{data/MC}$	Relative uncertainty on $C_{data/MC}$
Electronics, trigger, & DAQ inefficiency	1.0000	$< 0.1\%$
Veto for 1 ms after muon	0.9399	$< 0.1\%$
Light noise rejection	0.9994	$< 0.1\%$
OV-based veto	0.9994	$< 0.1\%$
IV-based veto	1.0000	0.16%
Vertex quality veto	0.9995	0.02%
^{9}Li/^{8}He likelihood veto	0.9949	0.01%
IBD selection efficiency	1.0000	0.22%
Multiplicity cut	0.9788	$< 0.1\%$
Proton number	1.0020	0.92%
H fraction	1.0150	0.21%
Spill current	1.0000	0.27%
Overall	0.9293	1.0%

14.3.3 Cosmogenic ^{9}Li and ^{8}He decays

The Li+He background is evaluated with the same techniques used in the Gd analysis. In the H channel, the estimated rate is $0.95^{+0.57}_{-0.33}$ per day. This estimate is very close to the expected Gd-channel Li+He rate of $0.97^{+0.41}_{-0.16}$, which is surprising, given the larger active volume of the H analysis. To first order, the rate of Li+He events should be proportional to the number of C nuclei in the corresponding volume, suggesting that the H-channel rate should be roughly twice the Gd-channel rate. However, detector geometry or efficiency-related effects may change this proportionality for the NT and GC. In any case, the large errors on both rates make them reasonably compatible with a range of expected ratios. The same Li+He spectrum is used for the H and Gd analyses, as justified in Sec. 13.3.3.

14.4 Summary of candidates

Table 14.4 summarizes the observed and expected IBD candidates in the H selection. The reactor(s)-on data shows a 4% deficit with respect to the no-oscillation prediction, comparable to the deficit observed in the Gd analysis. In contrast to the Gd analysis, the reactors-off period contains more observed candidates than expected. Once again, this tension is likely due to imperfect background rate estimates.

Table 14.4: The number of observed and expected events in the reactor(s)-on and reactors-off periods of the H analysis. The correction factor $C^{tot}_{data/MC}$ and the nominal energy scale corrections have been applied to the signal MC. No oscillation has been applied. Systematic uncertainties, and treatments of statistical uncertainties, are described in the text.

Event type	Events in reactor(s)-on period	Events in reactor-off period
No-oscillation expectation		
Antineutrino signal	30051	2
Accidentals	1974	31
FN+SM	706	10
Li+He	433	11
Total	33163	54
Observation		
Total	31835	63

Chapter 15

Rate+Shape fits

The final stage of the Gd and H analyses is the extraction of $\sin^2 2\theta_{13}$. Multiple strategies have been developed to accomplish this task. The most sensitive approach is the Rate+Shape fit, which infers $\sin^2 2\theta_{13}$ from the normalization and energy spectrum of the observed IBD candidates. Involvement of both rate and shape information demands careful consideration of systematic uncertainties. Treatment of uncertainties, and other features of the Rate+Shape fit, have evolved over successive oscillation analyses. This chapter begins with the statistical foundation and motivation of the Rate+Shape approach. Next, the prescription for each systematic uncertainty in the Gd and H analyses is described. Finally, results of the Gd and H Rate+Shape fits are presented, along with informative variations.

15.1 Statistical formalism

15.1.1 Maximum likelihood estimation

The goal of the oscillation fit is to estimate an unknown parameter, $\sin^2 2\theta_{13}$, from a set of observations and an expected distribution which depends on that parameter. The field of statistics offers several approaches to this general problem. Frequentist techniques dominate high energy physics, for both practical and philosophical reasons, although Bayesian approaches are sometimes used. Double Chooz has favored frequentist methods because they are standard in neutrino oscillation experiments and straightforward to implement.

In the frequentist approach, parameter fitting can be performed through maximum likelihood

estimation (MLE). The likelihood of a parameter θ, given an observed quantity x, is expressed as $\mathcal{L}(\theta; x) = \mathcal{L}(\theta) = f_\theta(x)$. Here, $f_\theta(x)$ is the continuous PDF for x, depending on θ. For a set of N independent measurements, $\mathbf{x} = (x_1, ..., x_N)$, each following the PDF $f_\theta(x)$, the joint likelihood is:

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{i=1}^{N} f_\theta(x_i) \tag{15.1}$$

where $\boldsymbol{\theta}$ is a vector of m parameters. In the MLE approach, the best estimator for these parameters, $\hat{\boldsymbol{\theta}}$, is the $\boldsymbol{\theta}$ that maximizes Eq. 15.1. In statistical terms, this estimator is consistent, efficient, and approximately unbiased.

For sufficiently large N, measurements can be sorted into a histogram with b bins. The MLE method can then be applied to the binned data, $\mathbf{n} = (n_1, ..., n_b)$, where n_i is the number of measurements falling into the i^{th} bin. The expected values of the bin contents, which depend on $\boldsymbol{\theta}$, are denoted $\boldsymbol{\mu}(\boldsymbol{\theta})$. In this situation, the quantity to maximize is the likelihood ratio:

$$\lambda(\boldsymbol{\theta}) = \frac{\mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}))}{\mathcal{L}(\boldsymbol{n})} \tag{15.2}$$

Equivalently, the quantity $-2\ln\lambda(\boldsymbol{\theta})$ can be minimized. A binned MLE fit is convenient when the number of measurements is large and systematic uncertainties are complicated, as in the Double Chooz oscillation analyses.

In a special case of the binned MLE fit, the n_i are independent and distributed about means $\mu_i(\boldsymbol{\theta})$, with known variances σ_i^2. In this case, the MLE approach corresponds to the method of least squares. Minimizing $-2\ln\lambda(\boldsymbol{\theta})$ is equivalent to minimizing:

$$\chi^2(\boldsymbol{\theta}) = -2\ln\lambda(\boldsymbol{\theta}) + C = \sum_{i=1}^{N} \frac{(n_i - \mu_i(\boldsymbol{\theta}))^2}{\sigma_i^2} \tag{15.3}$$

where C is a constant describing the normalization of $\mathcal{L}$. If the μ_i undergo relatively small changes during the χ^2 minimization procedure, as in Double Chooz, this constant can be neglected. In a case where the n_i are not independent, Eq. 15.3 becomes:

$$\chi^2(\boldsymbol{\theta}) = (\mathbf{n} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \mathbf{M}^{-1} (\mathbf{n} - \boldsymbol{\mu}(\boldsymbol{\theta})) \tag{15.4}$$

where $\mathbf{M}$ is a $b \times b$ covariance matrix with elements $M_{ij} = \mathrm{cov}(n_i, n_j)$.

15.1.2 Covariance matrix construction

The covariance matrix $\mathbf{M}$ can describe both statistical and systematic uncertainties. If n_i is the event count in the i^{th} bin, it follows a Poisson distribution. For $n_i \gtrsim 10$, this pattern is well approximated by a Gaussian distribution with mean and variance equal to n_i. Such distributions are uncorrelated between bins. Consequently, a statistical uncertainty covariance matrix can be built, following the Neyman convention, as:

$$M_{ij}^{Neyman} = \delta_{ij}\sqrt{n_i n_j} \tag{15.5}$$

An alternative construction, following the convention of Pearson's χ^2 test, uses the expected bin contents:

$$M_{ij}^{Pearson} = \delta_{ij}\sqrt{\mu_i \mu_j} \tag{15.6}$$

For large n_i and an appropriate model for μ_i, the Neyman and Pearson treatments produce very similar results. In Rate+Shape fits, Double Chooz uses the Pearson convention. This choice is consistent with the general Rate+Shape fit philosophy, in which $\mathbf{n}$ is considered the exact outcome of fixed experiment and all uncertainties apply to $\boldsymbol{\mu}$.

For systematic uncertainties, $\mathbf{M}$ is constructed by propagating errors in the prediction model to the μ_i. If the model includes q parameters, denoted $\boldsymbol{\beta}$, and the μ_i are functions of $\boldsymbol{\beta}$, this propagation is approximately:

$$M_{ij}^{syst} \approx \sum_{k=0}^{q}\sum_{l=0}^{q} \frac{\partial \mu_i}{\partial \beta_k}\frac{\partial \mu_j}{\partial \beta_l}\bigg|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} V_{kl} \tag{15.7}$$

where $V_{kl} = \mathrm{cov}(\beta_k, \beta_l)$. Eq. 15.7 can be expressed more succinctly using the Jacobian matrix, $\mathbf{J}$, where $J_{ik} = \partial\mu_i/\partial\beta_k|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}$:

$$\mathbf{M}^{syst} = \mathbf{J}\mathbf{V}\mathbf{J}^T \tag{15.8}$$

This approximation uses the first term in a Taylor series expansion of $\boldsymbol{\mu}(\boldsymbol{\beta})$ around $\hat{\boldsymbol{\theta}}$ and is exact if that function is linear. If $\boldsymbol{\mu}(\boldsymbol{\beta})$ is not close to linear, a covariance matrix approximation may not be feasible. Pull terms, introduced in Sec. 15.1.6, offer an alternative treatment for systematic uncertainties.

The covariance matrix for a linear $\boldsymbol{\mu}(\boldsymbol{\beta})$ has a simple form. Suppose that $\mu_i = \kappa g_i(\beta_l \neq \kappa)$, such that κ is an overall normalization factor and the expression in parentheses indicates that g_i is

a function of the other β_l. If κ is uncorrelated to the other β_l, Eq. 15.6 simplifies to:

$$M_{ij}^{norm} = (\sigma_\kappa g_i)(\sigma_\kappa g_j) \tag{15.9}$$

where $\sigma_\kappa^2 = V_{\kappa\kappa}$.

The matrices in Eqs. 15.6–15.9 depend on the μ_i (or the constituent g_i), which are functions of $\boldsymbol{\theta}$. Thus, the $\mathbf{M}$ in Eq. 15.4 is typically a function of $\boldsymbol{\theta}$. As a result, $\mathbf{M}$ must be reweighted as $\boldsymbol{\theta}$ changes during the χ^2 minimization procedure.

15.1.3 Extended likelihood treatment

For a bin with $n_i \lesssim 10$, Eq. 15.6 is not a good approximation for uncertainties due to Poisson statistics. A more accurate formulation is the extended likelihood:

$$-2\ln\lambda(\boldsymbol{\theta}) = 2\left(\mu_i(\boldsymbol{\theta}) - n_i + n_i \ln\frac{n_i}{\mu_i(\boldsymbol{\theta})}\right) \tag{15.10}$$

15.1.4 Goodness of fit

An advantage of the binned MLE strategy is its close connection to a goodness-of-fit test. If the best-fit model is true, and if the n_i follow Gaussian distributions, Wilks' theorem demonstrates that as N approaches infinity, $-2\ln\lambda(\hat{\boldsymbol{\theta}})$ (or, where equivalent, $\chi^2_{min} = \chi^2(\hat{\boldsymbol{\theta}})$) asymptotically approaches a χ^2 distribution with degrees of freedom equal to $d = b - m$ [116]. In this limit, the p-value of the data is given by:

$$p = 1 - F_d(\chi^2_{min}) \tag{15.11}$$

where $F_d(\chi^2)$ is the cumulative PDF of the χ^2 distribution with d degrees of freedom. As usual, the p-value specifies the probability of observing results which deviate from the best-fit model at least as much as the actual data, assuming the best-fit model is correct. MC techniques can provide goodness-of-fit tests which are valid under more general conditions.

15.1.5 Confidence intervals

Uncertainties on $\hat{\boldsymbol{\theta}}$ can be expressed through confidence intervals. From a frequentist perspective, the confidence interval for a probability of $(1-\alpha)$ covers the values of $\boldsymbol{\theta}$ that, if true, would lead to observations as extreme as the actual data in at least the fraction $(1-\alpha)$ of identical experiments.

Table 15.1: Values of $\Delta\chi^2$ corresponding to the coverage probability $(1-\alpha)$ for $m=1$ and $m=2$ fit parameters [10].

$(1-\alpha)$	$\Delta\chi^2$ for $m=1$	$\Delta\chi^2$ for $m=2$
0.6827	1.00	2.30
0.90	2.71	4.61
0.95	3.84	5.99
0.9545	4.00	6.18
0.99	6.63	9.21
0.9973	9.00	11.83

If the $\mu_i(\boldsymbol{\theta})$ follow Gaussian distributions, confidence intervals can be constructed with reference to a χ^2 PDF. By a corollary to Wilks' theorem, the statistic $\Delta\chi^2(\boldsymbol{\theta}) = \chi^2(\boldsymbol{\theta}) - \chi^2_{min}$ follows a χ^2 distribution with m degrees of freedom. Note that the function $\Delta\chi^2(\boldsymbol{\theta})$ is equivalent to $-2\ln\mathcal{L}(\boldsymbol{\theta}) + 2\ln\mathcal{L}(\hat{\boldsymbol{\theta}}) = 2\ln(\mathcal{L}_{max}/\mathcal{L}(\boldsymbol{\theta}))$; that is, the $\Delta\chi^2$ statistic corresponds to the natural logarithm of a likelihood ratio. For a coverage probability $(1-\alpha)$, the corresponding value of $\Delta\chi^2$ can be found from $F_d(\chi^2)$, where $d = m$. The interval, or region for $m > 1$, in which $\Delta\chi^2$ is less than or equal to that value is the $(1-\alpha)$ confidence interval. Table 15.1 lists values of $\Delta\chi^2$ corresponding to various $(1-\alpha)$ probabilities. The probability $(1-\alpha) = 0.6827$ defines the 1σ uncertainty interval, reflecting the Gaussian PDF. Unless otherwise noted, parameter uncertainties in this thesis are quoted at the 1σ level.

In Double Chooz oscillation analyses, the uncertainties on the $\mu_i(\boldsymbol{\theta})$ are generally dominated by Poisson statistics, which have a Gaussian distribution for sufficiently large μ_i, and spectrum normalization uncertainties, which are Gaussian. Thus, the $\Delta\chi^2$ values in Tab. 15.1 provide reasonable approximations for confidence interval boundaries. For some results, alternative confidence intervals have been derived with the MC-based method of Feldman and Cousins [117], which remains valid if the $\mu_i(\boldsymbol{\theta})$ do not follow Gaussian distributions. The Feldman-Cousins prescription also accounts for $\hat{\boldsymbol{\theta}}$ which occur close to definitional constraints, such as the requirement that $\sin^2 2\theta_{13}$ be non-negative. Outcomes of this method agree well with the standard $\Delta\chi^2$ intervals [78].

15.1.6 Pull parameters

The model for the $\mu_i(\boldsymbol{\theta})$ often depends on parameters which are not intrinsically important but which affect the determination of $\hat{\boldsymbol{\theta}}$. For example, in Double Chooz, the number of Li+He events in the predicted spectrum is partially correlated to the best-fit $\sin^2 2\theta_{13}$, but it is not an immediately useful quantity outside the experiment. Variables of this kind can be treated as pull parameters, otherwise known as nuisance parameters.

A pull parameter, here denoted η_α, can often be constrained by independent measurements or predictions. For instance, the Li+He rate is estimated by the method in Sec. 13.3.3. Such an estimate, $\overline{\eta_\alpha}$, can be included in the χ^2 as an additional data point. This point becomes the central value for η_α, with associated uncertainty σ_α. If the estimate is uncorrelated with the μ_i, it adds a pull term, or penalty term, to the χ^2 function:

$$(\chi^2_{pull})_\alpha = \frac{(\overline{\eta_\alpha} - \eta_\alpha)^2}{\sigma^2_\alpha} \tag{15.12}$$

Adding the data point $\overline{\eta_\alpha}$ to the χ^2 increases b by one, while including the new parameter η_α increases m by one. The degrees of freedom, $d = b - m$, are unchanged. If η_α has a linear effect on the μ_i, the pull term treatment is generally equivalent to a covariance matrix treatment of the same parameter (see [118] for a proof; exceptions can occur if multiple normalization-type pull parameters appear in the same χ^2 term). That is, one way to propagate uncertainty on pull parameters $\boldsymbol{\eta}$ is to generate $\mathbf{M}$ according to Eq. 15.7, with $\boldsymbol{\beta} = \boldsymbol{\eta}$; use the resulting $\mathbf{M}$ in Eq. 15.4, in which $\boldsymbol{\eta}$ is fixed at its central value; and minimize the resulting χ^2. Another approach is to minimize the following statistic:

$$\begin{aligned}\chi^2(\boldsymbol{\theta}, \boldsymbol{\eta}) =& \big(\mathbf{n} - \boldsymbol{\mu}(\boldsymbol{\theta}, \boldsymbol{\eta})\big)^T \mathbf{M}^{-1}(\boldsymbol{\theta}, \boldsymbol{\eta}) \big(\mathbf{n} - \boldsymbol{\mu}(\boldsymbol{\theta}, \boldsymbol{\eta})\big) \\ &+ (\overline{\boldsymbol{\eta}} - \boldsymbol{\eta})^T \mathbf{V}^{-1} (\overline{\boldsymbol{\eta}} - \boldsymbol{\eta})\end{aligned} \tag{15.13}$$

where $\mathbf{V}$ is the covariance matrix for the pull parameters, such that $V_{kl} = \mathrm{cov}(\eta_k, \eta_l)$. Note that $\mathbf{M}$ may now depend on $\boldsymbol{\eta}$ as well as $\boldsymbol{\theta}$.

Including pull parameters typically changes the procedure for defining confidence intervals. If $\boldsymbol{\eta}$ were treated on a par with $\boldsymbol{\theta}$, confidence intervals would be defined following Tab. 15.1, where m counts the dimensions of both $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$. The result would be confidence regions for the joint $\boldsymbol{\eta}$-$\boldsymbol{\theta}$ probability of $(1 - \alpha)$. Since $\boldsymbol{\eta}$ is not valuable in itself, a more useful region corresponds to the $(1 - \alpha)$ probability for $\boldsymbol{\theta}$, regardless of the value of $\boldsymbol{\eta}$. This region is constructed using the m that

counts only the dimensions of $\boldsymbol{\theta}$. Such a construction, often called the profile likelihood method, conserves equivalences between the pull parameter and covariance matrix approaches.

In cases where both the pull term and covariance matrix treatments are possible, the choice is usually between extra information and computational efficiency. If a pull parameter η_i is used, its value at χ^2_{min} reveals the best fit, $\hat{\eta}_i$. The uncertainty on this value, often smaller than the input uncertainty, can be found by evaluating σ_α while minimizing over all other parameters. Typically, the price of this information is computation speed. Neither $\hat{\eta}_i$ nor its uncertainty can be extracted from the covariance matrix treatment, but because its χ^2 includes one fewer parameter than the pulls-based χ^2, minimization is generally faster.

The pull parameter treatment is also possible in cases where the covariance matrix treatment is not. One such case includes an η_α which affects μ_i nonlinearly. Another case includes an η_α with asymmetric uncertainties, $\sigma_\alpha^{up} \neq \sigma_\alpha^{down}$. For the latter case, the pulls covariance matrix $\mathbf{V}$ must be dynamically redefined during the χ^2 minimization procedure. For $\eta_\alpha \geq \overline{\eta_\alpha}$, the matrix is built using σ_α^{up}; otherwise, the matrix is built using σ_α^{down}.

15.2 Rate+Shape strategy

15.2.1 Motivation

For the same set of observations and the same prediction model, various χ^2 statistics can be defined using different choices of binning scheme. In the simplest setup, a Rate-Only fit, all observations are grouped into a single bin. For Double Chooz, this means integrating IBD candidates over the entire data-taking period and over the whole prompt energy range. The best-fit $\sin^2 2\theta_{13}$ is determined by comparing just two numbers: the total number of observed and predicted events. Daya Bay and RENO used Rate-Only fits in their initial $\sin^2 2\theta_{13}$ analyses [29, 30], and Double Chooz employs them as cross checks. An advantage of these measurements is their involvement of purely normalization uncertainties, which are generally easier to quantify than spectrum shape uncertainties. The drawback is that they discard information about the prompt spectrum and other potentially useful qualities of the IBD candidates.

While a deficit in the overall candidate rate is the strongest signal of oscillation in Double Chooz, information in the prompt energy spectrum can significantly enhance $\sin^2 2\theta_{13}$ precision. A

Rate+Shape χ^2 preserves much of this information by dividing candidates into multiple prompt energy bins. The superior precision of a Rate+Shape fit has been recognized since at least the previous Gd-based analysis. In that publication, the Rate+Shape result for $\sin^2 2\theta_{13}$ had a relative uncertainty of about 60%, while a Rate-Only analysis of the same dataset produced a relative uncertainty of about 80%. Rate+Shape fits provided the primary oscillation results for all subsequent publications through [86].

In Double Chooz fits, the power of the Rate+Shape statistic is subtle. Unlike in the KamLAND experiment, for example (see Fig. 3.3), the Double Chooz prompt spectrum is not dramatically distorted by oscillation. Less than one oscillation wavelength appears in the range 0.5 MeV $< E_{vis} <$ 20 MeV. Moreover, because the far detector baseline is shorter than $L_{max} = \frac{\pi}{2}\langle E_\nu \rangle / (1.267|\Delta m^2_{32}|)$ (as discussed in Sec. 7.2), the first oscillation maximum occurs below the prompt energy threshold. As a result, the effect of $\sin^2 2\theta_{13}$ on the prompt spectrum is not sharply distinct from a simple normalization shift, especially when blurred by statistical fluctuations. Thus, the signal-related information in the Rate+Shape fit is only modestly more revealing than the basic rate comparison, and normalization uncertainties from the reactor flux and detection efficiency still contribute greatly to the $\sin^2 2\theta_{13}$ error.

The stronger capability of the Rate+Shape fit is in identifying and constraining backgrounds. The prompt spectra of accidentals and FN+SM (Figs. 13.4 and 13.6) are markedly different from the IBD spectrum (for example, Fig. 6.3). The prompt spectrum of Li+He is more similar in shape, but its mean energy is a few MeV higher. These differences help the Rate+Shape fit to discriminate signal and background, while the Rate-Only fit faces total degeneracy. For example, the Rate+Shape fit can distinguish a deficit of observed candidates which is concentrated in the 1–3 MeV region from a deficit spread over the 2–8 MeV region; the former can be effectively fit with oscillation, while the latter may be better fit by a reduction in the Li+He normalization. This type of distinction, impossible in a Rate-Only fit, allows tighter constraints on backgrounds. Smaller background uncertainties are correlated to better precision on $\sin^2 2\theta_{13}$. Results presented at the end of this chapter demonstrate this effect.

Conceivably, binning candidates in dimensions beyond prompt energy could enhance Rate+Shape precision. One potentially useful dimension is the reactor power at the time when data was taken. Since the signal rate scales with reactor power, while the background rate does not, dividing data

by reactor power can, in principle, improve signal versus background discrimination. The previous Gd-based Rate+Shape fit explored this possibility by dividing IBD candidates into two time bins, corresponding to one- and two-reactor-on periods; no reactors-off data had yet been taken. This division complicated the fitting algorithm while decreasing the 1σ uncertainty on $\sin^2 2\theta_{13}$ by less than 1%, compared to a fit with a single time bin. Studies for the next analyses projected similarly small gains from separating one- and two-reactor data, so this strategy was discontinued. However, reactors-off data does receive a separate bin in the current Gd and H fits, as explained in Sec. 15.2.4. Meanwhile, multiple reactor(s)-on bins are used in the RRM fit described in Ch. 16.

15.2.2 Definition of χ^2 statistic

The Rate+Shape χ^2 statistic is defined according to Eq. 15.13. For the fits in this chapter, $\sin^2 2\theta_{13}$ is the only element in $\boldsymbol{\theta}$, so the statistic can be written:

$$\begin{aligned}\chi^2(\sin^2 2\theta_{13}, \boldsymbol{\eta}) =& \big(\mathbf{n} - \boldsymbol{\mu}(\sin^2 2\theta_{13}, \boldsymbol{\eta})\big)^T \mathbf{M}^{-1}(\sin^2 2\theta_{13}, \boldsymbol{\eta})\big(\mathbf{n} - \boldsymbol{\mu}(\sin^2 2\theta_{13}, \boldsymbol{\eta})\big) \\ &+ (\overline{\boldsymbol{\eta}} - \boldsymbol{\eta})^T \, \mathbf{V}^{-1} \, (\overline{\boldsymbol{\eta}} - \boldsymbol{\eta}) \\ &+ \chi^2_{off}(\sin^2 2\theta_{13}, \boldsymbol{\eta})\end{aligned} \tag{15.14}$$

The vector $\mathbf{n}$ counts IBD candidates observed in the reactors-on period, such that n_i is the number of candidates in the i^{th} prompt energy bin. The binning scheme is discussed in Sec. 15.2.3.

The vector $\boldsymbol{\mu}$ counts the predicted events, including antineutrino signal and all backgrounds. The μ_i depend on $\sin^2 2\theta_{13}$ and the pull parameters $\boldsymbol{\eta} = (\Delta m^2_{ee},\ a,\ b,\ c,\ R_{acc},\ R_{FN},\ R_{Li},\ F_0,\ F_1,\ F_2,\ \mu_{res})$ in the following way:

$$\begin{aligned}\mu_i =& \mu_i^{\nu}(\sin^2 2\theta_{13}, \Delta m^2_{ee}, a, b, c) \\ &+ \mu_i^{acc}(R_{acc}) \ + \ \mu_i^{FN}(R_{FN}, F_0, F_1, F_2) \ + \ \mu_i^{Li}(R_{Li})\end{aligned} \tag{15.15}$$

Here, μ_i^{α} is the number of events of type α expected in the i^{th} prompt energy bin. The number of antineutrino events, μ_i^{ν} depends explicitly on the oscillation parameters ($\sin^2 2\theta_{13}$ and the effective mass splitting, Δm^2_{ee}, defined in Sec. 15.3.3) and the energy scale parameters (a, b, c). It depends implicitly on the shape of the simulated prompt E_{vis} spectrum. The functional form of these dependencies is explained later in this chapter.

The accidental and Li+He predictions, μ_i^{acc} and μ_i^{Li}, are implicitly functions of the corresponding spectrum shapes. The same is true for the FN+SM count, μ_i^{FN}, in the Gd analysis. In the H analysis,

the FN+SM shape is explicitly parametrized by the three pull parameters F_0, F_1, and F_2. 15.3.7. Each of the expected background counts, μ_i^α (α = accidentals, FN+SN, Li+He), depends explicitly on the rate pull parameter R_α. The forms of these dependencies are also given later in this chapter.

The total covariance matrix $\mathbf{M}$ is the sum of matrices representing uncertainties from finite data statistics ($\mathbf{M}^{stat} = \mathbf{M}^{Pearson}$, as defined in Eq. 15.6), the reactor flux model ($\mathbf{M}^{rxtr}$), signal detection efficiency ($\mathbf{M}^{eff}$), and the spectrum shapes of the accidental and Li+He backgrounds ($\mathbf{M}^{acc}$ and $\mathbf{M}^{Li}$). These matrices are summed linearly because they represent uncorrelated uncertainties. The construction of each systematic uncertainty matrix is described in Sec. 15.3. Although Eq. 15.14 does not explicitly show this dependence, each matrix is a function of $\sin^2 2\theta_{13}$ and/or one or more pull parameter.

The remaining systematic uncertainties are treated with pull terms. Central values for the pull parameters, $\overline{\boldsymbol{\eta}}$, and elements of the pulls covariance matrix $\mathbf{V}$ are given in Sec. 15.3.

The term χ^2_{off}, which handles data from the reactors-off period, is explained in Sec. 15.2.4.

15.2.3 Prompt energy binning scheme

The choice of a prompt energy binning scheme balances several considerations. Since binning represents a loss of information, smaller bins are generally desirable. However, it is convenient to keep bins large enough to maintain $n_i \gtrsim 10$, so that Eq. 15.6 can be used for statistical uncertainties. In Double Chooz, the prompt E_{vis} bin size may also be limited by the granularity of the available $S_k(E_\nu)$ spectra discussed in Secs. 12.1.3 and 12.1.6. The detector resolution smears the E_ν to E_{vis} conversion enough that discrete steps are not visible in the E_{vis} spectrum, but as a conservative measure, E_{vis} bins are chosen to be at least as large as the $S_k(E_\nu)$ bins. For the current analyses, the available $S_k(E_\nu)$ have bins of width 0.25 MeV. Considering these factors, the majority of bins are sized at 0.25 MeV. Above 8 MeV, wider bins are chosen to ensure $n_i \gtrsim 10$. The binning schemes for the Gd and H fits are given in Tab. 15.2.

15.2.4 Reactors-off term

In Eq. 15.14, χ^2_{off} compares the number of IBD candidates observed in the reactors-off period, n_{off}, to the expectation in that period, μ_{off}. This term uses only rate information, since the small size of n_{off} leads to large statistical uncertainty on the prompt spectrum. Because $n_{off} < 10$ for the Gd

Table 15.2: Prompt energy binning schemes for the Gd- and H-channel Rate+Shape fits. Note that the Gd prompt energy spectrum begins at 0.5 MeV, while the H spectrum begins at 1.0 MeV. The Gd fit uses a total of $b = 40$ bins, while H uses $b = 38$ bins.

Coverage range (MeV)	Bin width (MeV)	Number of Gd bins	Number of H bins
0.5–1.0	0.25	2	0
1.0–8.0	0.25	28	28
8.0–10.0	0.50	4	4
10.0–12.0	1.00	2	2
12.0–20.0	2.00	4	4

dataset (see Tab. 13.4), the extended likelihood construction (Eq. 15.10) is used:

$$\chi^2_{off}(\sin^2 2\theta_{13}, \boldsymbol{\eta}) = 2\left(\mu_{off}(\sin^2 2\theta_{13}, \boldsymbol{\eta}) - n_{off} + n_{off} \ln\left(\frac{n_{off}}{\mu_{off}(\sin^2 2\theta_{13}, \boldsymbol{\eta})}\right)\right) \tag{15.16}$$

In the reactors-off period, expected events include the accidentals, Li+He, and FN+SM backgrounds and a small number of antineutrinos emitted by fission products with halflives of hours or longer. The number of these residual antineutrinos is treated with the pull parameter μ_{res}. The rate and energy spectrum of residual antineutrinos is predicted with a dedicated reactor simulation, described in [119]. Applying factors for detector acceptance and efficiency yields an expected number of residual antineutrino IBD interactions, $\overline{\mu_{res}}$, with uncertainty σ_{res}. Oscillation probabilities for the residual antineutrinos are computed using the average factor:

$$\omega_{res} = \left\langle \sin^2 \frac{1.267 \Delta m^2_{ee} L_i}{(E_\nu)_i} \right\rangle \tag{15.17}$$

where the average runs over the simulated residual antineutrinos and Δm^2_{ee} is defined in Sec. 15.3.3. The Gd- and H-channel values for $\overline{\mu_{res}}$, σ_{res}, and ω_{res} appear in Tab. 15.3.

The factors $f_\alpha = R^{off}_\alpha / R^{on}_\alpha$ are introduced to account for the difference between the expected rate of the background of type α in the reactor-off period (R^{off}_α) versus in the reactor(s)-on period (R^{on}_α). The small uncertainty on f_α and slight differences in uncertainties between R^{off}_α and R^{on}_α are neglected, as their effects in the fit are not significant. The total number of events expected in

Table 15.3: The expected number of IBD interactions from residual antineutrinos in the reactors-off period, $\overline{\mu_{res}} \pm \sigma_{res}$, and the average oscillation factor, $\omega_{res} = \langle \sin^2(1.267\Delta m^2_{ee} L/E_\nu) \rangle$, for both the Gd and H datasets.

Parameter	Gd	H
$\overline{\mu_{res}} \pm \sigma_{res}$	1.57 ± 0.47	2.34 ± 0.70
ω_{res}	0.75	0.85

the reactors-off period is thus:

$$\begin{aligned} \mu_{off}(\sin^2 2\theta_{13}, \boldsymbol{\eta}) = (1 - \sin^2 2\theta_{13}\ \omega_{res})\mu_{res} \\ + (f_{acc} R_{acc} + f_{FN} R_{FN} + f_{Li} R_{Li})\ T_{off} \end{aligned} \tag{15.18}$$

where T_{off} is the livetime of the reactors-off period.

Since it contains one additional data point (n_{off}) but no additional free parameters, χ^2_{off} adds one degree of freedom to the fit.

15.3 Systematic uncertainty treatments

15.3.1 Reactor flux uncertainty

Uncertainty on each parameter in Eq. 12.6 contributes uncertainty to the detectable E_ν spectrum. In turn, uncertainty on the E_ν spectrum propagates to the expected prompt E_{vis} spectrum. Error propagation for ϵ and N_p, which relate to signal detection, is described in Sec. 15.3.2. Uncertainties on the remaining parameters in Eq. 12.6, which relate to antineutrino generation in the reactors, are treated via the two-step procedure described in this section. In the first step, a covariance matrix is constructed for the E_ν spectrum. Since the Gd and H candidates are selected from the same dataset, this matrix, $\mathbf{M}^{rxtr}(E_\nu)$, is common to both analyses. In the second step, a multisim procedure (see App. A) translates $\mathbf{M}^{rxtr}(E_\nu)$ into a covariance matrix for the prompt E_{vis} spectrum.

To build $\mathbf{M}^{rxtr}(E_\nu)$, covariance matrices are first constructed for each source of uncertainty: baseline length ($\mathbf{M}^L_r$); thermal power ($\mathbf{M}^{P_{th}}_r$); mean energy per fission ($\mathbf{M}^{\langle E_f \rangle}_r$); mean cross section per fission, including the impact of the Bugey4 anchor, ($\mathbf{M}^{\langle \sigma_f \rangle}_r$); and the fission fractions for the four main isotopes ($\mathbf{M}^{\alpha_k}_r$). Uncertainties are assumed to be fully correlated between reactors B1

and B2. This treatment is certainly valid for the dominant source of uncertainty, the mean cross section per fission, since the same $\langle\sigma_f\rangle^{Bugey}$ and S_k are used for both reactors. Because $\langle E_f\rangle$ and the α_k for both reactors come from a common simulation, their uncertainties are also expected to be highly correlated. The baseline measurements are made with the same method, presumably sharing common systematics; moreover, uncertainties on L are almost negligible. Whether P_{th} uncertainties are entirely correlated between reactors is a subject of ongoing investigation. In the meantime, assuming fully correlated uncertainties is conservative for a single-detector analysis.

Each matrix is created according to Eq. 15.8, with the integral of Eq. 12.6 supplying $\boldsymbol{\mu}$. Uncertainties and correlations for each $\mathbf{V}$ are discussed in Sec. 12.1. These sources of uncertainty are considered uncorrelated, so the full covariance matrix, binned in E_ν, is:

$$\mathbf{M}^{rxtr}(E_\nu) = \mathbf{M}^{L} + \mathbf{M}^{P_{th}} + \mathbf{M}^{\langle E_f\rangle} + \mathbf{M}^{\langle\sigma_f\rangle} + \mathbf{M}^{\alpha_k} \tag{15.19}$$

All five sub-matrices in Eq. 15.19 add uncertainty to the signal normalization, as summarized in Tab. 15.4. Here, as elsewhere in this thesis, the signal normalization uncertainty in a matrix $\mathbf{M}$ is calculated in analogy to Eq. 15.9:

$$\sigma_{norm}(\mathbf{M}) = \frac{\sqrt{\sum_i^b \sum_j^b M_{ij}}}{\sum_i^b \mu_i^\nu} \tag{15.20}$$

Uncertainties on $\langle\sigma_f\rangle$ and the α_k also contribute shape uncertainty on the signal spectrum. The total normalization uncertainty from the reactor flux model is 1.7%, the largest such factor in both the Gd and H analyses. Without the Bugey4 anchor, this uncertainty would be approximately 2.7%.

Propagation of $\mathbf{M}^{rxtr}(E_\nu)$ into $\mathbf{M}^{rxtr}(E_{vis})$ is performed with the multisim method outlined in App. A. In brief, $\mathbf{M}^{rxtr}(E_\nu)$ is Cholesky decomposed and used to throw sets of correlated fluctuations on the predicted E_ν spectrum. For each throw, the new E_ν spectrum becomes a reweighting scheme for the predicted E_{vis} spectrum. A covariance matrix binned in E_{vis} is constructed from the E_{vis} spectrum variations. The multisim procedure is carried out separately for the Gd and H analyses, resulting in similar but distinct $\mathbf{M}^{rxtr}(E_{vis})$. These matrices are depicted in Figs. 15.3 and 15.4.

15.3.2 Detection efficiency uncertainty

As indicated in Secs. 13.2.5 and 14.2.5, the normalization of the signal MC receives the correction $C_{data/MC}$ before entering oscillation fits. The uncertainty on $C_{data/MC}$ quantifies how well the

Table 15.4: Signal normalization uncertainties contributed by parameters in the reactor flux model, for both the Gd and H analyses.

Source of uncertainty	Signal uncertainty (%)
Mean cross section per fission, with Bugey4 ($\langle\sigma_f\rangle$, $\langle\sigma_f\rangle^{Bugey}$)	1.4
Fractional fission rates (α_k)	0.8
Thermal power (P_{th})	0.5
Mean energy per fission ($\langle E_{th}\rangle$)	0.2
Baseline (L)	< 0.1
Total	1.7

efficiency of MC is known, relative to data, and is therefore equivalent to the uncertainty on ϵ in Eq. 12.6. The detection efficiency covariance matrix, $\mathbf{M}^{eff}$, is constructed according to Eq. 15.9, with $\mu_i = \mu_i^{\nu}$ and σ_κ given by the total uncertainty on $C_{data/MC}$ in Tab. 13.3 (Tab. 14.3) for the Gd (H) analysis. These matrices appear in Figs. 15.3 and 15.4.

15.3.3 Mass splitting uncertainty

As noted in Sec. 7.2, Double Chooz has very limited ability to measure the $|\Delta m^2|$ of the oscillations it observes. This $|\Delta m^2|$ can be treated as a free parameter in the Rate+Shape fit, but statistical scatter of the data points, along with the systematic effect discussed in Ch. 17, may cause the fit to favor unphysical values. This phenomenon was studied in the previous Gd-based analysis with MC-based frequentist tests. Using the multisim method (App. A), simulated datasets were created with $|\Delta m^2|$ at a representative value of 2.3×10^{-3} eV^2 and $\sin^2 2\theta_{13} = 0.1$. A Rate+Shape fit was performed on each dataset, with both oscillation parameters unconstrained. The most probable $|\Delta m^2|$ result was an order of magnitude larger than the actual $|\Delta m^2| = 2.3 \times 10^{-3}$ eV^2. This effect arose mainly from statistical fluctuations near the peak of the observed spectrum. Further tests showed that spurious $|\Delta m^2|$ results dominate even if the MC sample is doubled to approximately the current dataset size, and even if all systematic uncertainties are removed.

To avoid this issue, and to increase $\sin^2 2\theta_{13}$ precision, $|\Delta m^2|$ is constrained by a measurement made by the MINOS experiment. As mentioned in Sec. 3.3.2, MINOS observes the disappearance

of muon neutrinos or antineutrinos in an accelerator-produced beam. The $|\Delta m^2|$ controlling this oscillation is not exactly the same as that governing electron antineutrino disappearance, so an adjustment is made to the MINOS value, according to the following reasoning.

In MINOS, the muon neutrino disappearance probability can be approximated by the two-neutrino formula of Eq. 3.12, in which the first oscillation maximum occurs at $L = \frac{\pi}{2}E_\nu/|\Delta m^2_{32}|$. A slightly more accurate formula can be derived from the full three-neutrino model. If $P^{3\nu}_{\nu_\mu\to\nu_\mu}$ expresses the muon disappearance probability in this model, setting $\mathrm{d}(P^{3\nu}_{\nu_\mu\to\nu_\mu})/\mathrm{d}E_\nu = 0$ indicates the true location of the first oscillation maximum, $L = \frac{\pi}{2}E_\nu/\Delta m^2_{\mu\mu}$. In this treatment:

$$P_{\nu_\mu\to\nu_\mu} = 1 - \sin^2 2\theta_{23}\sin^2\left(\frac{\Delta m^2_{\mu\mu}L}{4E_\nu}\right) \tag{15.21}$$

where the effective mass splitting $\Delta m^2_{\mu\mu}$ is a mixture of $|\Delta m^2_{32}|$ and $|\Delta m^2_{31}|$. Thus, the $|\Delta m^2|$ measured by MINOS is not quite $|\Delta m^2_{32}|$ but $\Delta m^2_{\mu\mu}$, although the difference between the two is less than 2% [120].

Applying the same logic to electron antineutrino disappearance yields:

$$P_{\bar\nu_e\to\bar\nu_e} = 1 - \sin^2 2\theta_{13}\sin^2\left(\frac{\Delta m^2_{ee}L}{4E_\nu}\right) \tag{15.22}$$

where Δm^2_{ee} is its own mixture of Δm^2_{32} and Δm^2_{31}. To a good approximation (detailed in [121]), the effective mass splittings are:

$$\Delta m^2_{\alpha\alpha} = r_\alpha|\Delta m^2_{31}| + (1-r_\alpha)|\Delta m^2_{32}| \tag{15.23}$$

where:

$$r_\alpha = \frac{|U_{\alpha1}|^2}{|U_{\alpha1}|^2 + |U_{\alpha2}|^2} \tag{15.24}$$

with $U_{\alpha i}$ as defined in Eq. 3.3. The difference between Δm^2_{ee} and $\Delta m^2_{\mu\mu}$ is then (following [122]):

$$\begin{aligned}\Delta m^2_{ee} - \Delta m^2_{\mu\mu} &= \left(\cos 2\theta_{12} - \cos\delta\sin\theta_{13}\tan\theta_{23} + \mathcal{O}(\sin^2 2\theta_{13})\right)\left(|\Delta m^2_{32}| - |\Delta m^2_{31}|\right)\\ &\approx \pm(0.03\pm0.01)\times10^{-3}\ \mathrm{eV}^2\end{aligned} \tag{15.25}$$

where the numerical value is derived from global fits for the oscillation parameters [10]. The plus (minus) sign corresponds to the normal (inverted) mass hierarchy. When the current analyses began, the most precise measurement of $\Delta m^2_{\mu\mu}$ was the MINOS value of $2.41^{+0.09}_{-0.10}\times10^{-3}$ eV^2 [120]. From Eq. 15.25, the best estimate for Δm^2_{ee} is then $2.44^{+0.09}_{-0.10}\times10^{-3}$ eV^2 for the normal hierarchy

and $2.38^{+0.09}_{-0.10} \times 10^{-3}$ eV^2 for the inverted hierarchy. For nominal results, Double Chooz uses the normal hierarchy value. A fit using the inverted hierarchy value was performed for the Gd channel, producing very similar results.

In the Rate+Shape fit, Δm^2_{ee} is a pull parameter. Its central value is $\overline{\Delta m^2_{ee}} = 2.44 \times 10^{-3}$ eV^2, with asymmetric uncertainties $\sigma^{up}_{\Delta} = 0.09 \times 10^{-3}$ eV^2 and $\sigma^{down}_{\Delta} = 0.10 \times 10^{-3}$ eV^2.

15.3.4 Energy scale uncertainty

Energy-related uncertainty arises from systematic differences between the E_{vis} definitions in data (Eq. 11.7) and MC (Eq. 12.10). Since all background spectra are determined from data, this type of uncertainty impacts only the antineutrino signal. Energy-related uncertainties appear in all steps of the MC production, from scintillator modeling to readout electronics to the definition of E_{vis}. From a certain perspective, uncertainties also exist in the energy reconstruction for data, since the uniformity and stability corrections (described in Secs. 11.1.4 and 11.1.6) have statistical and systematic limitations.

In principle, errors could be propagated directly from all these sources. A simpler and approach is to gather all energy-related uncertainties into the errors on a small set of parameters used to correct the MC. In this perspective, E_{vis} in data is exactly defined by Eq. 11.7, and all uncertainties reflect how closely E_{vis} in MC approximates this fixed standard. Data-MC differences can exist in both energy scale and resolution, but sensitivity studies show that the resolution uncertainty in Double Chooz has negligible impact on oscillation fits. For this reason, the MC energy correction function treats only energy scale discrepancies.

Two energy scale correction functions, f_{QNL} and f_{LNL}, appear directly in the definition of E^{MC}_{vis} (Eq. 12.10). An additional function can be defined to correct data-MC discrepancies produced by the time stability calibration, applied to data through f_s, and the spatial uniformity maps adjustments applied to both data and MC through f_u. Because the stability and uniformity corrections are defined from H captures, they maintain data-MC agreement at the 2.2 MeV anchor point. However, they may introduce discrepancies at higher energies. In particular, the f_u may be energy dependent, and the energy dependence built into $(\delta f_s)^{gain}$ (Eq. 11.6) may not be perfectly representative. The function $f_{s/u}$ is designed to cover these effects. Since the impact of these effects is expected to be small, and their functional forms are not obvious, the simple linear model of $f_{s/u} = b_{s/u}$ is selected.

The central value of $b_{s/u}$ is taken to be one. Uncertainty on $b_{s/u}$ is not straightforward to quantify, and conservative approaches are adopted. Uncertainty due to the stability correction is estimated from the RMS of energy calibration points in data, after the stability correction is applied. For the Gd analysis, the RMS was evaluated for alpha decays of ^{212}Po ($\sim$ 1 MeV) and spallation neutron captures on Gd ($\sim$ 8 MeV). A linear interpolation was performed between these points, and that line was convoluted with the PDF of the signal MC. The integral of the convolution, a spectrum-weighted average RMS, is 0.50%. For the H analysis, the RMS of H captures energy (2.2 MeV) was also included in the interpolation, resulting in a spectrum-weighted average RMS of 0.34%. Uncertainty from the uniformity corrections is estimated as the RMS of the relative differences between the data and MC maps, $f_u^{data}(\rho, z)$ and $f_u^{MC}(\rho, z)$. For both the Gd and H maps, this value is 0.25%. The stability and uniformity uncertainties are uncorrelated, so the total uncertainty on $b_{s/u}$ is their quadratic sum. For the Gd (H) analysis, this value is 0.56% (0.42%). Central values, uncertainties, and correlations for a_{LNL}, b_{LNL}, b_{QNL}, and c_{QNL} are given in Sec. 12.2.3.

Including all three correction functions, the visible energy of prompt MC signals in the Rate+Shape fit, $E_{vis}^{MC,R+S}$, is defined as:

$$\begin{aligned} E_{vis}^{MC,R+S} &= E_{vis}^{MC} \;\cdot\; f_{QNL}(E_{vis}^{MC}) \;\cdot\; f_{LNL}(E_{vis}^{MC}) \;\cdot\; f_{s/u} \\ &= E_{vis}^{MC} \;(b_{QNL} + c_{QNL} E_{vis}) \left(\frac{a_{LNL}}{E_{vis}} + b_{LNL} \right) b_{s/u} \end{aligned} \tag{15.26}$$

where E_{vis}^{MC} is defined in Eq. 12.7. Collecting coefficients on like powers of E_{vis}^{MC}, Eq. 15.26 can also be expressed as:

$$E_{vis}^{MC,R+S} = a + b\; E_{vis}^{MC} + c\; (E_{vis}^{MC})^2 \tag{15.27}$$

where (a, b, c) are pull parameters defined by:

$$\begin{aligned} a &= a_{LNL} \cdot b_{s/u} \cdot b_{QNL} \\ b &= a_{LNL} \cdot b_{s/u} \cdot c_{QNL} + b_{LNL} \cdot b_{s/u} \cdot b_{QNL} \\ c &= b_{LNL} \cdot b_{s/u} \cdot c_{QNL} \end{aligned} \tag{15.28}$$

This representation highlights how the Rate+Shape energy model is effectively a second-order Maclaurin expansion:

$$E_{vis}^{MC,R+S} - E_{vis}^{MC} = a\; (E_{vis}^{MC})^0 + (b-1)(E_{vis}^{MC})^1 + c\; (E_{vis}^{MC})^2 \tag{15.29}$$

Table 15.5: Central values, uncertainties, and correlations for the energy scale pull parameters in the Gd- and H-channel Rate+Shape fits. These values are derived from the central values of $(a_{LNL}, b_{LNL}, b_{s/u}, b_{QNL}, c_{QNL})$ via the multisim procedure (App. A). Uncertainties on $\overline{a}$ and $\overline{b}$ are significantly larger in the H analysis than in the Gd analysis, and parameter correlations are different, because the H-channel light nonlinearity model is much less constrained.

Parameter	Gd value	H value
$\overline{a} \pm \sigma_a$	-0.027 ± 0.006 MeV	0 ± 0.067 MeV
$\overline{b} \pm \sigma_b$	1.012 ± 0.007	1.004 ± 0.022
$\overline{c} \pm \sigma_c$	-0.0001 ± 0.0006 MeV^{-1}	-0.0001 ± 0.0006 MeV^{-1}
ρ_{ab}	-0.304	0
ρ_{bc}	0.007	-0.102
ρ_{ca}	-0.286	0

Central values for the parameters (a, b, c) are derived from Eq. 15.28, using the central values of $(a_{LNL}, b_{LNL}, b_{s/u}, b_{QNL}, c_{QNL})$. The covariance matrix for (a, b, c) can be derived from the covariance matrix for $(a_{LNL}, b_{LNL}, b_{s/u}, b_{QNL}, c_{QNL})$ using an approximation similar to Eq. 15.8. Alternatively, the new covariance matrix can be derived with the multisim technique (App. A). The two methods produce very consistent results. Since the multisim method accounts for higher-order moments and is therefore slightly more accurate, its results are used to define the energy scale uncertainties (σ_a, σ_b, σ_c) and correlations (ρ_{ab}, ρ_{bc}, ρ_{ca}). Central values, uncertainties and correlations for the energy scale parameters are listed in Tab. 15.5. These values are incorporated into a set of correlated pull terms in the Rate+Shape χ^2.

15.3.5 Simplified energy scale treatment

The nonlinear energy model described in Sec. 15.3.4 results in a computationally intense χ^2 minimization procedure, for reasons to be explained in Sec. 15.4.3. This procedure is used for the main Rate+Shape results, but it is too slow and resource-heavy to use for additional diagnostic fits. A simplified energy scale treatment allows the much more efficient fitting algorithm described in Sec. 15.4.2. In this treatment, a and c are fixed at their nominal central values, and all uncertainties are

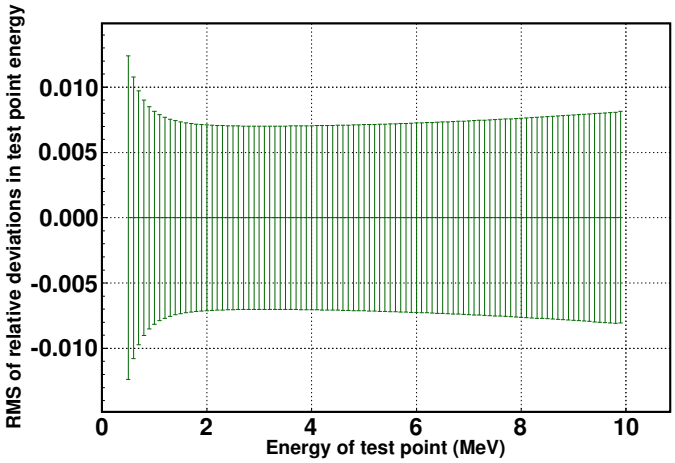

Figure 15.1: The spread in relative deviations in test point energies, $\delta E_{test}/E_{test} = (a + (b-1)E_{test} + c\ (E_{test})^2)/E_{test}$, for one million throws of (a, b, c) drawn from the Gd-channel values and their covariances. Test points are spaced by 0.1 MeV and cover the approximate prompt E_{vis} range for the signal MC. At each point, the RMS is plotted above and below the central horizontal axis. The spectrum-weighted average RMS supplies the effective linear uncertainty for the Gd energy scale, $\sigma_b^{eff} = 0.007$.

approximated as error on b. This effective linear model matches the approaches of previous Double Chooz analyses.

In the linear model, the uncertainty on b is adjusted to accommodate some of the variation allowed in the nonlinear model. The linear coefficient b cannot truly model nonlinear distortions of the energy spectrum. However, a modified σ_b can approximate the "uncertainty envelope" that exists in the nonlinear fit. The appropriate σ_b modification is determined from the following procedure: Many sets of correlated (a, b, c) are thrown using the multisim procedure (App. A). Each set is applied to a series of closely spaced energy test points, E_{test}. The relative shift in each test point, $\delta E_{test}/E_{test} = (a + (b-1)E_{test} + c\ (E_{test})^2)/E_{test}$ is recorded. The RMS of these differences is determined for each bin; an example set of RMS values is shown in Fig. 15.1. A weighted average is taken over these RMSs, with weights supplied by the signal MC PDF. This average, σ_b^{eff}, is defined as the effective linear uncertainty of the nonlinear model. Figure 15.2 illustrates how the uncertainty envelope of the effective linear model approximates that of the nonlinear model. Depending on the

Table 15.6: Central values and uncertainties for the background rate pull parameters in the Gd and H analyses. Uncertainty on the parameter R_{acc} includes only systematic uncertainty on the accidental rate; statistical uncertainty is handled in a covariance matrix.

Parameter	Gd rate (d^{-1})	H rate (d^{-1})
$\overline{R_{acc}} \pm \sigma_{acc}$	0.0701 ± 0.0026	4.334 ± 0.008
$\overline{R_{FN}} \pm \sigma_{FN}$	0.604 ± 0.051	1.55 ± 0.15
$\overline{R_{Li}} \pm \sigma_{Li}$	$0.97^{+0.41}_{-0.16}$	$0.95^{+0.57}_{-0.33}$

size of σ_a and σ_c and the energy parameter correlations, σ_b^{eff} may be larger or smaller than σ_b (or if $\sigma_a = \sigma_c = 0$, the procedure recovers $\sigma_b^{eff} = \sigma_b$). For the Gd analysis, $\sigma_b^{eff} = 0.007$; for the H analysis, $\sigma_b^{eff} = 0.035$.

15.3.6 Background rate uncertainties

Systematic uncertainties on the accidental, Li+He, and FN+SM rates are implemented through the pull parameters R_{acc}, R_{FN}, and R_{Li+He}. These parameters are applied in Eq. 15.15 so that:

$$\mu_i^{\alpha} = R_{\alpha} T_{on} \int_{E_i^{low}}^{E_i^{high}} f_{\alpha}(E_{vis})\ \mathrm{d}E_{vis} \tag{15.30}$$

where T_{on} is the livetime of the reactors-on period, E_i^{low} (E_i^{high}) is the lower (upper) bound of the i^{th} prompt energy bin, and f_{α} is the PDF of the prompt spectrum for event type α.

Central values and uncertainties for the R_{α} are summarized in Tab. 15.6, following the discussions in Secs. 13.3 and 14.3. Note that the accidental rates in Sec. 13.3.1 and 14.3.1 are quoted with both statistical uncertainty, from the off-time window measurements, and systematic uncertainty, from the off- to on-time correction factors. The accidental pull term covers only the systematic uncertainty; the statistical uncertainty is addressed in the following section.

15.3.7 Background spectrum shape uncertainties

Determination of the accidental and Li+He spectrum shapes is described in Secs. 13.3.1 and 13.3.3, respectively. Each of these measurements carries statistical uncertainty, which means that the spectrum shape in the Rate+Shape fit is not known exactly. In the H analysis, the FN+SM

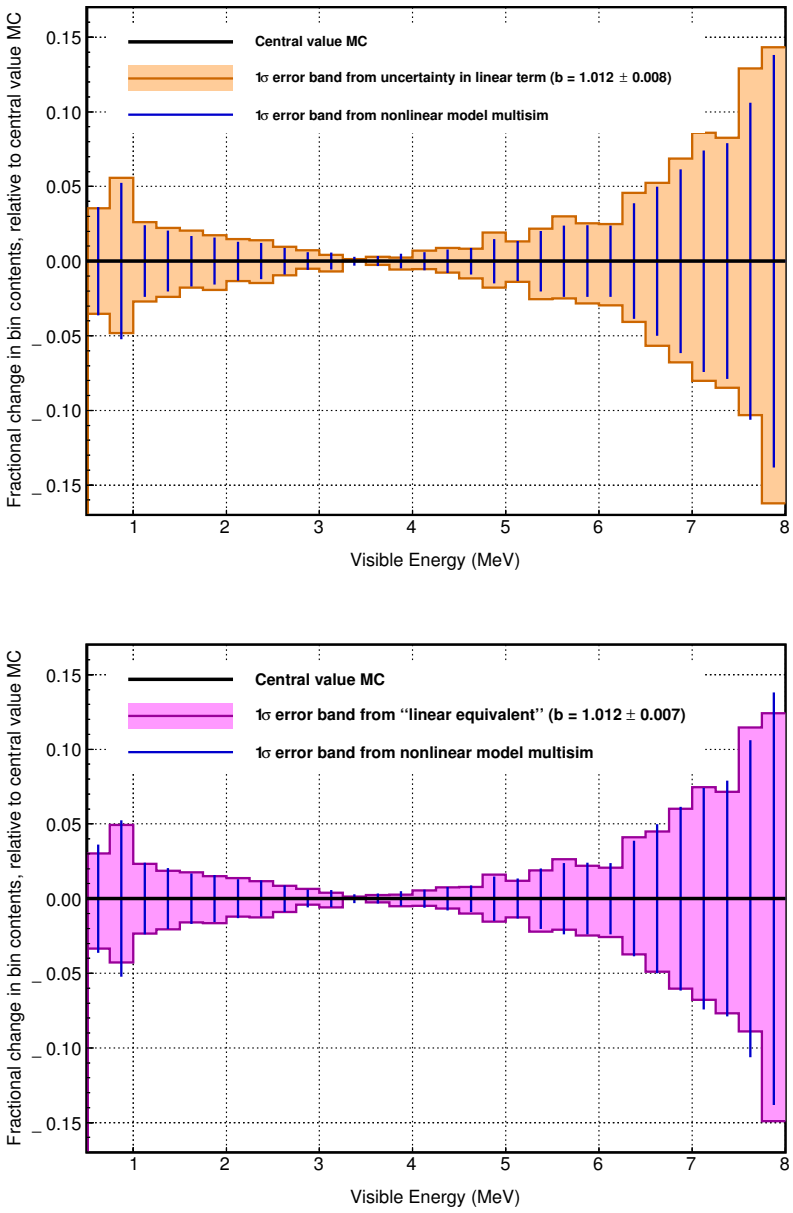

Figure 15.2: Comparisons of uncertainty envelopes around the central value energy scale in the Gd analysis. Each plot shows the fractional change in predicted signal events per bin allowed by $\leq 1\sigma$ deviations in the energy scale. The orange (magenta) region corresponds to $\sigma_b = 0.008$ ($\sigma_b^{eff} = 0.007$). The blue bars indicate the $\leq 1\sigma$ allowed region for the nonlinear energy model and are derived from a multisim procedure (App. A). Note that the orange region extends slightly beyond the region indicated by the blue bars, while the magenta region is nearly coextensive. Thus, σ_b^{eff} provides a better approximation of the nonlinear model uncertainties than σ_b alone.

spectrum also carries uncertainty from the functional fit for F_0, F_1, and F_2. The accidental and Li+He shape uncertainties are handled with covariance matrices $\mathbf{M}^{acc}$ and $\mathbf{M}^{Li}$, while the FN+SM shape parameters are constrained with pull terms.

In the accidental spectrum measured from off-time windows, the contents of each bin, y_i, has a relative statistical uncertainty of $1/\sqrt{y_i}$. These uncertainties are uncorrelated between bins. The predicted on-time accidentals spectrum shares these relative uncertainties. Therefore, a covariance matrix for the predicted spectrum can be constructed according to:

$$M_{ij}^{acc} = \delta_{ij} \left(\frac{\mu_i^{acc}}{\sqrt{y_i}} \right) \left(\frac{\mu_j^{acc}}{\sqrt{y_j}} \right) \tag{15.31}$$

In particular, for the nonzero elements:

$$M_{ii}^{acc} = \frac{(\mu_i^{acc})^2}{y_i} \tag{15.32}$$

It is convenient to cast this expression entirely in terms of the on-time spectrum and relative statistical uncertainty on the accidentals rate. Letting $y_{tot} = \sum_i^b y_i$ and $\mu_{tot}^{acc} = \sum_i^b \mu_i^{acc}$, the off-time bin contents can be expressed as:

$$y_i = \left(\frac{\mu_i^{acc}}{\mu_{tot}^{acc}} \right) y_{tot} \tag{15.33}$$

The total number of events in the off-time sample can be derived from the statistical uncertainty, σ_{acc}^{stat}, on the accidentals rate, R_{acc}:

$$y_{tot} = \frac{1}{(\sigma_{acc}^{stat}/R_{acc})^2} \tag{15.34}$$

From Eqs. 15.32–15.34:

$$M_{ii} = \mu_i^{acc} \mu_{tot}^{acc} \left(\frac{\sigma_{acc}^{stat}}{R_{acc}} \right)^2 \tag{15.35}$$

This formula is used to build the $\mathbf{M}^{acc}$, which appear in Figs. 15.3 and 15.4. In addition to expressing the spectrum shape uncertainty due to statistical fluctuations in each bin, $\mathbf{M}^{acc}$ represents the overall rate uncertainty from the off-time measurement. Note that using Eq. 15.35 in Eq. 15.20, with μ_i^{ν} replaced with μ_i^{acc}, results in the expected $\sigma_{norm}(\mathbf{M}^{acc}) = \sigma_{acc}^{stat}$.

Spectrum shape uncertainty is derived somewhat differently for Li+He. The Li+He rate does not come from the same measurement as the Li+He shape, so statistical uncertainty on the spectrum affects only the latter. Constructing the Li+He shape covariance matrix begins with constructing

Table 15.7: Central values and uncertainties for the fast neutron shape parameters in the H analysis. The FN+SM shape is fixed in the Gd fit.

Parameter	H value
$\overline{F_0} \pm \sigma_{F_0}$	12.52 ± 1.36
$\overline{F_1} \pm \sigma_{F_1}$	0.042 ± 0.015
$\overline{F_2} \pm \sigma_{F_2}$	0.79 ± 1.39

the covariance matrix for the $\mathcal{L}_{Li}$-vetoed spectrum shown in Fig. 13.6. Uncertainties on this spectrum are purely statistical, including the statistics of the final spectrum and subtracted background. The matrix $\mathbf{M}^{Li-veto}$ is then decomposed into its normalization, shape, and mixed components, as defined in App. A. The sum of the shape and mixed components, scaled to match the predicted number of Li+He events, becomes $\mathbf{M}^{Li}$. Like $\mathbf{M}^{acc}$, $\mathbf{M}^{Li}$ contains sizable diagonal elements, corresponding to uncertainty which is uncorrelated across bins. The very small, negative off-diagonal elements of $\mathbf{M}^{Li}$ reflect the subtraction of the small normalization component in $\mathbf{M}^{Li-veto}$. Since the same Li+He spectrum is used for the Gd and H channels, $\mathbf{M}^{Li}$ has the same form in both analyses. These matrices are shown in Figs. 15.3 and 15.4.

Sensitivity studies have show that the small amount of slope uncertainty present in the Gd-channel FN+SM spectrum would have negligible impact on a Rate+Shape fit. Consequently, no FN+SM shape uncertainty is included in the Gd fit. In the H fit, uncertainties on the three shape parameters F_0, F_1, and F_2 are treated with correlated pull terms. Table 15.7 lists central values and uncertainties for these parameters.

15.3.8 Summary of uncertainties

The covariance matrices used in the Gd and H fits are pictured in Figs. 15.3 and 15.4, respectively. The signal normalization impact of each matrix can be evaluated with Eq. 15.20, producing the results in Tab. 15.8 (background shape uncertainties are not listed, as they do not affect normalization). The same table lists normalization uncertainties from the background rate pull terms. The effect of the mass splitting and energy scale pull parameters is almost completely confined to the spectrum shape, with negligible impact on the signal normalization.

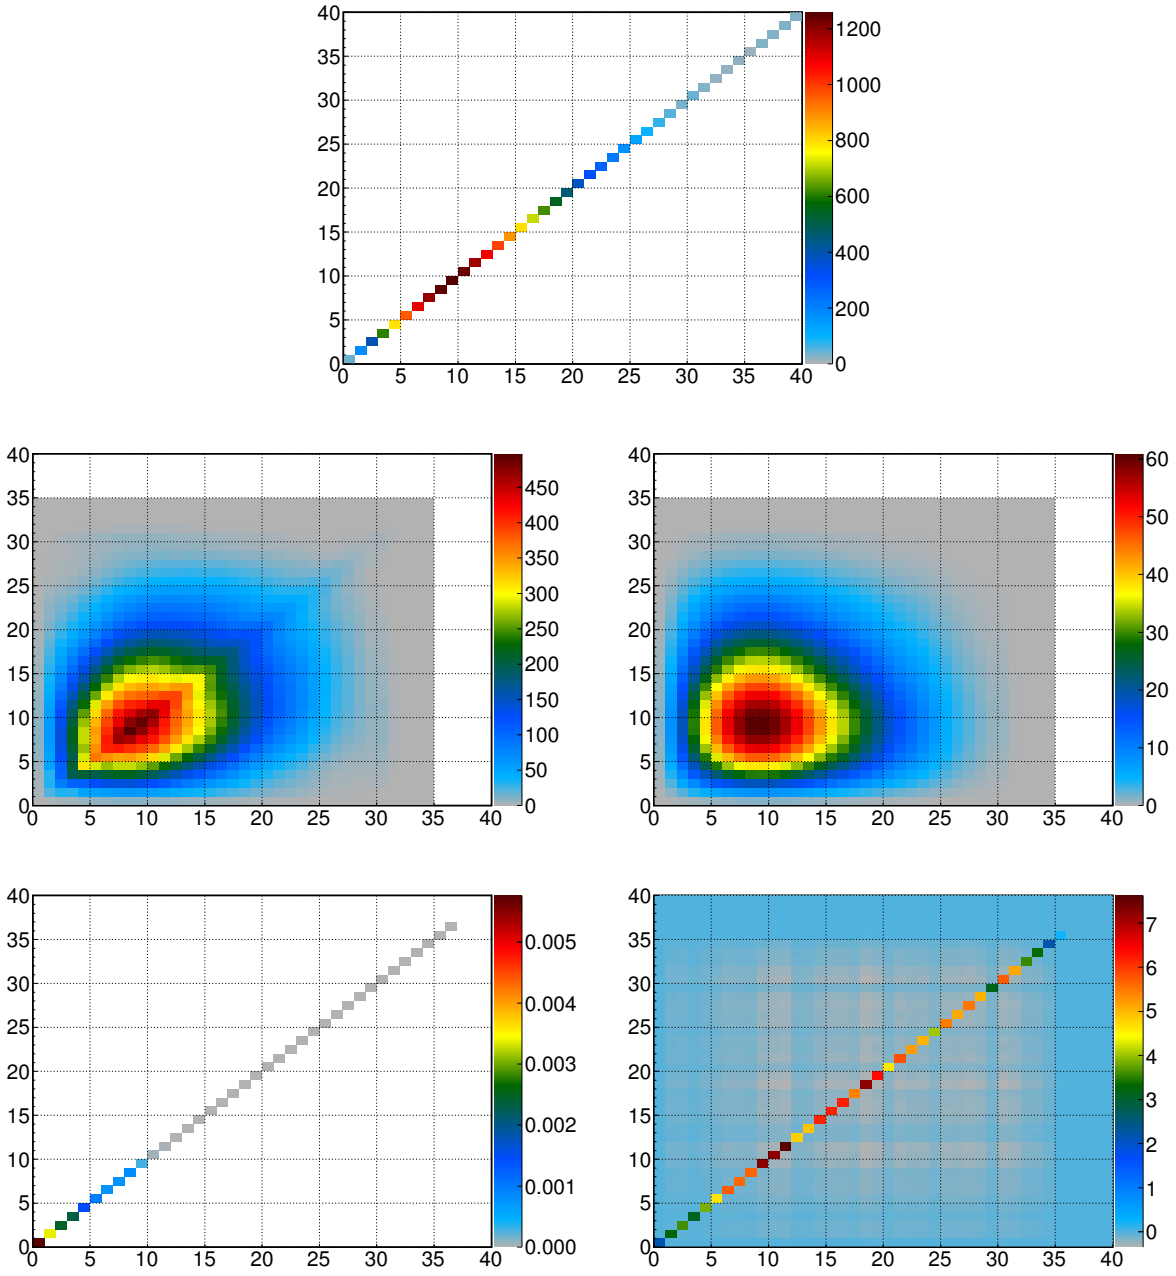

Figure 15.3: Covariance matrices in the Gd analysis, covering uncertainties from finite statistics (top), reactor flux (center left), detection efficiency (center right), accidental spectrum shape (bottom left), and Li+He spectrum shape (bottom right). Each matrix is weighted by the nominal, no-oscillation prediction.

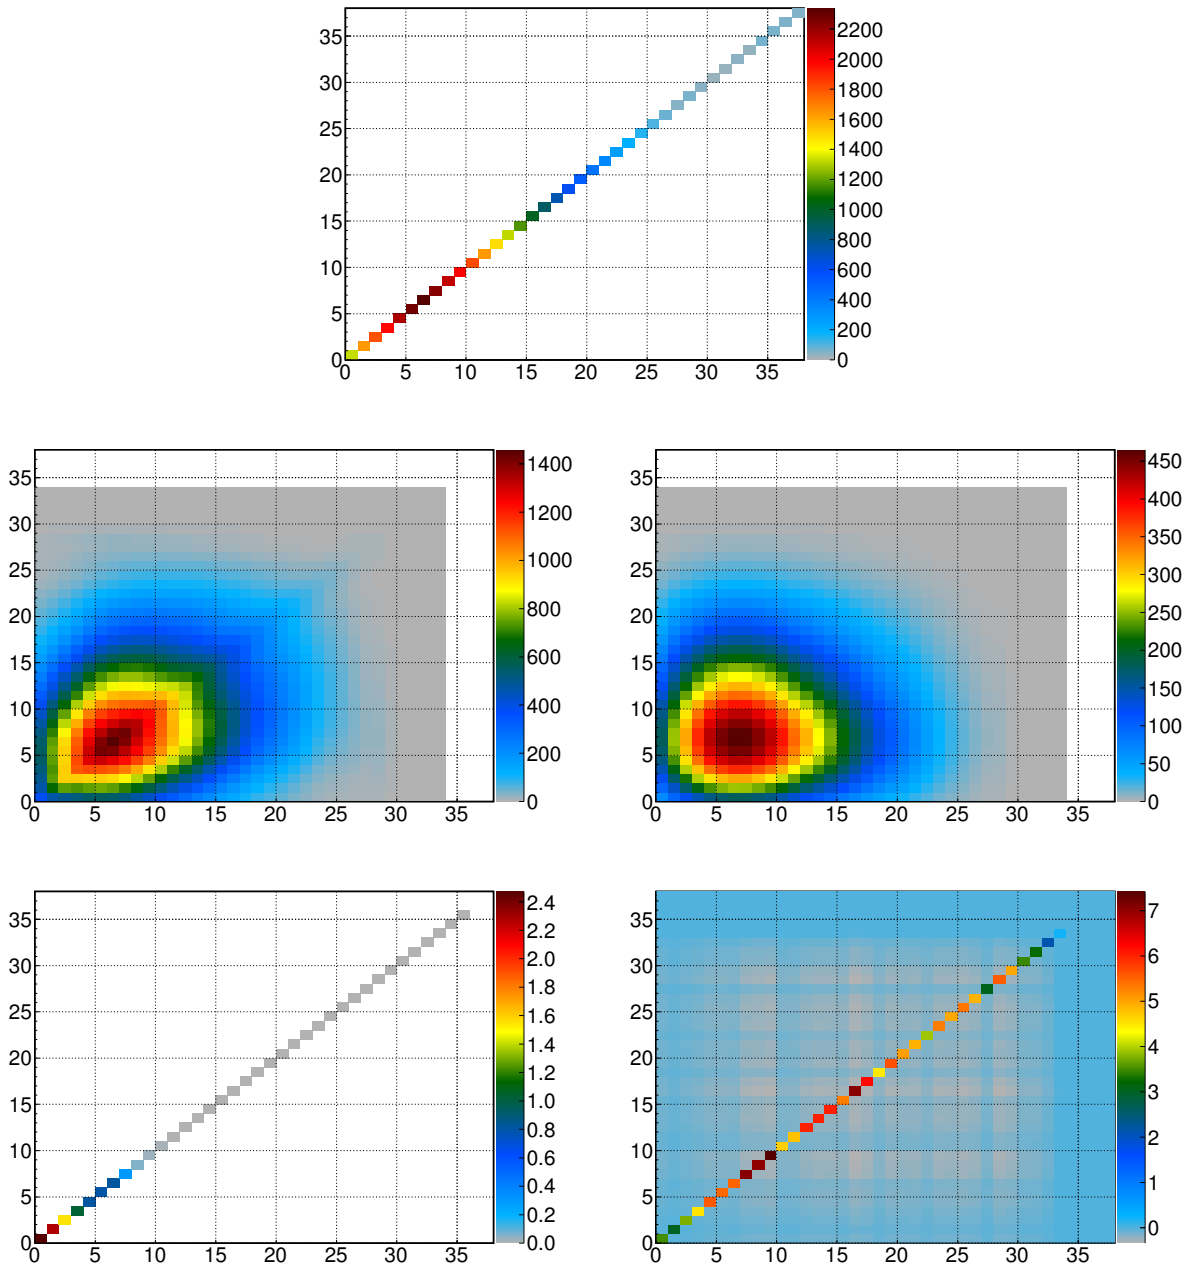

Figure 15.4: Covariance matrices in the H analysis, covering uncertainties from finite statistics (top), reactor flux (center left), detection efficiency (center right), accidental spectrum shape (bottom left), and Li+He spectrum shape (bottom right). Each matrix is weighted by the nominal, no-oscillation prediction.

Table 15.8: Signal normalization uncertainties contributed by each covariance matrix in the Gd- and H-channel Rate+Shape fits.

Source of uncertainty	Gd signal uncertainty (%)	H signal uncertainty (%)
Reactor flux	1.7	1.7
Signal detection efficiency	0.6	1.0
Li+He rate	$+1.1, -0.4$	$+0.9, -0.5$
FN+SM rate	0.1	0.2
Accidental rate	< 0.1	< 0.1
Statistics	0.8	0.6

15.4 χ^2 minimization and confidence interval estimation

After each component in the Rate+Shape χ^2 (Eq. 15.14) has been defined, this statistic can be minimized to find the best estimators for $\sin^2 2\theta_{13}$ and the pull parameters. Minimization consists of testing different values of $\sin^2 2\theta_{13}$ and $\boldsymbol{\eta}$ in the χ^2 function, according to some prescribed algorithm, until a global minimum is located. At each test point, the predicted spectrum $\boldsymbol{\mu}$ and covariance matrix $\mathbf{M}$ must be recalculated according to the current values of $\sin^2 2\theta_{13}$ and $\boldsymbol{\eta}$. This reweighting procedure is described in Sec. 15.4.1. Two different minimization algorithms are employed, depending on the energy scale treatment. These algorithms are outlined in Secs. 15.4.2 and 15.4.3. Appendix B describes the software package used to carry out these procedures.

15.4.1 Spectrum and matrix reweighting

At each χ^2 test point, the predicted bin contents μ_i are determined according to Eq. 15.15. The background counts, μ_i^α (α = accidental, Li+He, FN+SM), are calculated with Eq. 15.30. The signal count, μ_i^ν, is as a weighted sum of the signal MC events falling into the i^{th} bin. The weight w_i of each event includes the statistical scaling factor $1/100$, the MC correction factor $C_{data/MC}$, and the survival probability, $P_{\bar{\nu}_e \to \bar{\nu}_e}$. The latter factor is calculated using the simulated neutrino energy

and baseline for that event. Thus,

$$\begin{aligned}\mu_i^{\nu} &= \sum_j^{N_i} w_j \\ &= C_{data/MC} \cdot \frac{1}{100} \cdot \sum_j^{N_i} P_{\bar{\nu}_e \to \bar{\nu}_e}((E_\nu)_j, L_j, \sin^2 2\theta_{13}, |\Delta m^2_{ee}|)\end{aligned} \tag{15.36}$$

where N_i is the number of signal MC events with $E_i^{low} \leq E_{vis}^{MC,R+S} < E_i^{high}$. The survival probability follows Eq. 15.22.

Covariance matrices are reweighted in order to maintain the relationships in Eqs. 15.7 and 15.6, despite changes in μ_i resulting from shifts in $\sin^2 2\theta_{13}$ and the pull parameters. The statistical uncertainty matrix, $\mathbf{M}^{stat}$, is reweighted according to:

$$M_{ij}^{stat} = \sqrt{\frac{\mu_i}{\overline{\mu_i}} \frac{\mu_j}{\overline{\mu_j}}} \, \overline{M_{ij}^{stat}} \tag{15.37}$$

Each systematic covariance matrix, $\mathbf{M}^{\alpha}$ (α = reactor flux, detection efficiency, accidental shape, Li+He shape), is reweighted according to[1]:

$$M_{ij}^{\alpha} = \frac{\mu_i^{\alpha}}{\overline{\mu_i^{\alpha}}} \frac{\mu_j^{\alpha}}{\overline{\mu_j^{\alpha}}} \, \overline{M_{ij}^{\alpha}} \tag{15.38}$$

where $\overline{\mu_i^{\alpha}}$ is the nominal, no-oscillation prediction for α-type events in the i^{th} bin, and $\overline{M_{ij}^{\alpha}}$ is the covariance matrix constructed for that nominal prediction.

The reweighting scheme in Eq. 15.38 conserves relative uncertainties in each bin. This treatment is not exactly correct for $\mathbf{M}^{rxtr}$ when energy scale pull parameters are used. Movement of a, b, and c changes the $E_\nu \to E_{vis}^{MC}$ relationship for signal events, which changes the relative reactor flux shape uncertainties in each prompt E_{vis} bin. A fastidious treatment of this effect could proceed iteratively, with a series of $E_\nu \to E_{vis}^{MC}$ multisims being performed to generate new reactor matrices as the fit moves through (a,b,c) space. However, this time-intensive procedure would likely have negligible impact on fit results. The reason is that $\lesssim 2\sigma$ changes in the the energy scale parameters cause only small deformations in the signal MC spectrum, relative to the scale of energy-dependence in the reactor flux shape uncertainties. That conclusion is supported by a test in which $\mathbf{M}^{rxtr}$ was regenerated using best-fit values for (a,b,c). A fit performed with this alternative $\mathbf{M}^{rxtr}$ found the

[1] In the Gd fit, $\mathbf{M}^{rxtr}$ and $\mathbf{M}^{eff}$ were reweighted by $\boldsymbol{\mu}$, rather than $\boldsymbol{\mu}^{\nu}$. The difference between these weights is, on average, a few percent. A later reevaluation using $\boldsymbol{\mu}^{\nu}$ produced consistent results.

same best-fit point as in the original setup, to within 1% in each parameter. Thus, for computational simplicity, Eq. 15.38 is the only reweighting scheme applied to $\mathbf{M}^{rxtr}$.

15.4.2 Fully MINUIT-based fit

If the χ^2 statistic is a smooth function of $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$, with no local minima beyond the global minimum, it can be minimized with routines in the MINUIT package [123], a widely used tool in high energy physics. Here, smooth means continuously differentiable with respect to each element in $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$. These conditions are met in Rate+Shape fits with a well constrained Δm^2 and with energy scale uncertainty confined to the linear parameter b (i.e., $\sigma_a = \sigma_c = 0$), provided that b uncertainty is treated with the following special procedure.

During a χ^2 minimization, it would seem natural to recalculate the $E_{vis}^{MC,R+S}$ for each signal MC event, using the current value of b, in the same way that $P_{\bar{\nu}_e \to \bar{\nu}_e}$ is recalculated based on the current values of $\sin^2 2\theta_{13}$ and Δm^2_{ee}. However, that approach creates a problem if the $E_{vis}^{MC,R+S}$ of an event crosses a bin boundary. Such a shift moves the event to different bin, causing a discrete change in μ_i^ν (specifically, through a change in the integer N_i in Eq. 15.36). This change creates a discontinuity in $\partial\chi^2/\partial b$. The minimization algorithms in MINUIT rely on gradient estimations and do not perform properly when the χ^2 gradient is not continuous. Note that such discontinuities do not arise from $\sin^2 2\theta_{13}$ or any pull parameters other than a, b, and c.

Discontinuities in $\partial\chi^2/\partial b$ can be seen as a result of the finite statistics of the signal MC. If the signal MC could be modeled as a continuous PDF, like the FN+SM shape in the H fit, these features would not occur. In principle, it might be possible to construct a continuous PDF for the signal MC. Applying oscillation effects to this PDF would require a model for the relationship between E_ν and E_{vis}^{MC}. Since this relationship is not analytic, an MC-based mechanism such as an $E_\nu \to E_{vis}^{MC}$ migration matrix would be needed.

A simpler solution is to model changes in b directly through their impact on the μ_i^ν, rather than through their impact on the energy of individual events. This treatment relies on the fact that b has a nearly linear effect on the μ_i and can therefore be approximated by a Gaussian treatment. The method begins with the construction of a bin modulation vector, $\mathbf{s}^\nu$. One version of this vector, $(\mathbf{s}^\nu)_+$ is built with:

$$(s_i^\nu)_+ = \frac{\mu_i^\nu(b) - \mu_i^\nu(b+\sigma_b)}{\mu_i^\nu(b)} \tag{15.39}$$

That is, $(s_i^\nu)_+$ is the relative change in μ_i^ν due to a 1σ upward change in b. An analogous $(\mathbf{s}^\nu)_-$ vector can be constructed for a 1σ downward change in b. As an illustration, the magenta region in Fig. 15.2 corresponds to $(s_i^\nu)_+$ and $(s_i^\nu)_-$ for σ_b^{eff} in the Gd analysis. Because b has a close to linear effect on μ_i, the absolute magnitudes of the $(s_i^\nu)_+$ are nearly the same as those of the $(s_i^\nu)_-$. Moreover, a 2σ upward change in b affects the μ_i in almost the same way as $2(s_i^\nu)_+$. To average the small differences between $|(s_i^\nu)_+|$, $|(s_i^\nu)_-|$, and effects of changes in b other than 1σ, a $(\mathbf{s}^\nu)_{ms}$ vector is built through the multisim procedure in App. A. The definition of $(\mathbf{s}^\nu)_{ms}$ is implicitly dependent on $\sin^2 2\theta_{13}$ and Δm^2_{ee}, through the μ_i^ν. For $\sin^2 2\theta_{13}$ analyses with well constrained Δm^2_{ee}, this dependence is very small and can be neglected.

Changes in b are then propagated through application of $(\mathbf{s}^\nu)_{ms}$ to the signal MC spectrum. That is, a $k\sigma$ change in b results in the μ_i^ν:

$$\mu_i^\nu = \left(\sum_j^{N_i} w_j\right)\left(1 + k\cdot|(s_i^\nu)_{ms}|\right) \tag{15.40}$$

with w_i defined as in Eq. 15.36 and N_i counting events with $E_i^{low} \leq \left(\overline{a} + \overline{b}\ E_{vis}^{MC} + \overline{c}\ (E_{vis}^{MC})^2\right) < E_i^{high}$. The μ_i^ν depend smoothly on k, so $\partial\chi^2/\partial b$, by way of $\partial\chi^2/\partial k$, becomes continuous.

The resulting χ^2 can be minimized with MINUIT. The MIGRAD algorithm is used to find $\chi^2_{min} = \chi^2(\widehat{\sin^2 2\theta_{13}}, \hat{\boldsymbol{\eta}})$. In most cases, the MINOS algorithm is used to find the 1σ errors. Following the profile likelihood approach, the $+1\sigma$ and -1σ errors for $\sin^2 2\theta_{13}$ are defined at $\chi^2(\sin^2 2\theta_{13}, \hat{\hat{\boldsymbol{\eta}}}) = \chi^2_{min} + 1$, where the double hat indicates minimization over all parameters in $\boldsymbol{\eta}$. Similarly, the $\pm1\sigma$ errors for the η_i are defined at $\chi^2(\eta_i, \widehat{\widehat{\sin^2 2\theta_{13}}}, \hat{\hat{\eta}}_j(j\neq i)) = \chi^2_{min} + 1$. When the approximation of symmetric 1σ errors is acceptable, MIGRAD is followed by HESSE instead of the slower MINOS.

15.4.3 Lattice-based fit

As noted, the MINUIT algorithms are unreliable if the χ^2 function is not smooth. They are also limited to finding local minima. That is, the algorithms locate the minimum immediately downhill from the starting point in parameter space, which is not necessarily the global minimum. An example of the former case is a fit using the nonlinear energy model of Sec. 15.3.4; an example of the latter is a fit with unconstrained Δm^2. In both cases, the problem for MINUIT is that the local χ^2 gradient is not a reliable guide toward the global minimum.

The solution is a χ^2 minimization procedure that does not depend on gradients. One such procedure is a χ^2 scan, in which the χ^2 function is sampled at discrete, regularly spaced points across a wide range of fit parameter space. The test point lattice need only cover the dimensions in which the χ^2 is not smooth or has local minima. Minimization in the other dimensions can be performed at each test point with MINUIT. The test point with the smallest χ^2 provides χ^2_{min} and estimators for the fit parameters. Confidence regions can be constructed by scanning the lattice for the appropriate $\Delta\chi^2$ contours. Lattice-based χ^2 scans are computationally inefficient, especially when the test points are closely spaced to allow high fit precision. Efficiency can be improved by performing an initial scan at low resolution and performing higher resolution scans near the global minimum and confidence region boundaries.

In the main fits for the current analyses, a χ^2 scan is used for the a, b, and c dimensions. This strategy was chosen following an unsuccessful campaign to develop a fully MINUIT-based minimization for the nonlinear energy model. The s_i^ν technique discussed in the previous section is not viable for a and c, because their effects on the μ_i are significantly nonlinear. Furthermore, the effects of multiple energy scale parameters cannot be modeled independently. In principle, the s_i^ν method could be expanded to include a three-dimensional lattice of $s_i^\nu(\delta a, \delta b, \delta c)$, each expressing the relative change in μ_i^ν due to changes of $(\delta a, \delta b, \delta c)$ in the energy scale parameters. If a function $f(\delta a, \delta b, \delta c)$ could be developed to smoothly interpolate between the lattice points, it could replace $k \cdot |(s_i^\nu)_{ms}|$ in Eq. 15.40, making χ^2 differentiable with respect to a, b, and c.

Multiple interpolation schemes were attempted for $f(\delta a, \delta b, \delta c)$. The most promising candidate was a tricubic interpolation, in which both f and its gradient are continuous across the entire lattice. Still, the finite statistics of the MC cause a small amount of statistical scatter in the $s_i^\nu(\delta a, \delta b, \delta c)$ values, and interpolation between some points can create shallow local minima or slightly distorted gradients. These features are limited to particular regions of parameter space, so the interpolation method can be used for some applications, such as the plots in Fig. 15.5. However, the method is not sufficiently robust for overall χ^2 minimization.

Ultimately, the χ^2 functions are minimized through a simple lattice scan, with no interpolation. First, a coarse scan is performed in the three energy scale dimensions, with lattice spacing equal to approximately 0.1σ and boundaries at approximately $\pm 3\sigma$ in each dimension. At each test point, the χ^2 is is minimized over the other pull parameters and $\sin^2 2\theta_{13}$ with MIGRAD. A finer scan is

performed in the region of the χ^2 minimum. The resolution of this scan is limited by the signal MC statistics, but 100× data statistics permits determination of χ^2_{min} at the desired precision.

The $\pm 1\sigma$ confidence intervals for $\sin^2 2\theta_{13}$ and the pull parameters is also determined from lattice scans. The approximate boundaries of each interval are determined from MINUIT-based tests and coarse four-dimensional lattice scans, where the fourth dimension is the parameter of interest. Fine four-dimensional scans are then performed in the appropriate regions to precisely determine the confidence interval boundaries.

Determining χ^2_{min} and all $\pm 1\sigma$ errors for either the Gd or H analysis requires χ^2 evaluations at $\mathcal{O}(10^6)$ test points. Each evaluation requires $\mathcal{O}(1$ minute) of CPU time, such that the entire minimization requires $\mathcal{O}(1$ day) on a $\sim$ 500 CPU computing cluster. This procedure is feasible for the main fits, but it is not practical for additional diagnostic fits. Those fits instead use fully MINUIT-based strategy involving the simplified energy model, which require less than one hour on a single CPU. Improving χ^2 minimization efficiency, whether through a MINUIT-based PDF-driven method (as outlined in Sec. 15.4.2) or a more sophisticated non-MINUIT algorithm, is a priority for future Double Chooz analyses.

15.5 Gd-channel results

15.5.1 Best fit parameters

The Rate+Shape fit for the Gd analysis yields a best estimate of $\sin^2 2\theta_{13} = 0.090^{+0.032}_{-0.029}$ at $\chi^2_{min} = 52.2$. The fit contains $d = b - m + 1 = 40$ degrees of freedom (d.o.f.), where $b = 40$ for the number of prompt energy bins; $m = 1$ for the single unconstrained parameter, $\sin^2 2\theta_{13}$; and the additional degree of freedom comes from the reactors-off term. The χ^2_{min}/d.o.f. corresponds, via Eq. 15.11, to a p-value of 9%; reasons for this relatively small value are discussed below.

Best-fit values for all fit parameters, including the pull parameters, are listed in Tab. 15.9. Note that the R_{Li} and R_{FN} are significantly constrained by the fit. These results are based on the Δm^2_{ee} value corresponding to the normal mass hierarchy. A fit using the inverted hierarchy value produces $\sin^2 2\theta_{13} = 0.092^{+0.033}_{-0.029}$, with $\chi^2_{min} = 52.2/40$ d.o.f.

According to the MC-based frequentist technique of Feldman and Cousins [117], the data excludes the no-oscillation hypothesis at the 99.9% confidence level, corresponding to 3.2σ in a one-

Table 15.9: Results of the Gd-channel Rate+Shape fit. The first eight rows list pull parameters; their definitions are given in the text.

Parameter	Input central value	Input error	Output best fit	Output error
a	-0.027	0.006	-0.026	$+0.006, -0.005$
b	1.012	0.008	1.011	$+0.004, -0.006$
c	-0.0001	0.0006	-0.0006	$+0.0007, -0.0005$
R_{FN} (d^{-1})	0.604	0.051	0.568	$+0.038, -0.037$
R_{Li} (d^{-1})	0.97	$+0.41, -0.16$	0.74	0.13
R_{acc} (d^{-1})	0.0701	0.0026	0.0703	0.0026
μ_{res}	1.57	0.47	1.48	0.47
Δm^2_{ee} (10^{-3} eV^2)	2.44	$+0.09, -0.10$	2.44	$+0.09, -0.10$
$\sin^2 2\theta_{13}$	—	—	0.090	$+0.032, -0.029$
χ^2_{min}/d.o.f.	—	—	52.2/40	—

sided significance test.

15.5.2 Parameter correlations

Figure 15.5 shows contours corresponding to $\Delta\chi^2 = 1$ and $\Delta\chi^2 = 4$ for select pairs of fit parameters. These plots illustrate the low level of correlation between most parameters in the Rate+Shape fit. A mild positive correlation is evident for ($\sin^2 2\theta_{13}$, R_{Li}), which is intuitively reasonable: if a fixed-normalization spectrum has a larger share of Li+He, it must contain fewer signal events and thus match a larger $\sin^2 2\theta_{13}$. Likewise, the mild anti-correlation of ($\sin^2 2\theta_{13}$, Δm^2_{ee}) is sensible: a larger Δm^2_{ee} moves the oscillation maximum further below the prompt energy threshold, reducing the deficit visible in the prompt spectrum. The strongest relationship is an anti-correlation for ($\sin^2 2\theta_{13}$, b). This too is consistent with expectations, since an increase in b tends to shift lower-energy events into the oscillation region, which correlates to a smaller $\sin^2 2\theta_{13}$. Although increases in a or c also move events to higher energies, the low correlations for ($\sin^2 2\theta_{13}$, a) and ($\sin^2 2\theta_{13}$, c) suggest that these shifts deform the spectrum in a manner quite distinct from oscillation.

15.5.3 Prompt spectrum observations

Figures 15.6 and 15.7 show the prompt energy spectrum of IBD candidates, along with the no-oscillation and best-fit predictions. Figure 15.6 highlights the agreement between data and prediction in the background-dominated regions. In Fig. 15.7, the best-fit backgrounds are subtracted from both data and prediction, and the expanded vertical scale reveals an energy-dependent deficit of data compared to the no-oscillation prediction. This deficit is more obvious in Fig. 15.8, which shows the ratio of the background-subtracted data to the no-oscillation prediction. In the region of maximum oscillation, $\sim$ 1–3 MeV, this ratio is well described by θ_{13}-driven oscillation.

Above the main oscillation region, the data is not as well matched by the best-fit curve. An excess appears in the $\sim$ 4–6 MeV region, flanked by somewhat less significant deficits. These features are the primary reason for the relatively large χ^2_{min}/d.o.f. of the Rate+Shape fit, although tension with the Li+He pull term and reactors-off term also contribute, as discussed below. Importantly, the main data-MC discrepancies occur outside the region which is most sensitive to θ_{13}-driven oscillations. A Rate+Shape fit using only the most sensitive prompt energy region, $0.5 > E_{vis} > 4$ MeV, yields a central value of $\sin^2 2\theta_{13} = 0.087$, consistent with the standard result. In this restricted E_{vis} fit, the

$\sin^2 2\theta_{13}$ uncertainty interval is about 10% larger than the standard result due to weaker constraints on the backgrounds. Investigations of the non-θ_{13}-related spectrum features are the subject of Ch. 17.

15.5.4 Variations on background constraints

Variations on the standard fit were performed to assess the impact of two types of background constraints: the reactors-off data and the Li+He and FN+SM pull terms. If the reactors-off data is not included in the fit, the best-fit $\sin^2 2\theta_{13}$ barely changes, increasing by about 1% relative to the standard result. The $\sin^2 2\theta_{13}$ uncertainty interval is unaffected at the quoted precision. These results show that the reactors-off data has only a minor influence on the Rate+Shape fit. A contrasting effect is visible in the RRM fit (Ch. 16), which does not contain prompt spectrum information.

Removing the Li+He and FN+SM pull terms from the Rate+Shape fit, but keeping the associated pull parameters, is equivalent to assuming infinite *a priori* uncertainty on R_{Li} and R_{FN}. In this fit configuration, background constraints come only from the IBD candidates themselves. When this setup is tested in the Gd analysis, the best-fit FN+SM is nearly same as in the standard fit. The best-fit Li+He rate is about 20% lower than the standard result. This new Li+He result disagrees with $\overline{R_{Li}}$ at about the 2σ level, suggesting that prompt E_{vis} (plus reactors-off) information and ΔT_{μ} information favor different Li+He rates. Efforts are ongoing to understand whether this discrepancy results from statistical fluctuations or systematic bias. In the pull-term-free fit, the lower Li+He rate is correlated to a $\sim 5\%$ decrease in the best-fit for $\sin^2 2\theta_{13}$, relative to the standard result. Meanwhile, the uncertainty interval on $\sin^2 2\theta_{13}$ is only a few percent larger. Evidently, the strongest $\sin^2 2\theta_{13}$ constraint in the Rate+Shape fit is the prompt spectrum of the IBD candidates.

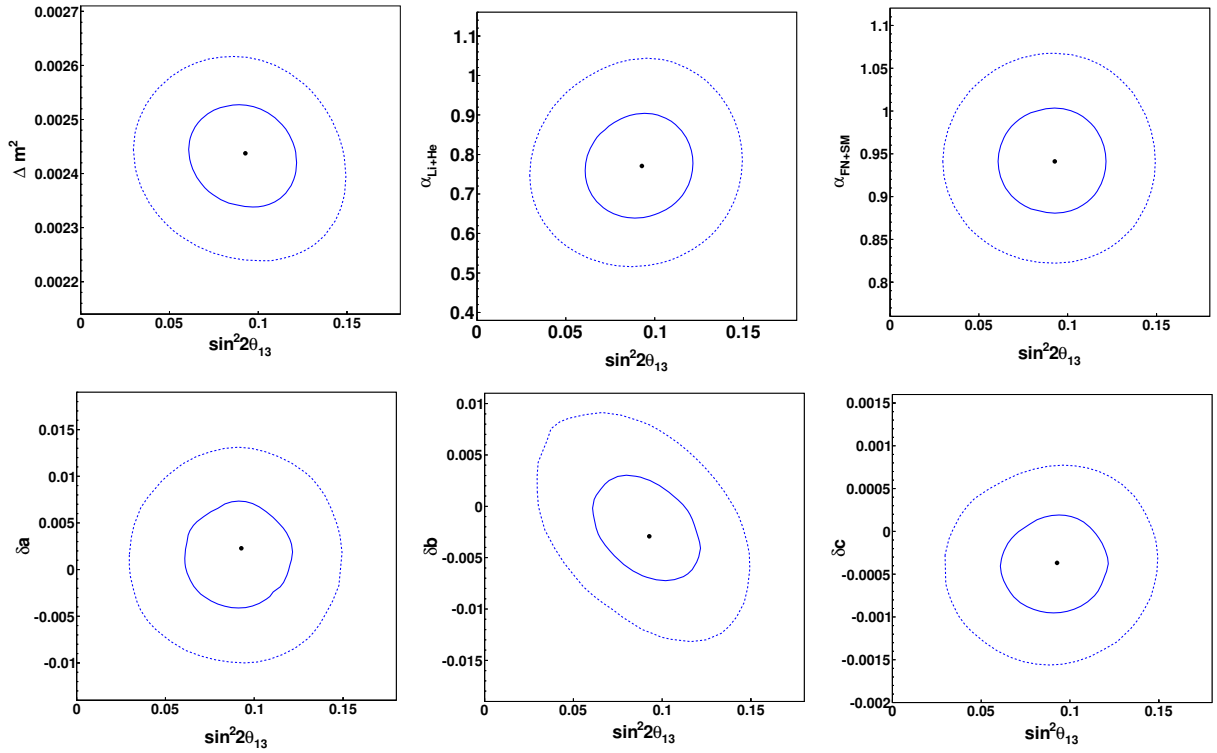

Figure 15.5: Contours corresponding to $\Delta\chi^2 = 1$ and $\Delta\chi^2 = 4$ for pairs of parameters in the Gd-channel Rate+Shape fit. The pairs contain $\sin^2 2\theta_{13}$ and, from left to right, top to bottom: Δm^2_{ee}, $R_{Li}/\overline{R_{Li}}$, $R_{FN}/\overline{R_{FN}}$, $(a-\overline{a})$, $(b-\overline{b})$, $(c-\overline{c})$. The projection of each region onto the horizontal and vertical axes gives the 1σ and 2σ uncertainty intervals for the corresponding parameters, according to the profile likelihood approach. Note that the contours do not represent the 68% and 95% confidence level regions for the pairs of parameters, taken together. Those regions would be drawn at $\Delta\chi^2 = 2.30$ and $\Delta\chi^2 = 6.18$, respectively.

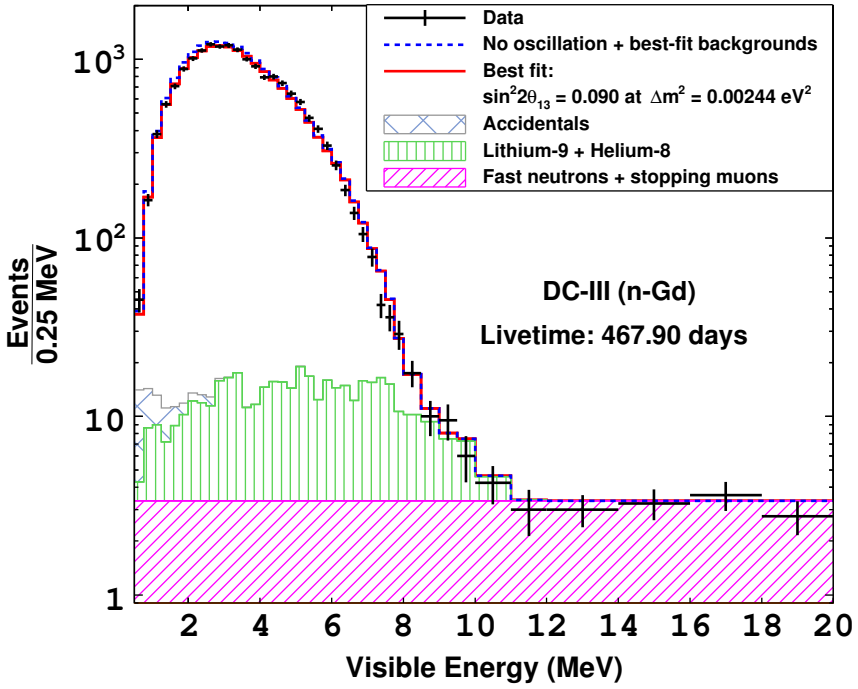

Figure 15.6: The prompt energy spectrum of Gd-channel IBD candidates (black points with statistical error bars), superimposed on the no-oscillation prediction (dashed blue line) and best fit (red line). The stacked histograms shows the background contributions from accidentals (blue crosshatched), Li+He (green vertically hatched), and FN+SM (magenta hatched).

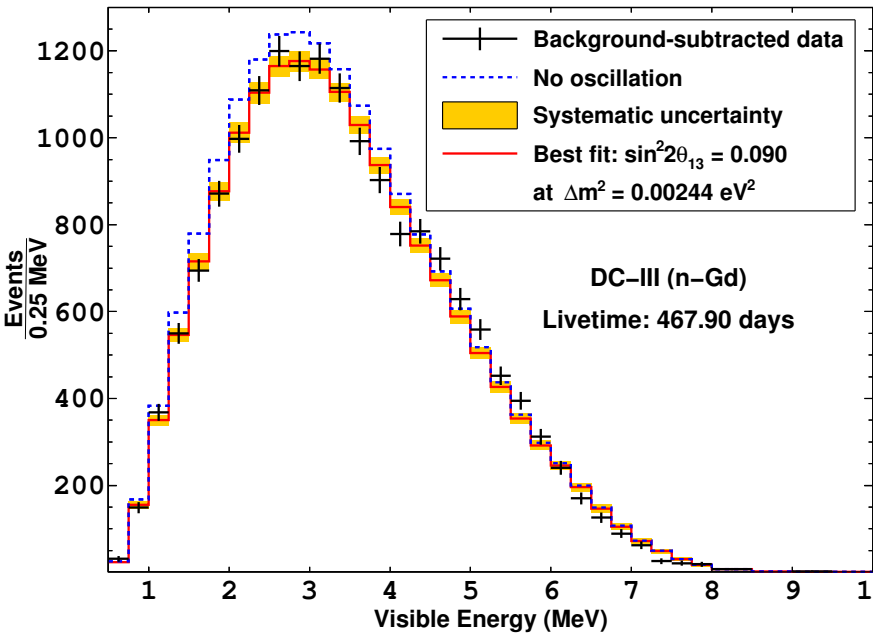

Figure 15.7: The background-subtracted prompt energy spectrum of Gd-channel IBD candidates (black points with statistical error bars), superimposed on the no-oscillation prediction (dashed blue line) and best fit (red line). Gold bands indicate systematic uncertainty in each bin.

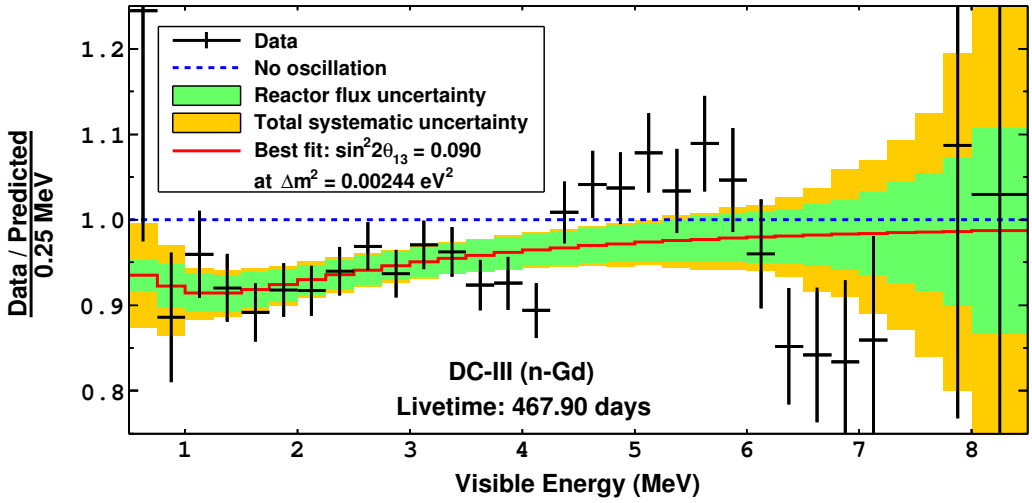

Figure 15.8: The ratio of background-subtracted Gd-channel data to the no-oscillation prediction (black points with statistical error bars), superimposed on corresponding ratios for the no-oscillation prediction (dashed blue line) and best fit (red line). Gold bands indicate the total systematic uncertainty in each bin; green bands indicate the systematic uncertainty from reactor flux modeling.

15.6 H-channel results

Figure 15.9 shows the prompt spectrum of H-channel IBD candidates. The Rate+Shape fit for this dataset produces a best estimate of $\sin^2 2\theta_{13} = 0.124^{+0.030}_{-0.039}$ at $\chi^2_{min} = 69.5/38$ d.o.f. By Eq. 15.11, this χ^2_{min}/d.o.f. corresponds to a p-value of 0.1%. Such a low p-value might be expected if the H-channel spectrum exhibited data-MC discrepancies similar to those in the > 4 MeV region of the Gd-channel spectrum, since the H analysis involves smaller statistical uncertainties. Indeed, features consistent with the Gd-channel spectrum are observed in the > 4 MeV portion of the H-channel prompt energy spectrum, as shown in Fig. 15.10. Figure 15.11 highlights the similarity of the Gd- and H-channel spectra.

As in the Gd analysis, the impact of the reactors-off data and background pull terms was examined for the H analysis. Again, the reactors-off data has minimal influence on the fit results. The Li+He pull term is somewhat more influential in the H analysis than in the Gd analysis: a fit pull-term-free H-channel fit yields a 20% higher $\sin^2 2\theta_{13}$ than the standard fit. Once again, the best-fit Li+He from this pull-term-free fit disagrees, albeit less significantly, with $\overline{R_{Li}}$. In a reversal of the Gd-channel comparison, the pull-term-free result is larger than $\overline{R_{Li}}$. Interestingly, these opposing tendencies are consistent with the intuition (advanced in Sec. 14.3.3) that the Li+He rate in the H analysis should be significantly larger than in the Gd analysis.

Figure 15.12 shows $\Delta\chi^2 = 1$ and $\Delta\chi^2 = 4$ contours for select pairs of parameters in the H analysis. The correlations are qualitatively similar to those in the Gd analysis.

Table 15.10: Results of the H-channel Rate+Shape fit. The first eight rows list pull parameters, whose definitions are given in the text.

Parameter	Input central value	Input error	Output best fit	Output error
a	0	0.067	-0.008	$+0.028, -0.020$
b	1.004	0.022	0.997	$+0.007, -0.009$
c	-0.0001	0.0006	-0.0006	$+0.0006, -0.0005$
R_{FN} (d^{-1})	1.55	0.15	1.62	0.10
R_{Li} (d^{-1})	0.95	$+0.57, -0.33$	1.60	$+0.22, -0.24$
R_{acc} (d^{-1})	4.334	0.011	4.334	0.011
μ_{res}	2.34	0.70	2.40	0.70
Δm^2_{ee} (10^{-3} eV^2)	2.44	$+0.09, -0.10$	2.44	$+0.09, -0.10$
$\sin^2 2\theta_{13}$	—	—	0.124	$+0.030, -0.039$
χ^2_{min}/d.o.f.	—	—	69.5/38	—

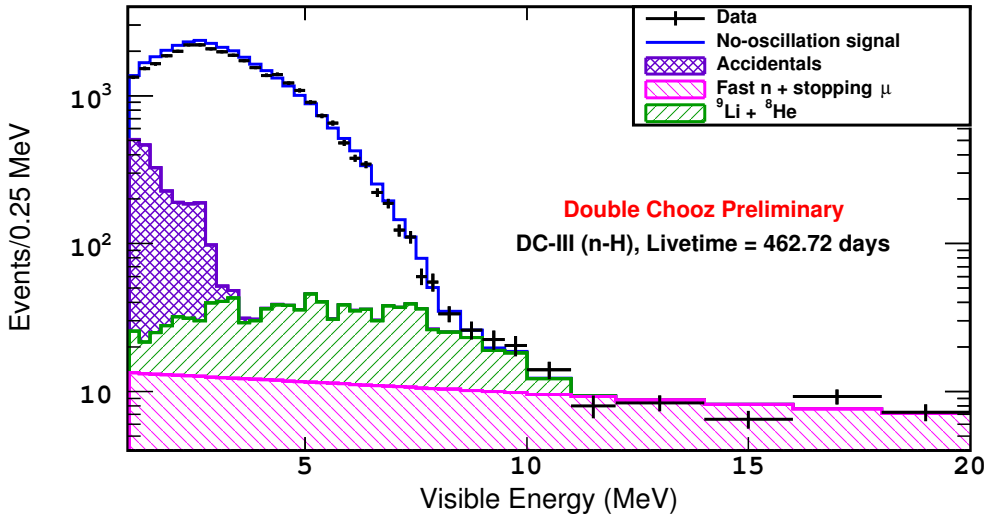

Figure 15.9: The prompt energy spectrum of H-channel IBD candidates (black points with statistical error bars), superimposed on the no-oscillation prediction (dashed blue line) and best fit (red line). The stacked histograms shows the background contributions from accidentals (purple crosshatched), Li+He (green hatched), and FN+SM (magenta hatched).

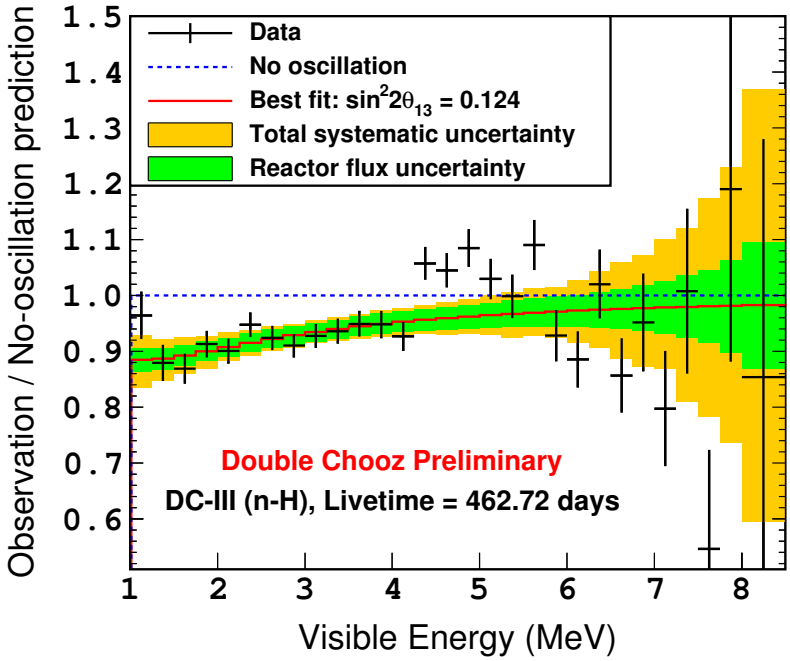

Figure 15.10: The ratio of background-subtracted H-channel data to the no-oscillation prediction (black points with statistical error bars), superimposed on corresponding rations for the no-oscillation prediction (dashed blue line) and best fit (red line). Gold bands indicate the total systematic uncertainty in each bin; green bands indicate the systematic uncertainty from reactor flux modeling.

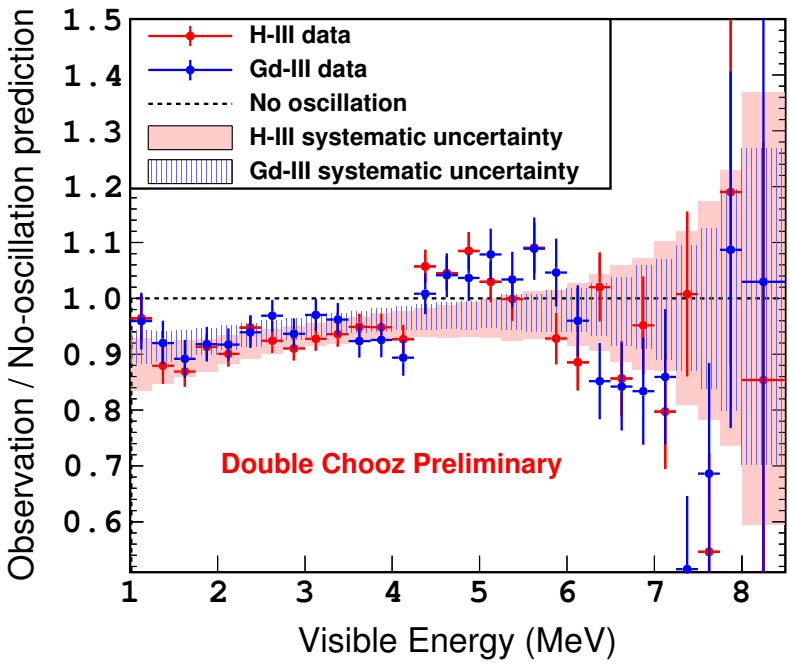

Figure 15.11: The ratio of background-subtracted Gd-channel (H-channel) data to the no-oscillation prediction, shown as blue (red) points with statistical error bars. The blue vertically hatched (light red) bands indicate total systematic uncertainty for the Gd (H) analysis. Note that the MC correction factor $C_{data/MC}$, best-fit energy scales, and best-fit background rates are determined separately in the Gd and H analyses. Differences in these parameters may cause systematic shifts between the data points in this plot.

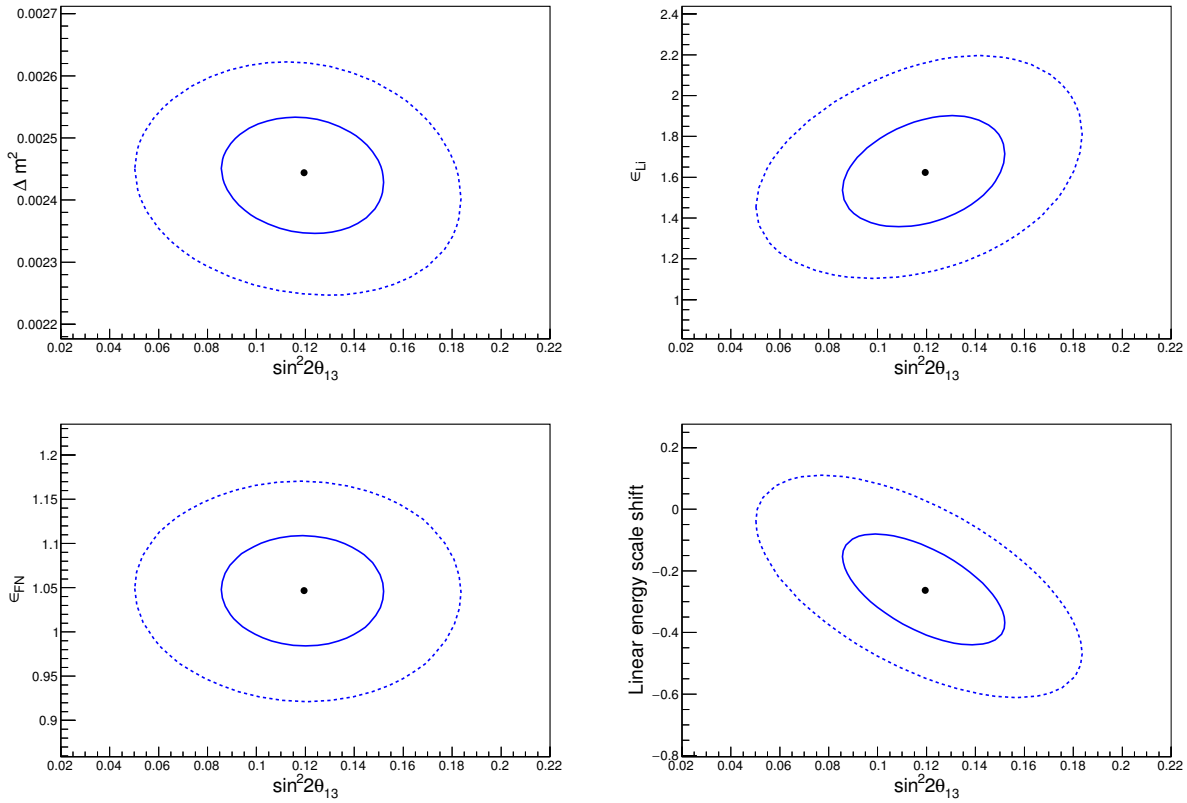

Figure 15.12: Contours corresponding to $\Delta\chi^2 = 1$ and $\Delta\chi^2 = 4$ for pairs of parameters in the H-channel Rate+Shape fit, using the simplified energy scale treatment. The pairs contain $\sin^2 2\theta_{13}$ and, from left to right, top to bottom: Δm^2_{ee}; $R_{Li}/\overline{R_{Li}}$; $R_{FN}/\overline{R_{FN}}$; k from Eq. 15.40, which is fully correlated with b. The projection of each region onto the horizontal and vertical axes gives the 1σ and 2σ uncertainty intervals for the corresponding parameters, according to the profile likelihood approach. As in Fig. 15.5, these contours do not represent the 68% and 95% confidence level regions for the pairs of parameters.

15.7 Perspective on a combined Gd+H fit

A combined Rate+Shape fit to the Gd and H datasets would be possible if correlations could be assessed between all sources of systematic uncertainty. The same analysis could quantify the compatibility between the individual Gd and H fits. Inter-analysis correlations are easy to understand for two sources of uncertainty. The dominant source, the reactor flux model, is certainly correlated between the analyses. Statistical uncertainties are certainly uncorrelated. Correlations for the remaining sources of uncertainty are more challenging to quantify. When such a reckoning was attempted for the previous Gd- and H-based analyses, estimating correlations for the detection efficiency, energy scale, and Li+He rate uncertainties proved especially difficult. In light of these complications, a combined Rate+Shape fit has not been undertaken for the current analyses. Because it does not involve spectrum shape information, the RRM technique is a simpler platform for a combined Gd+H fit. Results of a combined RRM fit are presented in the next chapter.

Chapter 16

Reactor Rate Modulation fits

While the Rate+Shape approach offers valuable sensitivity to $\sin^2 2\theta_{13}$, its reliance on prompt spectrum information has certain disadvantages. A model must be assumed for the types and prompt energy spectra of background populations. Careful assumptions must also be made about the shape of the signal spectrum, especially if it is provided by a reactor simulation. Compared to the Rate+Shape strategy, the Reactor Rate Modulation approach represents the opposite compromise: weaker $\sin^2 2\theta_{13}$ sensitivity but greater model independence. This chapter presents the concept for the RRM fits in Double Chooz, an outline of their implementation, and results for the Gd analysis, H analysis, and a Gd+H combination.

16.1 Concept

As remarked in Sec. 15.2.1, correlation with reactor power conditions can help to distinguish signal from background. Although this effect adds little leverage to a Rate+Shape fit, where powerful signal and background distinctions already exist in the prompt spectrum, it is important in a Rate-Only analysis. That principle is the basis of the RRM fit. In this approach, data is divided into periods of varying reactor power. In each period, the number of detected IBD candidates is compared to the number of signal events that would be expected if no oscillation occurred. An observed deficit that scales with reactor power is an indication of nonzero $\sin^2 2\theta_{13}$.

Focusing on the binning scheme of the RRM fit underscores its complementarity to the Rate+Shape strategy. In the Rate+Shape fit, events are integrated over all time and binned according to prompt

energy. In the RRM fit, events are integrated over all prompt energies but divided into multiple bins according to the time at which they occurred. For the Gd and H datasets, the essential division is between reactor(s)-on and reactors-off data. Division according to single-reactor versus double-reactor operation is a slight enhancement. Beyond that, the number of reactor power bins is mostly a matter of convention.

Integrating over prompt energy removes dependence on the signal spectrum shape. As a result, systematic uncertainties on that shape, namely energy scale and reactor flux shape uncertainties, do not factor into the RRM fit. The RRM strategy also obviates the need for any kind of background model, beyond the assumption that backgrounds do not scale with reactor power. The total background rate can simply be treated as a free parameter, constant across all reactor power bins. In this basic version of the RRM fit, the only sources of systematic uncertainty are the reactor flux normalization and detection efficiency.

An RRM-style fit can also be performed with the total background rate constrained by the sum of the rate estimates in Secs. 13.3 and 14.3. This constraint improves $\sin^2 2\theta_{13}$ precision at the cost of a background model assumption.

16.2 Definition of χ^2 statistic

16.2.1 Individual Gd and H fits

The χ^2 statistic for the RRM fit is, like the Rate+Shape statistic, a version of Eq. 15.13. The reactors-on data is divided into $b = 6$ bins, with three bins covering various levels of single-reactor operation and three covering various levels of double-reactor operation. One additional bin is included for reactors-off data. As in the Rate+Shape χ^2, the reactors-off bin receives the extended likelihood treatment of Eq. 15.10.

Across all bins, signal normalization is modulated by three pull parameters, each corresponding to a different source of systematic uncertainty. The parameter κ_{eff} accounts for detection efficiency effects and is constrained by the relative uncertainty on $C_{data/MC}$, as listed in Tab. 13.3 (Tab. 14.3) for the Gd (H) analysis. The parameter κ_{rxtr} covers normalization uncertainties from the reactor flux simulation. Reactor uncertainty varies somewhat across the reactor(s)-on bins, since, as noted in Sec. 12.1.1 the P_{th} measurements have larger uncertainties for lower values. Variations between

bins are treated with scaling factors for κ_{rxtr}. Each scaling factor is defined as $r_i = (\sigma_{rxtr})_i/\sigma_{rxtr}$, where $(\sigma_{rxtr})_i$ is the relative flux normalization uncertainty in the i^{th} reactor power period and $\sigma_{rxtr} = 0.0175$ is the uncertainty when both cores operate at full power. The scaling factor treatment assumes that reactor uncertainties are fully correlated between bins, consistent with the conservative approach taken in Sec. 15.3.1.

The third pull parameter, κ_{res}, accounts for uncertainty on the number of residual antineutrinos, as defined in Sec. 15.2.4. While the Rate+Shape fit neglects the extremely small contribution of residual antineutrinos in the reactor(s)-on periods, they are included in the RRM fit. The relative uncertainty on the residual antineutrino rate in the reactor(s)-on and reactors-off is $\sigma_{res} = 0.30$. In the reactor(s)-on periods, κ_{res} is scaled by factors of $s_i = (\mu_{res})_i/\mu_i$, where $(\mu_{res})_i$ is the number of residual antineutrino events predicted in the i^{th} bin and $\mu_i^{\nu,null}$ is the total number of antineutrino events predicted in the i^{th} bin in the absence of oscillations.

In a Rate-Only method such as the RRM fit, it is unnecessary to model oscillation on an event-by-event basis. Average oscillation factors ω_i defined in analogy to Eq. 15.17, provide a simpler and equivalent treatment. Statistical uncertainty in the RRM fit is treated according to the Neyman convention (Eq. 15.5).

The basic RRM fit includes $m = 2$ free parameters: $\sin^2 2\theta_{13}$ and the total background rate, B. The χ^2 can thus be written:

$$\chi^2(\sin^2 2\theta_{13}, B) = \sum_{i=0}^{b} \frac{\left(n_i - \mu_i(\sin^2 2\theta_{13}, B)\right)^2}{n_i} + \chi^2_{off} + \sum_{\alpha} \frac{(\kappa_\alpha - 1)^2}{\sigma_\alpha^2} \tag{16.1}$$

where α runs over reactor flux, detection efficiency, and residual antineutrinos. In this expression, the μ_i are given by:

$$\mu_i = \left(1 - \sin^2 2\theta_{13}\ \omega_i\right)\left(\kappa_{eff} + r_i\kappa_{rxtr} + s_i\kappa_{res}\right)\mu_i^{\nu,null} + BT_i \tag{16.2}$$

where T_i is the total livetime for data in the i^{th} bin. As in the Rate+Shape fit, Eq. 15.16 defines χ^2_{off}. In this case, $\boldsymbol{\eta}$ contains only B and μ_{off} is given by Eq. 16.2 with $r_{off} = 0$ and $s_{off} = 1$.

Independent background constraints can be included in the χ^2 by adding a pull term on B:

$$\frac{(B - \overline{B})^2}{\sigma_B^2} \tag{16.3}$$

where $\overline{B} = \overline{R_{acc}} + \overline{R_{FN}} + \overline{R_{Li}}$ and $\sigma_B^2 = \sigma_{acc}^2 + \sigma_{Li}^2 + \sigma_{FN}^2$, using the values in Tab. 15.6.

16.2.2 Combined Gd+H fit

If the Gd and H analyses were completely uncorrelated, a combined fit could be performed by minimizing the sum of the Gd and H χ^2 statistics, each defined by Eq. 16.1. A single $\sin^2 2\theta_{13}$ would be common to both fits, while separate B^{Gd} and B^H parameters would describe the background rate in each analysis. Separate σ_α^{Gd} and σ_α^H, with their own pull terms, would cover systematic uncertainties.

In reality, the reactor uncertainties are completely correlated between Gd and H, so a common σ_{rxtr}, with an associated pull term, covers the reactor flux uncertainty for both analyses. Similarly, the residual antineutrino uncertainties are considered fully correlated and are treated with a single σ_{res}. As noted in Sec. 15.7, it is difficult to precisely quantify the correlation between the Gd and H detection efficiency uncertainties. Consideration of each factor in $C_{data/MC}$ suggests that the maximum possible correlation is about 30%. To account for this correlation, the χ^2 can include separate σ_{eff}^{Gd} and σ_{eff}^H which are constrained by correlated pull terms.

If background constraints are included in the combined fit, their correlations must also be considered. The FN+SM rate uncertainties are dominated by the statistics of the IV-tagged and extended-range IBD events, which are independent samples in the Gd and H analyses, so the σ_{FN} are essentially uncorrelated. The systematic portion of the accidental rate uncertainty is probably somewhat correlated between Gd and H, but σ_{acc} is a small enough of fraction of σ_B that this correlation can be neglected. The remaining question is the correlation of the σ_{Li}. Scrutiny of all contributors to σ_{Li} indicates that the correlation is between zero and 50%. Variations on the combined Gd+H RRM fit are performed to test the impact of σ_B and σ_{eff} correlations.

16.3 Results

Results of the Gd, H, and combined Gd+H fits are listed in Tab. 16.1. The combined fits shown in this table assume no inter-analysis correlation of the detection efficiency and background rate uncertainties. Variations in which the maximum correlation is assumed produce nearly identical results, with best-fit parameters agreeing up to the percent level. Note that studies for the previous Gd- and H-based analyses indicated that correlations tend to have larger impacts on Rate+Shape fits.

Figure 16.1 displays selected $\Delta\chi^2$ projections for the Gd and H fits. These plots depict the significant correlation between $\sin^2 2\theta_{13}$ and B, which contrasts with the lower levels of parameter correlation evident in the Rate+Shape fit (see Figs. 15.5 and 15.12). Compatibility of the RRM and Rate+Shape $\sin^2 2\theta_{13}$ findings is not straightforward to quantify, due to subtleties of the inter-fit correlations. On face, the results are consistent, especially when background rate constraints are included in the RRM fits. The larger $\sin^2 2\theta_{13}$ uncertainty intervals in the RRM analysis underscore the observation that reactor power is a somewhat less informative dimension than prompt energy.

Figure 16.2 shows the combined best fit to the Gd and H datasets and illustrates the attractive simplicity of this measurement.

Table 16.1: Results of RRM fits to the Gd, H, and combined Gd+H datasets. Gd results are from [86]; results for H and the combination are from the forthcoming H publication.

Parameter	Gd best fit	H best fit	Gd+H best fit
	Background model–independent fit		
$\sin^2 2\theta_{13}$	0.060 ± 0.039	$0.123^{+0.042}_{-0.043}$	0.078 ± 0.037
B^{Gd} (d^{-1})	$0.93^{+0.43}_{-0.36}$	—	$1.24^{+0.43}_{-0.39}$
B^{H} (d^{-1})	—	8.28 ± 0.87	$7.45^{+0.70}_{-0.68}$
χ^2_{min}/d.o.f.	1.9/5	6.2/5	10/11
	Including background pull term(s)		
$\sin^2 2\theta_{13}$	$0.090^{+0.034}_{-0.035}$	$0.098^{+0.038}_{-0.039}$	0.090 ± 0.033
B^{Gd} (d^{-1})	$1.56^{+0.18}_{-0.16}$	—	$1.57^{+0.18}_{-0.16}$
B^{H} (d^{-1})	—	7.29 ± 0.49	7.22 ± 0.43
χ^2_{min}/d.o.f.	4.2/6	8.1/6	12/13

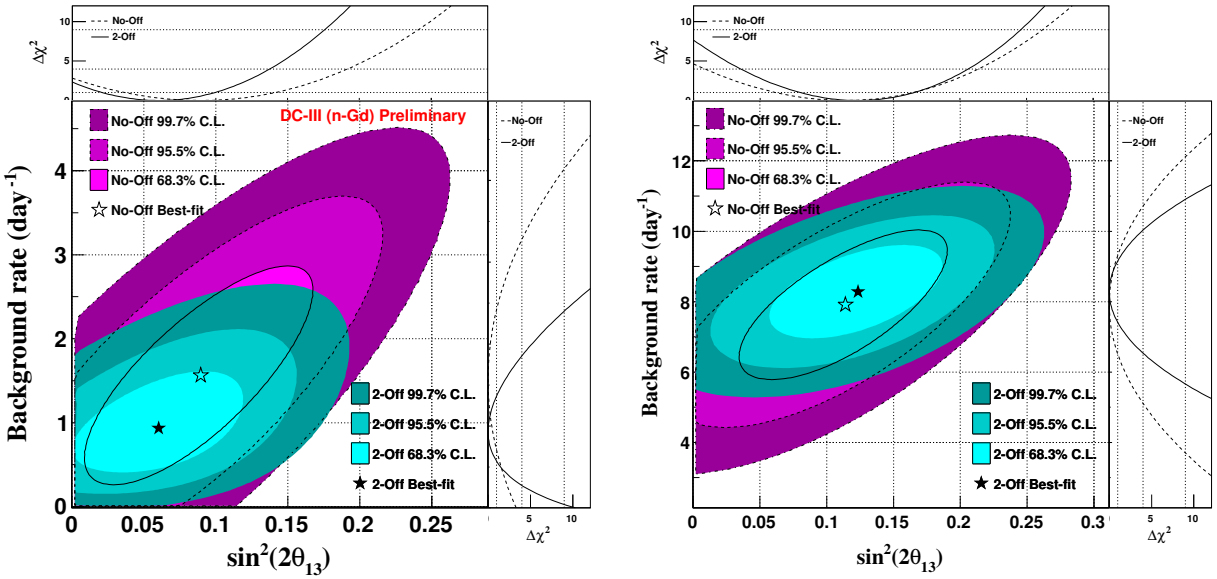

Figure 16.1: Confidence intervals for $\sin^2 2\theta_{13}$ and B in the background model–independent RRM fits to the Gd (left) and H (right) datasets. The central panes show the 68.3%, 95.5%, and 99.7% CL regions for the standard fit (shades of turquoise). The best fit point is indicated with a filled star. To show the impact of the reactors-off data point, analogous regions are also drawn for a fit which excludes that point (shades of magenta). The best fit for this variation is indicated with an open star. The smaller panes to the top (right) show $\Delta\chi^2$ as function of $\sin^2 2\theta_{13}$ (B), continuously minimized over the other fit parameters. According the profile likelihood prescription, 1σ, 2σ, and 3σ CL intervals on $\sin^2 2\theta_{13}$ and B are defined at $\Delta\chi^2 = 1$, 4, and 9, respectively (marked with dotted lines). Left figure from [86]; right figure from the forthcoming H publication.

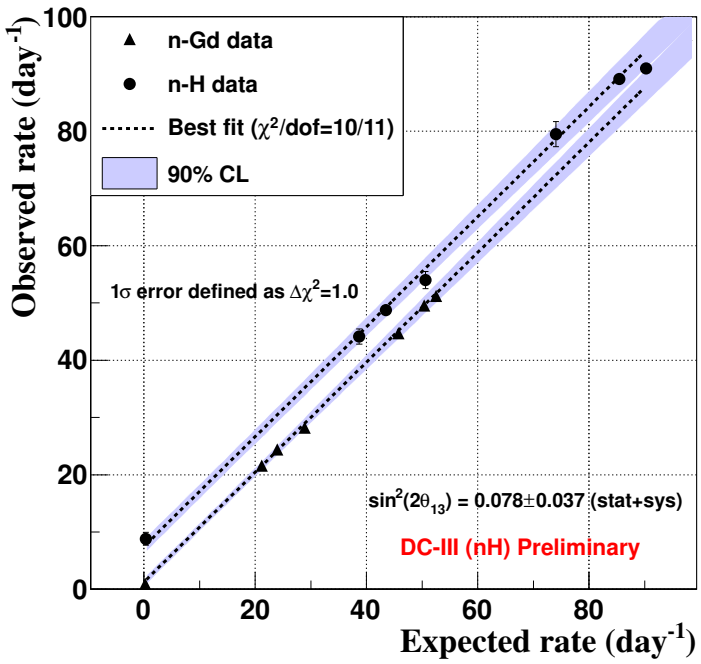

Figure 16.2: The observed rate of IBD candidates versus the no-oscillation expectation, for the Gd dataset (triangles, with statistical error bars too small to see) and the H dataset (circles, with barely visible statistical errors). Dashed lines indicate the best combined fit, with systematic uncertainty indicated by the blue region. The lines have a slope of slightly less than one, corresponding to nonzero $\sin^2 2\theta_{13}$. The vertical intercept of each line corresponds to the background rate.

Chapter 17

Reactor spectrum features

The consistency of Rate+Shape and RRM results reinforces the idea that the data-MC discrepancies observed above 4 MeV are mostly irrelevant to the Rate+Shape measurement of $\sin^2 2\theta_{13}$. Still, the presence of strong, similar features in both the Gd and H spectra demands some explanation. If the features arise from an instrumental effect in Double Chooz, such as energy dependence in signal detection efficiency or a poorly modeled energy scale, correcting this effect will strengthen future analyses. If the cause is more general, such as a new type of background or a bias in the reactor model, that revelation could have wider experimental and theoretical impact. This chapter summarizes examinations of multiple hypotheses in Double Chooz, along with related observations from other experiments. The later sections concentrate on the most likely explanation, the reactor spectrum model, and its relation to Double Chooz measurements.

17.1 Investigations in Double Chooz

A discrepancy between the data and MC energy scales can cause patterns of excess and deficit like those observed in Double Chooz. However, that phenomenon seems unlikely here because the Double Chooz MC energy scale, especially in the Gd analysis, matches the data very well (see Fig. 12.5). Spallation neutron captures on C provide a calibration point at 5 MeV, in the middle of the observed spectrum distortions. The peak energy from neutron captures on C is the same to within 0.5% in data and MC. The energy resolution also agrees extremely well at this point. These observations leave no room for a significant energy scale discrepancy.

Another explanation might be energy dependence in the signal detection efficiency, as modeled through $C_{data/MC}$. This hypothesis is difficult to motivate, since nearly all factors in $C_{data/MC}$ are clearly independent of prompt E_{vis}. A possible exception is the ANN cut efficiency, but for that cut, energy dependence was carefully studied and found to be insignificant. Moreover, the ANN cut is used only in the H analysis. Observations of the cosmogenic ^{12}B beta decay spectrum also undercut the detection efficiency and energy scale hypotheses. The beta spectrum of ^{12}B, which like ^{9}Li and ^{8}He is created by muon-initiated processes involving C nuclei, has a well known shape with a maximum around 7 MeV. The ^{12}B spectrum seen in Double Chooz is statistically limited, but comparisons with the MC spectrum show no major deviations.

The $\sim$ 4–6 MeV excess might be explained by a previously unknown background population. This possibility would not account for the deficits observed alongside the excess, but those might be the result of some secondary effect. Several new backgrounds were considered, including inelastic scattering of neutrons off C nuclei, neutron captures on Fe, and ^{13}C(α,n)^{16}O reactions (described in [124]). Studies confirm that the rates of these events should be vanishingly small in Double Chooz, and none of their prompt spectra match the observed excess.

The background hypothesis is also disfavored by two tests showing that the excess is correlated with reactor power. By the logic of Ch. 16, this correlation implies that the feature comes from reactor antineutrinos. One simple test, detailed in [86], focused only on the 4.25–6 MeV region. The excess in that region was shown to scale with the number of reactors operating, following the same trend as the bulk of the signal spectrum. Consistent correlations appear in the Gd and H selections.

A more sophisticated test uses a variation of the RRM fit. In addition to the usual reactor power bins, data is divided into five prompt energy bins. The binning scheme intentionally targets the observed excess and deficits. Separate κ_{rxtr} and B are defined for each bin. Three fits are performed, each with $\sin^2 2\theta_{13}$ fixed at a value measured in Daya Bay [125]. In the first fit, the κ_{rxtr} in each bin are constrained by the reactor model, while B is left free. As might be expected, the B move to fit the observed excesses and deficits, while the κ_{rxtr} remain close their input values. The second fit is the inverse of the first, with the κ_{rxtr} set free while the B are constrained by the background model. Again, as expected, the κ_{rxtr} move to fit the observed features while the B remain bound near their central values. The crucial distinction is that the χ^2_{min} of the second fit is significantly lower than that of the first. This difference indicates that changes to the reactor flux

model can better fit the data than changes to the background model. That is, the magnitude of the excess and deficits scales with reactor power, like signal, rather than remaining constant across reactor power bins, like a background.

If the spectrum features are indeed signal-related, and not explained by energy scale or efficiency effects, they might be the result of some previously unknown, nonstandard antineutrino interaction. However, there are no convincing theoretical models for such an interaction, and none have been observed in other reactor antineutrino experiments. The undulating structure of the spectrum distortion almost invites theories about nonstandard neutrino oscillations, but these could not explain a significant excess over the expectation. Furthermore, a suitable Δm^2, which would be of order ~ 0.01 eV^2, is inconsistent both with standard three-neutrino oscillations and the anomalous results discussed in Ch. 5.

A remaining hypothesis is inaccuracy in the reactor spectrum modeling. This explanation is appealing because the S_k constructions described in Sec. 12.1.3 are acknowledged to be difficult, and revisions continue to be developed [46, 96]. Problems could exist in the original beta spectrum measurements, in the conversion from beta to antineutrino spectra, or both. This hypothesis would explain why the features are essentially the same in Gd and H datasets, since both are compared to the same reactor model. It could also explain the curious affinity between the data-MC ratios in Double Chooz (Fig. 15.11) and CHOOZ (Fig. 3.5), as both experiments used similar S_k.

17.2 Observations in other experiments

Soon after Double Chooz reported anomalous features in the Gd-channel spectrum, RENO and Daya Bay announced similar findings in their near and far detectors [126, 127]. Although details of the observed shapes vary slightly, all three experiments share a clear excess in the $\sim$ 4–6 MeV region, compared to the predicted reactor antineutrino spectrum. These experiments have different readout electronics, different reconstruction and selection algorithms, and distinct background compositions due to their different detector depths. However, all use the same antineutrino flux predictions, based on [96]. Consequently, the most compelling explanation for the excess lies in these predictions.

17.3 Reactor model studies

Recent work in the neutrino, nuclear, and reactor physics communities has investigated why the predicted antineutrino flux does not match the observed spectrum shape. One study includes a new implementation of the summation method introduced in Sec. 12.1.3 [128]. This study finds an excess in the $\sim$ 4–6 MeV region, compared to the conversion method spectra, but is not explicit about the issues which may have led to this discrepancy.

A later study points out that a $\sim$ 4–6 MeV excess appears if nuclear data from one canonical database is used, but not if an alternative database is used [129]. This work gives detailed attention to several possible causes of the excess: antineutrino flux from processes other than fissions in reactor cores, which are not included in the predicted spectra; aspects of forbidden beta transitions that were not treated properly in conversion schemes; improper estimations of the ^{238}U contribution to the total spectrum; the harder neutron spectrum in commercial PWRs, compared to the extremely thermal spectrum at ILL; and systematic problems in the original ILL measurements.

No definitive conclusion has been reached. New measurements from reactor antineutrino experiments, including the Double Chooz analysis anticipated in Ch. 17, will continue to shape these investigations. Proposals have also been made to repeat the ILL beta spectrum measurements, although this prospect is more remote. Beyond benefiting the current generation of reactor antineutrino experiments, deeper insights into the reactor spectrum may advance fundamental nuclear physics, searches for sterile neutrinos, and perhaps the nonproliferation applications mentioned in Sec. 19.2. Enhanced knowledge of fission product decay schemes may also have practical value to commercial reactor operators.

17.4 Robustness of Double Chooz $\sin^2 2\theta_{13}$ measurements

The new spectrum features have been an intriguing surprise for Double Chooz. Their discovery would be less exciting if it undermined the main physics measurements of the experiment. Fortunately, the Double Chooz capability to measure $\sin^2 2\theta_{13}$, even in a single-detector configuration, seems to be minimally affected by the apparent issues with the reactor spectrum prediction.

The Rate+Shape fit appears to be protected by the fact that the spectrum features occur at higher energies than the main oscillation signal. Still, features above 4 MeV could bias the fits for

the background populations, which in turn could bias the result for $\sin^2 2\theta_{13}$. One assurance against this possibility is noted in Sec. 15.5.3: a fit to the prompt energy region below 4 MeV is consistent with a fit to the entire prompt spectrum. Further support comes from a set of Rate+Shape tests in which a new event population was included in the Gd-III predicted spectrum. This hypothetical population, which could be a new background or an additional component of the signal spectrum, was modeled with a Gaussian spectrum shape centered around 5 MeV (with variations from 4.6 MeV to 5.6 MeV tested), with a width of 0.5 MeV (with variations from 0.2 MeV to 0.7 MeV tested). The normalization was left free. Predictably, fits with this new component found better χ^2_{min} than the standard fit, but the impact on other parameters was not large. At most, the best fit $\sin^2 2\theta_{13}$ changed by about 0.3σ. These tests are somewhat contrived, but they suggest that a more realistic spectrum prediction above 4 MeV would not greatly impact the Rate+Shape results.

The RRM results also provide strong corroboration for the Rate+Shape results. As emphasized in Ch. 16, the RRM fit is completely insensitive to the shape of the prompt signal spectrum. The fit does depend on the reactor flux normalization, but in Double Chooz, this value is provided by Bugey4 (Sec. 12.1.5). The Bugey4 anchor is the integral of an observed IBD spectrum, so it is essentially unaffected by reactor modeling problems. Even in the unlikely event that the spectrum feature are not due to reactor modeling, the basic RRM fit is still insulated from bias. Issues with detection efficiency or energy scale would not affect the RRM measurements, and an unknown background would simply be counted in B.

The preceding arguments suggest that current Rate+Shape results are insensitive to the unexpected spectrum features. However, certain qualifications are in order. Assuming the reactor hypothesis is correct, the reactor flux uncertainties have been substantially underestimated, at least above 4 MeV. Inaccuracy in that region naturally raises questions about the uncertainty budget below 4 MeV. It is possible that the $\sin^2 2\theta_{13}$ confidence regions given in Secs. 15.5 and 15.6 are somewhat optimistic. At present, this concern cannot be systematically addressed, as no comprehensive alternatives exist to the S_k estimations and uncertainties used thus far.

Data from the near detector will soon improve this situation. Since near detector background rates and oscillation effects are distinct from the far detector, the appearance of familiar features in the near detector data will solidify the reactor model explanation. Moreover, in a two-detector fit for $\sin^2 2\theta_{13}$, reliance on the reactor model is automatically diminished. Following the method

in Sec. 18.4, dependence can be entirely removed.

Part IV

Prospects for two-detector analyses

Chapter 18

Two-detector $\sin^2 2\theta_{13}$ precision

Figures 15.8 and 15.10 spotlight the dominance of reactor flux uncertainties in analyses that use only the far detector (and as discussed in Ch. 17, the uncertainty estimates in those plots may not even be as large as warranted). Since reactor-related uncertainty is irreducible in a single-detector analysis, the current Gd-channel results are close to the best precision attainable in that configuration. Of course, as laid out in Sec. 6.4, the full design of Double Chooz was chosen to overcome exactly this limitation. Double Chooz expects to produce its first two-detector $\sin^2 2\theta_{13}$ analysis within the next year. This chapter presents an MC-based projection the precision of such an analysis. The projection uses a Gd-channel Rate+Shape analysis and assumes backgrounds and systematics similar to the current Gd analysis. The impact of including H-channel data and the possibility of reactor model–independent fit are also briefly discussed.

18.1 Framework of study

While comparing far detector (FD) data directly to near detector (ND) data is an intuitive approach to a two-detector fit, it is also possibles to compare both FD and ND data to their respective predictions, accounting for correlations in the reactor model and other systematics. In principle, these two approaches can be made equivalent, but the second is more readily implementable in the existing Rate+Shape framework. This study constructs predicted spectra for the ND and FD, builds a joint Rate+Shape χ^2 statistic in which $\mathbf{n} = \boldsymbol{\mu}$, and estimates the 1σ uncertainty on $\sin^2 2\theta_{13}$ by evaluating the $\Delta\chi^2 = 1$ interval around the simulated value.

The two-detector χ^2 statistic is defined as an extension of the single-detector χ^2 in Eq. 15.14. In the two-detector case, the number of prompt energy bins in $\mathbf{n}$ and $\boldsymbol{\mu}$, and the number of rows and columns in $\mathbf{M}$, is increased to cover both the ND and FD spectra. This study uses the binning scheme of the current Gd analysis, so that $b = 2 \times 40 = 80$. The reactors-off point is not included here, since it has minimal impact on Rate+Shape precision.

This study includes pull parameters for the accidental, Li+He, and FN+SM rates in each detector. The corresponding pull terms are treated as uncorrelated. This treatment is justified by the same arguments given in Sec. 16.2.2. Some correlation may be realistic for the Li+He rate, but assuming minimal correlations between the near and far detectors is a conservative approach.

As in the Gd analysis, the covariance matrix $\mathbf{M}$ covers uncertainties from finite data statistics, reactor and detection efficiency, and the spectrum shapes for the Li+He and accidental backgrounds. The matrix takes the form a 2×2 block matrix, with the diagonal blocks representing uncertainties in the individual detectors and off-diagonal blocks representing correlations. Figure 18.1 illustrates this structure. The construction of each matrix is described in the following sections.

18.2 Inputs

18.2.1 Antineutrino signal

The exact normalization and shape of the signal spectrum for future two-detector analyses depend on future reactor dynamics. They also depend on the the efficiency of selection cuts, which will likely progress beyond the current Gd analysis. Until these details are known, the current Gd-channel prediction is a reasonable starting point. In the study detailed in this chapter, both the FD and ND spectra are constructed from the current signal MC with the Gd-channel selection applied, according to the following prescriptions.

One year of FD signal: Since the Gd-channel MC sample includes 100× the expected signal, the sample is scaled the MC by a factor of 1/100. The Gd-channel MC correction factor is also applied. Next, to produce a signal sample corresponding to one year of data-taking, the sample is scaled by a factor of (1 year)/(data-taking time of current analysis) $\approx$ (1 year)/(1.76 years) = 0.57. Here, data-taking time is the total time elapsed between the start and end of the Gd-channel dataset, which is longer than the analysis runtime due to calibration campaigns and detector maintenance.

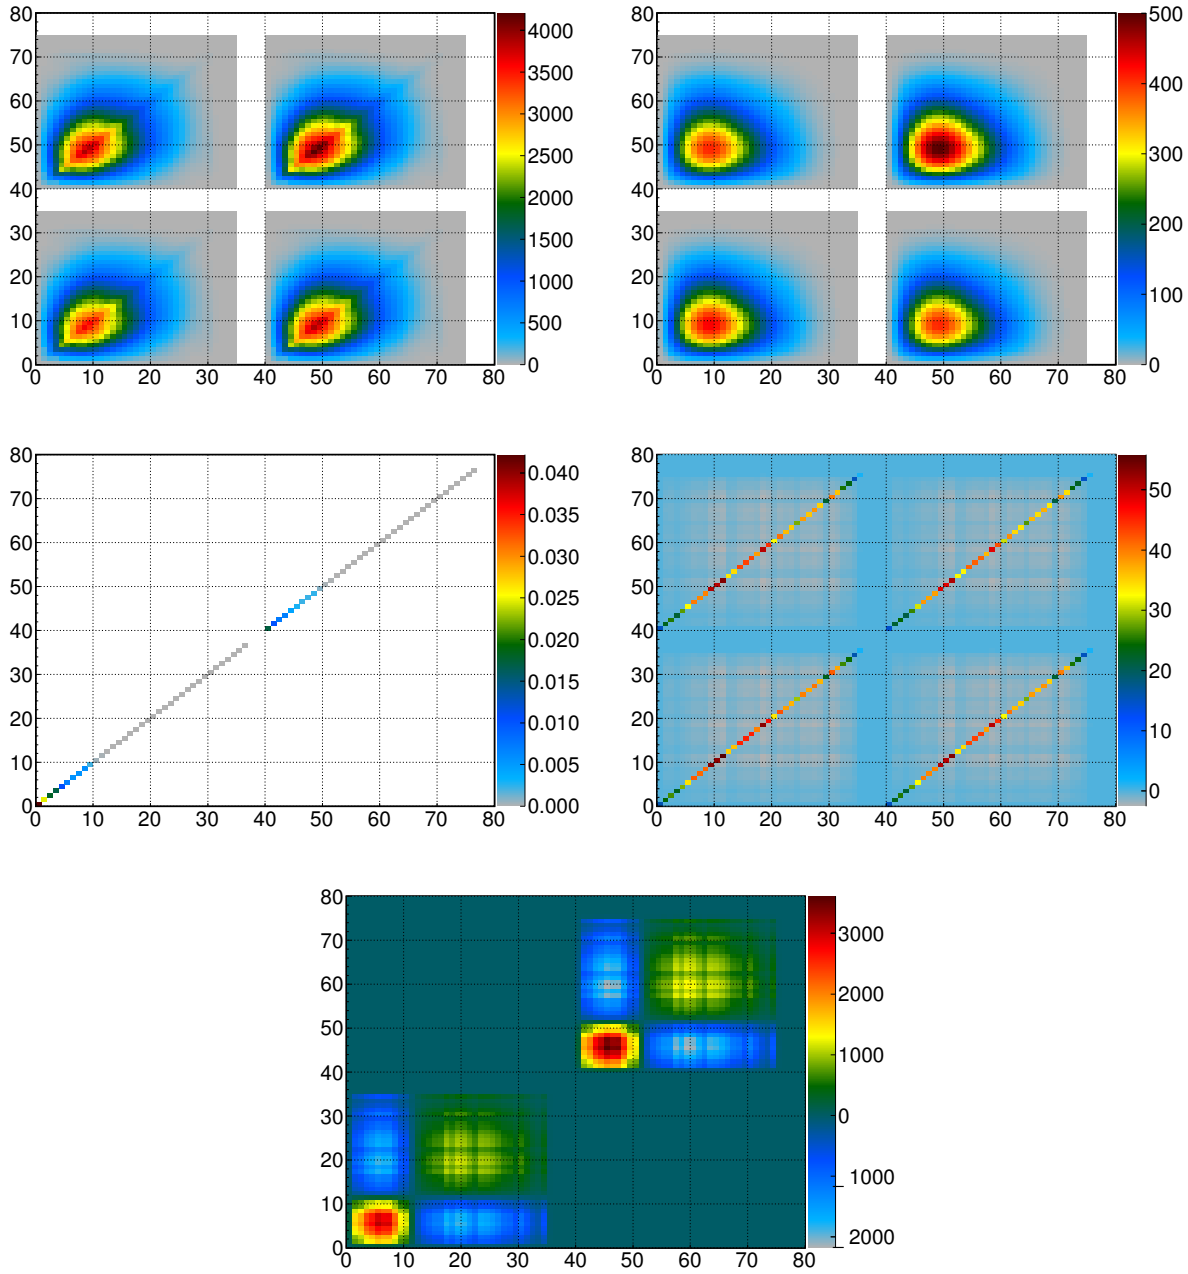

Figure 18.1: Systematic covariance matrices for the two-detector $\sin^2 2\theta_{13}$ sensitivity projection. The matrices represent uncertainties on the reactor flux model (top left), signal detection efficiency (top right), accidental spectrum shape (middle left), Li+He spectrum shape (middle right), and energy scale (bottom). In each matrix, the lower left (upper right) block covers the FD (ND) bins. The matrices are weighted by the nominal, no-oscillation prediction for one year of combined ND+FD data-taking, plus all of the previously taken FD data.

This scaling assumes that future FD data-taking will collect the same amount of signal events per unit time as the previous period of data-taking. That is, the average reactor flux, detector uptime, and selection efficiency are presumed to remain the same. This treatment also assumes that the current Gd MC is a good approximation of the IBD spectrum shape, ignoring the features discussed in Ch. 17. This treatment can be justified in multiple scenarios. By the time of a two-detector analysis in Double Chooz, more accurate reactor spectra may be available. The current spectrum is an acceptable stand-in, assuming the new spectrum shape uncertainties are similar in magnitude and the central value changes are not extreme. Even if new spectra are not available, the ND spectrum quickly becomes a more powerful constraint on the FD spectrum than the reactor model, so that the ultimate $\sin^2 2\theta_{13}$ precision is essentially independent of that model.

One year of ND signal: Like the FD sample, the ND signal sample begins with the current Gd-channel MC. In this case, the baseline of each event is adjusted to match the ND-reactor distances in Tab. 7.1. Specifically, the baseline of each event is shortened by 650 m (647 m) if the antineutrino originated in reactor B1 (B2). This step allows the small amount of oscillation in the ND to be factored into the weight of each event. Next, the weight of each event is increased to account for the larger solid angle of the ND with respect to the reactors. For events originating in B1 (B2), the scaling factor is $L^2_{\text{FD-B1}}/L^2_{\text{ND-B1}} = 5.75$ ($L^2_{\text{FD-B2}}/L^2_{\text{ND-B2}} = 8.08$), where L_{d-r} is the distance between the d^{th} detector and r^{th} reactor.

Next, the event weights are adjusted to account for increased muon veto deadtime expected at the shallower ND site. Measurements with a small OV prototype indicate that the cosmic ray muon flux is approximately 6.7× higher at the ND than at the FD. For the current Gd analysis, the muon veto caused a roughly 4.4% IBD signal inefficiency. For the ND sample, this factor can be scaled by ND-to-FD muon flux ratio, yielding 29% signal inefficiency. The ratio of livetime in the ND sample, compared to the current Gd dataset, is thus $(1 - 0.29)/(1 - 0.044) = 0.74$. Finally, the same statistical scaling, MC correction, and time scaling factors described for the FD sample are applied.

Table 18.1 summarizes the construction of the FD and ND signal samples, and Tab. 18.2 gives the number of predicted signal events in one year of data-taking.

Table 18.1: Summary of construction of one year of FD and ND signal. Each sample is constructed from current Gd-channel signal MC (with baselines shortened for the ND, as indicated by an asterisk) and the following scaling factors: standard MC statistics scaling, MC correction factor ($C_{data/MC}$), scaling for different reactor-detector baselines, scaling to give equivalent of one year of data-taking, scaling for different muon veto deadtime. The total scaling factor is the product of these five factors.

		Scaling factor					
Signal type	MC sample	MC stat.	$C_{data/MC}$	Solid angle	1 year	Muon veto	Total
FD	current Gd	0.01	0.92	none	0.57	none	0.005
ND from B1	" from B1*	0.01	0.92	5.75	0.57	0.74	0.022
ND from B2	" from B2*	0.01	0.92	8.08	0.57	0.74	0.031

Table 18.2: Number of signal events predicted for one year of ND and FD data-taking. Note that the fraction of events from B1 versus B2 depends on the operation of the reactors as well as the reactor-detector distances. These counts are based on the reactor cycles in the current analysis, in which B1 produced more antineutrinos than B2.

Signal type	Events from B1	Events from B2	Total signal events in one year
ND	21,760	28,722	50,482
FD	5,131	4,814	9,945

18.2.2 Backgrounds

For the FD spectrum, background rates and uncertainties are taken from the current Gd analysis. A symmetrized error is used for the Li+He rate. The accidental and FN+SM shapes are likely to differ somewhat due to the different muon profile, shielding, and detector surroundings at the ND, but the analysis impact should be small. The ND background rates are scaled from the FD rates according to the larger muon flux at the ND, as described below. The same relative uncertainties are used for the ND rates as for the FD.

Li+He: A relationship between Li+He, cosmogenic muon flux (Φ_μ), and average muon energy ($\langle E_\mu \rangle$) at a depth h is given in [119] as:

$$R_{Li}(h) \propto \Phi_\mu(h) \langle E_\mu(h) \rangle^{0.84} \tag{18.1}$$

Equation 18.1 implies the following relationship between the ND and FD rates of Li+He, denoted R_{Li}^{ND} and R_{Li}^{FD}:

$$R_{Li}^{ND} = R_{Li}^{FD} \frac{\Phi_\mu(h^{ND})}{\Phi_\mu(h^{FD})} \left(\frac{\langle E_\mu(h^{ND}) \rangle}{\langle E_\mu(h^{FD}) \rangle} \right)^{0.84} \tag{18.2}$$

Using the OV prototype measurements of $\Phi_\mu(h^{FD}) = 0.54$ m^{-2} s^{-1} and $\Phi_\mu(h^{ND}) = 3.60$ m^{-2} s^{-1}, along with estimates of $\langle E_\mu(h^{ND}) \rangle = 39.7$ GeV and $\langle E_\mu(h^{FD}) \rangle = 63.7$ GeV from [119], results in $R_{Li}^{ND} \approx 4.5 \times R_{Li}^{FD}$.

Fast neutrons: An analysis based on ΔT_μ indicates that about 70% of the FN+SM rate in the Gd analysis consists of fast neutrons. That gives a Gd-channel fast neutron rate of 0.42 d^{-1}. For the ND, this rate can be scaled following the same procedure as for Li+He. In this case, the exponent in Eq. 18.1 becomes 0.74 [119]. The result is $R_{FN}^{ND} \approx 4.7 \times R_{FN}^{FD}$.

Stopping muons: The remaining 30% of the Gd-channel FN+SM rate must be stopping muons, at a rate of 0.18 d^{-1}. No scaling relation for stopping muons is given in [119]. The rate should scale directly with muon flux, with perhaps some inverse relationship with average muon energy. Conservatively neglecting the latter effect yields $R_{SM}^{ND} \approx 6.7 \times R_{SM}^{FD}$.

Accidentals: The ND accidental rate depends on two factors. First, the rate of ambient radioactivity in the ND will largely determine the rate of prompt signals. This rate will differ from the FD in a way that is difficult to predict. Second, the rate of cosmogenic ^{12}B decays and spallation neutrons will make an important contribution to the rate of delayed signals. The rates of these cosmogenic events will scale with muon flux and energy, much like the Li+He and FN+SM rates.

Table 18.3: Background rates and spectrum shapes for the two-detector $\sin^2 2\theta_{13}$ sensitivity projection. As described in the text, FD rates are from the current Gd analysis (with uncertainties from that analysis) and ND rates are scaled from the current Gd analysis based on muon flux considerations (with relative uncertainties the same as in the current Gd analysis). Shapes and shape uncertainties for both ND and FD are from the current Gd analysis.

Background type	FD rate (d^{-1})	ND rate (d^{-1})	Spectrum shape	Shape unc.
Li+He	0.97 ± 0.29	4.37 ± 1.31	current Gd	current Gd
FN+SM	0.60 ± 0.05	3.18 ± 0.27	current Gd	current Gd
Accidentals	0.070 ± 0.005	0.21 ± 0.02	current Gd	current Gd

Overall, the rate of accidentals in the ND will have some relation to the FD-to-ND muon flux ratio, but the relationship will be weaker than for the purely cosmogenic backgrounds. A simple approximation is $R_{acc}^{FD} \approx 3 \times R_{acc}^{FD}$.

The spectrum shapes for all ND and FD backgrounds are taken from the current Gd analysis. In the two-detector phase, a higher-statistics Li+He spectrum can be measured by combining ND and FD data. This strategy implies full correlation of the Li+He shape uncertainty, as shown in the off-diagonal blocks of $\mathbf{M}^{Li}$ (Fig. 18.1). The accidental shape uncertainty is treated as uncorrelated, since it will be measured independently in each detector. As in the current Gd analysis, no shape uncertainty is included for the FN+SM spectrum.

Table 18.3 summarizes the background inputs for this study.

18.2.3 Detection efficiency uncertainty

This study assumes that the $C_{data/MC}$ uncertainty in each detector is 0.62%. That figure is based on the Gd-channel uncertainties for selection cut efficiencies, the Gd fraction, the spill current, and the proton number. The first three factors contribute 0.54%, while N_p contributes 0.30%. For simplicity, this study neglects the small contributions from remaining factors in the Gd-channel $C_{data/MC}$.

The N_p uncertainty is dominated by uncertainty on the scintillator composition, which is common to the ND and FD (see Sec. 13.2.2). This uncertainty must be correlated between the detectors. Correlations in the remaining 0.54% are less clear. A preliminary estimate, based on the amount

of Gd-channel uncertainty arising from data and MC statistics, suggests that about 0.20% may be uncorrelated between detectors. Under this assumption, the total correlated portion of the detection efficiency uncertainty is 0.58%.

18.2.4 Energy scale uncertainty

To speed computation, energy scale uncertainty is taken from the simplified treatment of Sec. 15.3.5. The Gd-channel value of $\sigma_b^{eff} = 0.007$ is used for each detector. As a faster substitute for the pull parameter treatment of Sec. 15.4.2, a covariance matrix is built to model these uncertainties. Matrix construction proceeds through the same multisim technique as the $(\Delta\boldsymbol{\mu}^{\nu})_{ms}$ derivation described in Sec. 15.4.2. Energy scale uncertainty is likely to be partially correlated between the ND and FD, since it will be determined through similar calibration campaigns. Once again, it is difficult to quantify the magnitude of this correlation. This study makes the conservative assumption of no correlation for energy scale uncertainties. The resulting $\mathbf{M}^{energy}$ is pictured in Fig. 18.1.

18.2.5 Reactor flux uncertainty

As noted, this study assumes that the current reactor model is adequate for estimating $\sin^2 2\theta_{13}$ sensitivity. Consequently, the reactor uncertainty in each detector can be reasonably described by $\mathbf{M}^{rxtr}$ from the current Gd analysis. Now, the question is the degree of correlation between the two detectors. Two scenarios would lead to full correlation. The flux uncertainties could be completely correlated between the reactors, in which case the ND and FD budgets would be fully correlated regardless of where they were positioned or when they took data. Alternatively, the detectors could take data at the same time and in isoflux positions, a concept introduced in Sec. 7.2.

Both of these conditions are partially, but not entirely, met in Double Chooz. The single-detector analyses assume that all reactor-related uncertainties are fully correlated between reactors. As explained in Sec. 15.3.1, this assumption is well justified for all major components of the reactor uncertainty budget except P_{th} and possibly the α_k. By these arguments, about 1.5% of the reactor flux uncertainty can safely be taken as correlated. The Double Chooz site does not have a perfectly isoflux geometry, but as depicted in Fig. 7.2, it is not too far off. A recent study by Double Chooz collaborators estimates that the site layout alone leads to $\sim 90\%$ correlation of reactor uncertainties. Taken together, the reactor model correlations and geometric effects suggest that at least $\sim 97\%$

of the reactor uncertainty is correlated between reactors. A 97% correlation, constant across the prompt energy bins, is assumed in this study.

18.2.6 Mass splitting

The mass splitting is constrained with the adjusted MINOS value of $\Delta m^2_{ee} = 2.44^{+0.09}_{-0.10} \times 10^{-3}$ eV^2, introduced in Sec. 15.3.3.

18.3 Results

Figure 18.2 shows the estimated Gd-only, Rate+Shape precision as a function of time since the start of FD operation. The systematic limitation of the single-detector analysis is clearly visible. On its own, five more years of far detector data would decrease $\sin^2 2\theta_{13}$ uncertainty by barely 15%. When near detector data is included, $\sin^2 2\theta_{13}$ precision eclipses that mark within just a few months. After four years of combined ND+FD running, according the assumptions made in this study, the $\sin^2 2\theta_{13}$ uncertainty interval will be about half the size of the current Gd-based interval. Since this study uses rather conservative assumptions, the actual uncertainty may well be smaller.

At this level of precision, Double Chooz will become a key check on the world-leading $\sin^2 2\theta_{13}$ measurement from Daya Bay. Due to its shallower detector locations, lower detector mass, later ND start-up, and smaller number of reactors, Double Chooz will probably always lag Daya Bay in absolute precision. However, with increasing significance, it will be possible to see whether the Double Chooz measurement of $\sin^2 2\theta_{13}$ is compatible with results from Daya Bay. Tension between these measurements could be a valuable sign of systematic bias in one or more of the analyses.

18.4 Possible enhancements

The two-detector analysis in this chapter reflects the achievement of latest Gd analysis, which is already one year old. In the past year, new techniques for background rejection and systematic error reduction have been explored. Some of these are used in the H analysis, while others are still under development. By the completion of the first two-detector analysis, these techniques are likely to decrease some backgrounds and improve the detection efficiency uncertainty beyond the levels

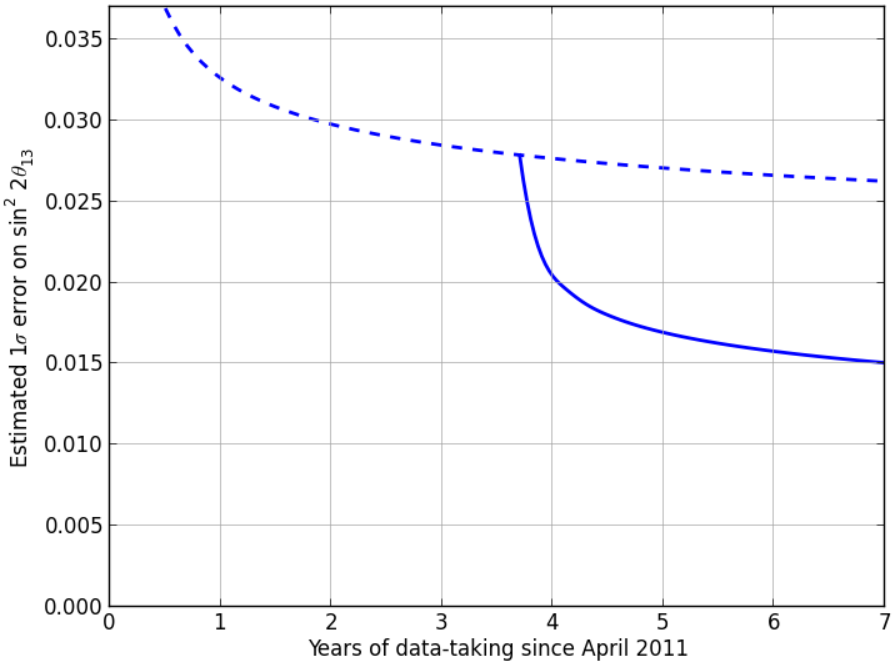

Figure 18.2: Estimated Gd-only, Rate+Shape precision as a function of time since the start of Double Chooz data-taking in April 2011. The solid curve shows precision from a combined ND+FD analysis, assuming physics-quality ND data begins on January 1, 2015. The dashed curve shows precision from an FD-only analysis.

assumed here. These enhancements will increase both single-detector and two-detector $\sin^2 2\theta_{13}$ precision.

One basic enhancement may be a re-optimization of the prompt energy binning scheme. The current scheme follows from the considerations in Sec. 15.2.3. If more finely binned S_k become available, or if the reactor model is removed as described later, a finer binning scheme may be feasible. Finer binning has been shown to have only modest impact on a single-detector analysis, but early tests suggest it has a larger impact on two-detector analyses. That finding is reasonable, since a higher-precision picture of the unoscillated antineutrino spectrum should provide a better constraint for the FD measurement.

Another substantial boost will be the inclusion of H-channel data. Figure 18.2 shows that $\sin^2 2\theta_{13}$ precision continues to increase significantly for many years after the ND begins operation, indicating that the measurement is statistically limited. As demonstrated with the H-based analysis in this theses, the H channel can provide about twice the signal of the Gd channel in the same time span. Now that the systematics of the H measurement approach those of Gd, adding the H channel is nearly as useful as tripling the data-taking time. Of course, as noted many times throughout this thesis, combining the Gd and H channels at the oscillation fit level requires difficult evaluations of the inter-analysis correlations. A simpler approach is to combine the channels at the selection level, adjusting the requirements for ΔT, delayed E_{vis}, and other parameters so that Gd and H signals can be selected together. This strategy may be employed in future analyses.

Finally, inputs related to the reactor model will probably evolve beyond those used here. As noted, this study uses a preliminary estimate for the inter-detector correlations of reactor uncertainties. Continuing studies will provide more perspective on these correlations. The uncertainty models within each detector may also change if new S_k, with updated uncertainties, become available. Although these models become progressively less influential as ND data accumulates, they have some impact in the early stages of two-detector analyses. Another possible update is removal of the Bugey4 anchor. That change would would slightly decrease $\sin^2 2\theta_{13}$ precision in early two-detector analyses but make the measurement more self-contained. One further option would be remove the reactor model almost entirely, so that even in early stages the ND provides the only constraint on the reactor flux. A mechanism for this approach is described in the following chapter.

Chapter 19

Measurements beyond $\sin^2 2\theta_{13}$

Beyond increasing $\sin^2 2\theta_{13}$ precision, the two-detector configuration creates opportunities for several new measurements. This chapter gives a qualitative discussion of two possibilities. The first is a search for sterile neutrino signals, introduced in Ch. 5. The second is a precise measurement of the reactor antineutrino spectrum, which will help to address the issues discussed in Ch. 17. In Double Chooz, both of these analyses gain distinctive advantages from the site geometry and periods of data-taking in which only one reactor is operating.

19.1 Sterile neutrinos

The $\sin^2 2\theta_{13}$ analyses in this thesis were designed to be insensitive to sterile neutrinos. All fits use an oscillation model with only one mixing angle and a single mass splitting, constrained near the expected Δm^2_{ee}. The Bugey4 anchor ensures that flux predictions match the IBD rate observed close to reactors, eliminating the $\sim 6\%$ discrepancy behind the so-called reactor antineutrino anomaly.

Expanding the oscillation model and removing the Bugey4 anchor will give Double Chooz some sensitivity to sterile neutrinos. A Rate+Shape fit using the survival probability of Eq. 5.6 can test $3+1$ models with broadly ranging Δm^2_{41}. The far detector offers limited sensitivity on its own, due to its relatively low signal rate and relatively long baseline. At ~ 1 km from the reactors, oscillations from $\Delta m^2_{41} \gtrsim 0.1$ eV^2 average to an energy-independent normalization shift proportional to $\sin^2 2\theta_{14}$. The $\sim 2.5\%$ flux normalization uncertainty in a Bugey4-free model hinders sensitivity to the small $\sin^2 2\theta_{14}$ favored by current observations.

Data from the near detector will expand the sterile neutrino parameter space in which Double Chooz is sensitive. The shorter baseline of the near detector enables it to observe or exclude signals from higher Δm^2_{41}. A joint fit to the near and far detectors will combine information from both baselines and provide even greater coverage. Even for $\Delta m^2_{41} \gtrsim 1$ eV^2, for which the near detector would also see only a normalization shift, the statistics boost from the near detector will enhance sensitivity. In any two-detector fit, as in the fit from the previous chapter, the nearly-isoflux layout of the Double Chooz site will suppress reactor-related uncertainties.

Sterile neutrino studies will also benefit from a unique opportunity afforded by the Double Chooz site. Because the Chooz power station has only two reactors, and both require periodic refueling and service, there are long periods in which only one reactor is operating. Single-reactor periods comprise about a quarter of the previously analyzed far detector runtime. In such periods, all signal can be attributed to a single reactor, removing ambiguity about the reactor-detector baseline. Moreover, as indicated in Tab. 7.1, each detector is located at different baselines with respect to reactors B1 and B2. The relative difference is especially notable for the near detector. Precise baseline knowledge, plus the ability to observe signal from two distinct baselines, is valuable in searches for oscillations at large Δm^2_{41}. The dual-baseline approach has been explored in a study by Double Chooz collaborators [130].

Inevitably, its relatively long baselines will prevent Double Chooz from achieving the sterile neutrino sensitivity of upcoming short-baseline experiments. Nonetheless, any new observations or limits may have some value at this ever-shifting frontier of neutrino phenomenology.

19.2 Reactor antineutrino spectrum

If the features discussed in Ch. 17 result from reactor modeling issues, they should emerge almost identically in near detector data. Near-far consistency in these features will further disfavor the cosmogenic background hypothesis, since the signal to background ratio is somewhat different for each detector. Assuming the reactor hypothesis explains the anomalous spectrum features, the near detector data will also allow them to be measured with useful precision. A possible strategy for this measurement is described in the following paragraphs.

Much as pull parameters describe the best-fit shape for the FN+SM spectrum in H-III, a set

of pull parameters could be used to determine the best-fit, no-oscillation signal spectrum shape. The number of pull parameters for the signal spectrum would likely be larger than for the FN+SM spectrum. A basic parametrization would include one pull parameter for each bin in the predicted E_ν spectrum, or, if the $E_\nu \to E_{vis}$ relation is sufficiently similar for the two detectors, each bin in the E_{vis} spectrum. For simplicity, the following discussion uses the latter convention. Here, the new parameters, $\boldsymbol{\zeta}$, are defined as multiplicative factors on the signal prediction in each bin, so that:

$$\mu_i^\nu = \zeta_i \sum_j^{N_i} w_j \tag{19.1}$$

with N_i and w_j defined as in Eq. 15.36. In this scheme, $\overline{\zeta_i} = 1$. The parameter set $\boldsymbol{\zeta}$ can be built into the Rate+Shape χ^2, producing:

$$\begin{aligned}\chi^2(\sin^2 2\theta_{13}, \boldsymbol{\eta}, \boldsymbol{\zeta}, \kappa) = &\left(\mathbf{n} - \boldsymbol{\mu}(\sin^2 2\theta_{13}, \boldsymbol{\eta}, \boldsymbol{\zeta}, \kappa)\right)^T \mathbf{M}^{-1}(\boldsymbol{\eta}) \left(\mathbf{n} - \boldsymbol{\mu}(\sin^2 2\theta_{13}, \boldsymbol{\eta}, \boldsymbol{\zeta}, \kappa)\right) \\ &+ \left(\overline{\boldsymbol{\eta}} - \boldsymbol{\eta}\right)^T \mathbf{V}^{-1} \left(\overline{\boldsymbol{\eta}} - \boldsymbol{\eta}\right) \\ &+ \left(\overline{\boldsymbol{\zeta}} - \boldsymbol{\zeta}\right)^T \mathbf{U}^{-1} \left(\overline{\boldsymbol{\zeta}} - \boldsymbol{\zeta}\right) \\ &+ \chi^2_{off}(\sin^2 2\theta_{13}, \boldsymbol{\eta}, \kappa)\end{aligned} \tag{19.2}$$

Here, κ is a pull parameter covering signal normalization uncertainties common to each bin, namely reactor flux normalization and detection efficiency uncertainties. The parameter set $\boldsymbol{\eta}$ includes all pull parameters other than κ and $\boldsymbol{\zeta}$. Since they are now covered with pull terms, all signal-related uncertainties are removed from $\mathbf{M}$.

The term in the third line of Eq. 19.2 is optional. It can be included in order to constrain $\boldsymbol{\zeta}$ with a reactor model, in which case $\mathbf{U}$ represents the shape uncertainty in that model. This matrix can be derived from the $\mathbf{M}^{rxtr}$ described in Sec. 15.3.1 by extracting the component describing shape uncertainties (according to App. A) and fractionalizing it with respect to the μ_i. For a reactor model–independent fit, the term can be excluded. In that case, $\boldsymbol{\zeta}$ becomes a set of free parameters, and κ must be removed to avoid redundancy.

Equation 19.2 can be used for single-detector fits, but $\boldsymbol{\zeta}$ must be constrained to avoid total degeneracy with the other fit parameters. Even if $\boldsymbol{\zeta}$ is constrained, remaining degeneracy limits the usefulness of a single-detector fit. Including both detectors in the fit breaks most of this degeneracy, since each has different background populations and oscillation effects. A two-detector fit with free $\boldsymbol{\zeta}$ will reveal the best-fit, no-oscillation IBD spectrum, in the absence of any reactor model constraints.

Implicit in a fit with free $\boldsymbol{\zeta}$ is the assumption that, except for oscillation effects, the IBD spectra in the near and far detectors are identical. In reality, small differences may exist in the energy scale or other aspects of the $E_\nu \rightarrow E_{vis}$ relation. Placing pull parameters on the E_ν spectrum, instead of E_{vis}, would allow for such differences. However, the assumption of identical E_ν spectrum would remain. In Double Chooz, that assumption is substantially validated by the nearly-isoflux site geometry. It is guaranteed to be true in the periods when only one reactor is operating. In experiments with more reactors and less ideal detector positions, measurements of the antineutrino spectrum retain greater dependence on reactor simulations. For this reason, Double Chooz measurements will add unique information to the continuing study of the reactor antineutrino spectrum.

Chapter 20

Conclusions

This thesis has presented the final set of electron antineutrino disappearance measurements made in the single-detector configuration of Double Chooz. The Gd-based analysis yields the most precise oscillation measurement yet made in the experiment, $\sin^2 2\theta_{13} = 0.090^{+0.032}_{-0.029}$. The H-based analysis provides a statistically independent validation and demonstrates the potential for high-precision IBD detection without a Gd dopant. Two types of oscillation fits exploit different dimensions of the data. The Rate+Shape approach uses the highly informative prompt energy spectrum to disentangle backgrounds and oscillated signal. The RRM approach uses a more minimal model to extract $\sin^2 2\theta_{13}$ from signal variation with reactor power. Agreement of these complementary methods over two disjoint datasets provides strong support for the Double Chooz oscillation results.

Although the single-detector Double Chooz analyses do not lead the world in $\sin^2 2\theta_{13}$ precision, their distinct systematics and attention to detail enhance confidence in the results of more sensitive experiments. Overall, global efforts to measure $\sin^2 2\theta_{13}$ with reactor antineutrinos have been a remarkable success. Within barely a year, beginning with the first analysis from Double Chooz, θ_{13} has risen from the least to the most well-known neutrino mixing angle. The confirmation of a relatively large $\sin^2 2\theta_{13}$ has accelerated pursuits of the mass hierarchy and CP violation in neutrino oscillations. Already, joint considerations of the reactor antineutrino experiments and T2K hint at a large value of the CP-violating phase δ.

Beyond measurements of $\sin^2 2\theta_{13}$, analyses described in this thesis have helped to reveal unexpected features in the reactor antineutrino spectrum. These features have inspired new lines of work in the nuclear and reactor physics communities. A more precise understanding of the reactor

antineutrino spectrum will benefit neutrino oscillation experiments, nuclear physics models, and potentially even practical technologies.

The start of the near detector, some months before this thesis was written, opened a new era for Double Chooz. In the coming years, data from the near detector will enable world-class $\sin^2 2\theta_{13}$ measurements capable of closely checking the leading Daya Bay results. The two-detector configuration will also allow sterile neutrino searches and a unique analysis of the reactor antineutrino spectrum. Through these measurements, and perhaps others not yet anticipated, Double Chooz will continue to advance physics within and beyond the neutrino sector.

Part V

Bibliography

Bibliography

[1] W. Pauli. Letter addressed Dear Radioactive Ladies and Gentlemen, 1956. Translated by Kurt Riesselmann.

[2] E. Fermi. Versuch einer Theorie der β-Strahlen. *Zeitschrift für Physik*, 88:161–177, 1934.

[3] F. L. Wilson. Fermi's Theory of Beta Decay. *Am.J.Phys*, 36:1150, 1968.

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

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

[6] S. M. Bilenky and J. Hosek. Glashow-Weinberg-Salam theory of electroweak interactions and the neutral currents. *Phys.Rept.*, 90:73–157, 1982.

[7] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys.Rev.Lett.*, 13:321–323, 1964.

[8] G. Guralnik, C. Hagen, and T. Kibble. Global Conservation Laws and Massless Particles. *Phys.Rev.Lett.*, 13:585–587, 1964.

[9] P. W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys.Rev.Lett.*, 13:508–509, 1964.

[10] K. Olive et al. Review of Particle Physics. *Chin.Phys.*, C38:090001, 2014.

[11] D. Harris and K. Riesselmann. Deconstruction: Standard Model discoveries. *Symmetry Breaking*, 06:30–31, 2009.

[12] B. Pontecorvo. Mesonium and anti-mesonium. *Sov.Phys.JETP*, 6:429, 1957.

[13] M. Gell-Mann and A. Pais. Behavior of Neutral Particles under Charge Conjugation. *Phys.Rev.*, 97:1387–1389, 1955.

[14] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the Unified Model of Elementary Particles. *Prog.Theor.Phys.*, 28:870–880, 1962.

[15] C. Giunti and C. W. Kim. Quantum Mechanics of Neutrino Oscillations. *Found.Phys.Lett.*, 14:213–229, 2001.

[16] A. Esteban-Pretel, S. Pastor, R. Tomas, G. G. Raffelt, and G. Sigl. Decoherence in supernova neutrino transformations suppressed by deleptonization. *Phys.Rev.*, D76:125018, 2007.

[17] J. Davis, Raymond, D. S. Harmer, and K. C. Hoffman. Search for Neutrinos from the Sun. *Phys.Rev.Lett.*, 20:1205–1209, 1968.

[18] R. Davis. A review of the Homestake solar neutrino experiment. *Prog.Part.Nucl.Phys.*, 32:13–32, 1994.

[19] Q. Ahmad et al. Direct Evidence for Neutrino Flavor Transformation from Neutral Current Interactions in the Sudbury Neutrino Observatory. *Phys.Rev.Lett.*, 89:011301, 2002.

[20] B. Aharmim et al. Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory. *Phys.Rev.*, C72:055502, 2005.

[21] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys.Rev.*, D17:2369–2374, 1978.

[22] S. Mikheev and A. Y. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov.J.Nucl.Phys.*, 42:913–917, 1985.

[23] A. Gando et al. Constraints on θ_{13} from A Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND. *Phys.Rev.*, D83:052002, 2011.

[24] Y. Fukuda et al. Evidence for Oscillation of Atmospheric Neutrinos. *Phys.Rev.Lett.*, 81:1562–1567, 1998.

[25] F. Boehm, J. Busenitz, B. Cook, G. Gratta, H. Henrikson, et al. Final results from the Palo Verde neutrino oscillation experiment. *Phys.Rev.*, D64:112001, 2001.

[26] M. Apollonio et al. Search for neutrino oscillations on a long baseline at the CHOOZ nuclear power station. *Eur.Phys.J.*, C27:331–374, 2003.

[27] K. Abe et al. Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam. *Phys.Rev.Lett.*, 107:041801, 2011.

[28] Y. Abe et al. Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment. *Phys.Rev.Lett.*, 108:131801, 2012.

[29] F. An et al. Observation of Electron-Antineutrino Disappearance at Daya Bay. *Phys.Rev.Lett.*, 108:171803, 2012.

[30] J. Ahn et al. Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. *Phys.Rev.Lett.*, 108:191802, 2012.

[31] Kismalac. Neutrino mass hierarchy. Licensed under Creative Commons Attribution-Share Alike 3.0 Unported.

[32] V. Aseev et al. An upper limit on electron antineutrino mass from Troitsk experiment. *Phys.Rev.*, D84:112003, 2011.

[33] P. Ade et al. Planck 2013 results. XVI. Cosmological parameters. *Astron.Astrophys.*, 571:A16, 2014.

[34] P. Minkowski. $\mu \to e\gamma$ at a Rate of One Out of 10^9 Muon Decays? *Phys.Lett.*, B67:421–428, 1977.

[35] M. Gell-Mann, P. Raymond, and R. Slansky. Complex Spinors and Unified Theories. In D. Freedman and P. V. Nieuwenhuizen, editors, *Supergravity*. North Holland, Amsterdam, 1979.

[36] T. Yanagida. Horizontal Symmetry and Masses of Neutrinos. *Prog.Theor.Phys.*, 64:1103, 1980.

[37] R. N. Mohapatra and G. Senjanovic. Neutrino Mass and Spontaneous Parity Violation. *Phys.Rev.Lett.*, 44:912, 1980.

[38] J. Schechter and J. Valle. Neutrino Masses in SU(2) x U(1) Theories. *Phys.Rev.*, D22:2227, 1980.

[39] R. Barate et al. Measurement of the Z resonance parameters at LEP. *Eur.Phys.J.*, C14:1–50, 2000.

[40] A. Palazzo. Constraints on very light sterile neutrinos from θ_{13}-sensitive reactor experiments. *JHEP*, 1310:172, 2013.

[41] A. de Gouvea and T. Wytock. Light Sterile Neutrino Effects at theta(3)-Sensitive Reactor Neutrino Experiments. *Phys.Rev.*, D79:073005, 2009.

[42] K. Abazajian, M. Acero, S. Agarwalla, A. Aguilar-Arevalo, C. Albright, et al. Light Sterile Neutrinos: A White Paper, 2012.

[43] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of anti-neutrino(electron) appearance in a anti-neutrino(muon) beam. *Phys.Rev.*, D64:112007, 2001.

[44] A. Aguilar-Arevalo et al. A Search for Electron Antineutrino Appearance at the $\Delta m^2 \sim 1$ eV^2 Scale. *Phys.Rev.Lett.*, 103:111801, 2009.

[45] A. Aguilar-Arevalo et al. Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations. *Phys.Rev.Lett.*, 105:181801, 2010.

[46] G. Mention, M. Fechner, T. Lasserre, T. Mueller, D. Lhuillier, et al. The Reactor Antineutrino Anomaly. *Phys.Rev.*, D83:073006, 2011.

[47] A. Pichlmaier, V. Varlamov, K. Schreckenbach, and P. Geltenbort. Neutron lifetime measurement with the UCN trap-in-trap MAMBO II. *Phys.Lett.*, B693:221–226, 2010.

[48] F. An et al. Search for a Light Sterile Neutrino at Daya Bay. *Phys.Rev.Lett.*, 113:141802, 2014.

[49] B. Armbruster et al. Upper limits for neutrino oscillations muon-anti-neutrino $\to$ electron-anti-neutrino from muon decay at rest. *Phys.Rev.*, D65:112001, 2002.

[50] P. Astier et al. Final NOMAD results on muon-neutrino → tau-neutrino and electron-neutrino → tau-neutrino oscillations including a new search for tau-neutrino appearance using hadronic tau decays. *Nucl.Phys.*, B611:3–39, 2001.

[51] P. Adamson et al. Active to Sterile Neutrino Mixing Limits from Neutral-Current Interactions in MINOS. *Phys.Rev.Lett.*, 107:011802, 2011.

[52] P. Adamson et al. Search for sterile neutrino mixing in the MINOS long baseline experiment. *Phys.Rev.*, D81:052004, 2010.

[53] F. Dydak, G. Feldman, C. Guyot, J. Merlo, H. Meyer, et al. A Search for Muon-neutrino Oscillations in the Δm^2 Range 0.3 eV^2 to 90 eV^2. *Phys.Lett.*, B134:281, 1984.

[54] Y. Declais, H. de Kerret, B. Lefievre, M. Obolensky, A. Etenko, et al. Study of reactor anti-neutrino interaction with proton at Bugey nuclear power plant. *Phys.Lett.*, B338:383–389, 1994.

[55] C. Cowan, F. Reines, F. Harrison, H. Kruse, and A. McGuire. Detection of the Free Neutrino: A Confirmation. *Science*, 124:103–104, 1956.

[56] C. R. Nave. HyperPhysics, 2014.

[57] P. Vogel and J. F. Beacom. Angular distribution of neutron inverse beta decay, $\bar{\nu}_e + p \rightarrow e^+ + n$. *Phys.Rev.*, D60:053003, 1999.

[58] T. Mueller, D. Lhuillier, M. Fallot, A. Letourneau, S. Cormon, et al. Improved Predictions of Reactor Antineutrino Spectra. *Phys.Rev.*, C83:054615, 2011.

[59] G. Leinweber et al. Neutron Capture and Total Cross-Section Measurements and Resonance Parameters of Gadolinium. *Nuc.Sci.Eng.*, 154:261–279, 2006.

[60] G. Mention, T. Lasserre, and D. Motta. A unified analysis of the reactor neutrino program towards the measurement of the θ_{13} mixing angle. *J.Phys.Conf.Ser.*, 110:082013, 2008.

[61] P. Hitchin. Excavating Chooz A, 2010. [Online; accessed 30-April-2015].

[62] T. Bezerra et al. A Global Fit Determination of Effective Δm^2_{31} from Baseline Dependence of Reactor ν Disappearance. *Phys.Lett.B*, 725:271–276, 2013.

[63] V. Fischer and M. Vivier. Geodetic survey's results, 2013. Internal Double Chooz communication 4934.

[64] International Atomic Energy Agency. Chooz B1, Chooz B2. Technical report, Power Reactor Information System, 2015. [Online; accessed 22-May-2015].

[65] T. Mueller. *Expérience Double Chooz: simulation des spectres antineutrinos issue des réacteurs*. PhD thesis, Université de Paris-Sud XI, 9 2010.

[66] F. Ardellier et al. Double Chooz: A Search for the neutrino mixing angle θ_{13}, 2006.

[67] A. Onillon and C. Jones. Reactor and $\bar{\nu}_e$ spectrum prediction for the Double Chooz experiment, 6 2012. Poster for Neutrino2012 conference.

[68] H. Back et al. Phenylxylylethane (PXE): A high-density, high-flashpoint organic liquid scintillator for applications in low-energy particle and astrophysics experiments. *Nucl.Instrum.Meth.A*, 2004.

[69] S. Abe et al. Precision Measurement of Neutrino Oscillation Parameters with KamLAND. *Phys.Rev.Lett.*, 100:221803, 2008.

[70] C. Aberle, C. Buck, B. Gramlich, F. Hartmann, M. Lindner, et al. Large scale Gd-beta-diketonate based organic liquid scintillator production for antineutrino detection. *JINST*, 7:P06008, 2012.

[71] P. Pfahler. *Realization of the low background neutrino Detector Double Chooz: From the development of a high-purity liquid and gas handling concept to first neutrino data*. PhD thesis, Technische Universität München, 12 2012.

[72] Hamamatsu Photonics K.K. Large photocathode area photomultiplier tubes. Technical report, Electron Tube Division, 2008. [Online; accessed 3-May-2015].

[73] T. Matsubara, T. Haruna, T. Konno, Y. Endo, M. Bongrand, et al. Evaluation of 400 low background 10-in. photo-multiplier tubes for the Double Chooz experiment. *Nucl.Instrum.Meth.*, A661:16–25, 2012.

[74] C. Bauer, E. Borger, R. Hofacker, K. Janner, F. Kaether, et al. Qualification Tests of 474 Photomultiplier Tubes for the Inner Detector of the Double Chooz Experiment. *JINST*, 6:P06008, 2011.

[75] E. Calvo, M. Cerrada, C. Fernandez-Bedoya, I. Gil-Botella, C. Palomares, et al. Characterization of large area photomutipliers under low magnetic fields: Design and performances of the magnetic shielding for the Double Chooz neutrino experiment. *Nucl.Instrum.Meth.*, A621:222–230, 2010.

[76] P. Wahnon. Absolute Light Level Measurements of the Light Injection Calibration System. Undergraduate thesis, University of Sussex, 4 2010.

[77] CAEN. A1535 data sheet. Technical report, Electronic Instrumentation, 2015. [Online; accessed 21-May-2015].

[78] Y. Abe et al. Reactor electron antineutrino disappearance in the Double Chooz experiment. *Phys.Rev.*, D86:052008, 2012.

[79] E. Caden. Front End Electronics for Double Chooz, 2011. Internal Double Chooz communication 2345.

[80] CAEN. VX1721 data sheet. Technical report, Electronic Instrumentation, 2015. [Online; accessed 21-May-2015].

[81] Y. Abe, T. Akiri, A. Cabrera, B. Courty, J. Dawson, et al. The Waveform Digitiser of the Double Chooz Experiment: Performance and Quantisation Effects on PhotoMultiplier Tube Signals. *JINST*, 8:P08015, 2013.

[82] P. Barrillon and others. MAROC: Multi-Anode ReadOut Chip. Technical report, IN2P3, LAL-Orsay, 2007. [Online; accessed 21-May-2015].

[83] Altera Corporation. Cyclone FPGA Family Data Sheet. Technical report, Cyclone series, 2008. [Online; accessed 21-May-2015].

[84] Y. Abe et al. First Measurement of θ_{13} from Delayed Neutron Capture on Hydrogen in the Double Chooz Experiment. *Phys.Lett.*, B723:66–70, 2013.

[85] Y. Abe et al. Background-independent measurement of θ_{13} in Double Chooz. *Phys.Lett.*, B735:51–56, 2014.

[86] Y. Abe et al. Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector. *JHEP*, 1410:086, 2014.

[87] Y. Abe et al. Precision Muon Reconstruction in Double Chooz. *Nucl.Instrum.Meth.*, A764:330–339, 2014.

[88] E. Tornu and S. Fortier. Improving Pressurized Water Reactor Performance Through Instrumentation: Application Case of Reducing Uncertainties on Thermal Power. Technical report, Electric Power Research Institute, 2001.

[89] O. Meplan et al. MURE: MCNP Utility for Reactor Evolution – Description of the methods, first applications and results. In *Nuclear Power for the XXIst Century: From basic research to high-tech industry*. ENC 2005: European Nuclear Conference, 2005.

[90] G. Marleau, A. Hébert and R. Roy. A User Guide for DRAGON. Technical report, Institut de génie nucléaire, École Polytechnique de Montréal, 2000. [Online; accessed 9-Jun-2015].

[91] C. Jones, A. Bernstein, J. Conrad, Z. Djurcic, M. Fallot, et al. Reactor Simulation for Antineutrino Experiments using DRAGON and MURE. *Phys.Rev.*, D86:012001, 2012.

[92] V. Kopeikin, L. Mikaelyan, and V. Sinev. Reactor as a Source of Antineutrinos: Thermal Fission Energy. *Phys.Atom.Nucl.*, 67:1892–1899, 2004.

[93] K. Schreckenbach, G. Colvin, W. Gelletly, and F. Von Feilitzsch. Determination of the antineutrino spectrum from U-235 thermal neutron fission products up to 9.5 MeV. *Phys.Lett.*, B160:325–330, 1985.

[94] F. Von Feilitzsch, A. Hahn, and K. Schreckenbach. Experimental Beta Spectra From Pu-239 And U-235 Thermal Neutron Fission Products And Their Correlated Anti-neutrinos Spectra. *Phys.Lett.*, B118:162–166, 1982.

[95] A. Hahn, K. Schreckenbach, G. Colvin, B. Krusche, W. Gelletly, et al. Anti-neutrino Spectra From ^{241}Pu and ^{239}Pu Thermal Neutron Fission Products. *Phys.Lett.*, B218:365–368, 1989.

[96] P. Huber. On the determination of anti-neutrino spectra from nuclear reactors. *Phys.Rev.*, C84:024617, 2011.

[97] N. Haag, A. Gütlein, M. Hofmann, L. Oberauer, W. Potzel, et al. Experimental Determination of the Antineutrino Spectrum of the Fission Products of ^{238}U. *Phys.Rev.Lett.*, 112(12):122501, 2014.

[98] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Dubois, et al. Geant4 Developments and Applications. *IEEE Trans.Nucl.Sci.*, 53:270, 2006.

[99] J. Apostolakis et al. GEANT4 Physics Lists for HEP. In *Nuclear Science Symposium Conference Record*. Institute of Electrical and Electronics Engineers, 2008.

[100] J. R. Granada. Slow-neutron scattering by molecular gases: A synthetic scattering function. *Phys.Rev.*, B31:4167, 1985.

[101] J. B. Birks. Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations. *Proc.Phys.Soc.*, A64:874, 1951.

[102] C. Aberle. *Optimization, simulation and analysis of the scintillation signals in the Double Chooz experiment*. PhD thesis, Ruperto-Carola University, 12 2011.

[103] C. Aberle, C. Buck, F. Hartmann, S. Schonert, and S. Wagner. Light output of Double Chooz scintillators for low energy electrons. *JINST*, 6:P11006, 2011.

[104] A. Stueken. Trigger Efficiency Distribution of Singles, 2013. Internal Double Chooz communication 4491-v3.

[105] E. Conover. *Muon-Induced Backgrounds in the Double Chooz Neutrino Oscillation Experiment*. PhD thesis, University of Chicago, 6 2014.

[106] P. Novella. DCIII nGd candidates: Blessed plots, 2014. Internal Double Chooz communication 5608-v4.

[107] Nuclear Energy Agency. TRIPOLI-4 version 8.1, 3D general purpose continuous energy Monte Carlo Transport code, 2013. [Online; accessed 12-Jul-2015].

[108] National Nuclear Data Center. Chart of Nuclides. Interactive database.

[109] TUNL Nuclear Data Evaluation Project. Energy Level Diagram, ^{9}Li.

[110] TUNL Nuclear Data Evaluation Project. Energy Level Diagram, ^{8}He.

[111] S. Abe et al. Production of Radioactive Isotopes through Cosmic Muon Spallation in KamLAND. *Phys. Rev.*, C81:025807, 2010.

[112] G. Bellini et al. Cosmogenic Backgrounds in Borexino at 3800 m water-equivalent depth. *JCAP*, 1308:049, 2013.

[113] S. Enomoto. *Neutrino Geophysics and Observation of Geo-Neutrinos at KamLAND*. PhD thesis, Tohoku University, 2 2005.

[114] A. Hoecker, P. Speckmayer, J. Stelzer, J. Therhaag, E. von Toerne, and H. Voss. Toolkit for Multivariate Data Analysis. *PoS*, ACAT:040, 2007.

[115] Y. Abe, J. dos Anjos, J. Barriere, E. Baussan, I. Bekman, et al. Ortho-positronium observation in the Double Chooz experiment. *JHEP*, 1410:32, 2014.

[116] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Annals Math.Statist.*, 9(1):60–62, 1938.

[117] G. J. Feldman and R. D. Cousins. A Unified Approach to the Classical Statistical Analysis of Small Signals. *Phys.Rev.*, D57:3873–3889, 1998.

[118] G. Fogli, E. Lisi, A. Marrone, D. Montanino, and A. Palazzo. Getting the most from the statistical analysis of solar neutrino oscillations. *Phys.Rev.*, D66:053010, 2002.

[119] Y. Abe et al. Direct Measurement of Backgrounds using Reactor-Off Data in Double Chooz. *Phys.Rev.*, D87(1):011102, 2013.

[120] P. Adamson et al. Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS. *Phys.Rev.Lett.*, 110(25):251801, 2013.

[121] H. Nunokawa, S. J. Parke, and R. Zukanovich Funchal. Another possible way to determine the neutrino mass hierarchy. *Phys.Rev.*, D72:013009, 2005.

[122] H. Minakata, H. Nunokawa, S. J. Parke, and R. Zukanovich Funchal. Determining Neutrino Mass Hierarchy by Precision Measurements in Electron and Muon Neutrino Disappearance Experiments. *Phys.Rev.*, D74:053008, 2006.

[123] F. James and M. Winkler. MINUIT User's Guide, 2004.

[124] J. Zhao, Z.-Y. Yu, J.-L. Liu, X.-B. Li, F.-H. Zhang, et al. $^{13}C(\alpha,n)^{16}O$ background in a liquid scintillator based neutrino experiment. *Chin.Phys.*, C38(11):116201, 2014.

[125] F. An et al. Spectral Measurement of Electron Antineutrino Oscillation Amplitude and Frequency at Daya Bay. *Phys.Rev.Lett.*, 112:061801, 2014.

[126] S.-H. Seo. New Results from RENO and The 5 MeV Excess. In *26th International Conference on Neutrino Physics and Astrophysics (Neutrino 2014) Boston, Massachusetts, United States, June 2-7, 2014*, 2014.

[127] W. Zhong, for the Daya Bay collaboration. Measurement of the reactor antineutrino flux and spectrum at Daya Bay. Talk at ICHEP 2014 conference.

[128] D. Dwyer and T. Langford. Spectral Structure of Electron Antineutrinos from Nuclear Reactors. *Phys.Rev.Lett.*, 114(1):012502, 2015.

[129] A. Hayes, J. Friar, G. Garvey, D. Ibeling, G. Jungman, et al. The Origin and Implications of the Shoulder in Reactor Neutrino Spectra, 2015.

[130] M. Bergevin, C. Grant, and R. Svoboda. A potential sterile neutrino search utilizing spectral distortion in a two-reactor/one-detector configuration, 2013.

[131] M. Shaevitz. Separating Components of Error Matrices. Technical report, MiniBooNE, 2008.

[132] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nucl.Instrum.Meth.*, A389:81–86, 1997.

[133] W. Verkerke and D. P. Kirkby. The RooFit toolkit for data modeling. *eConf*, C0303241:MOLT007, 2003.

[134] M. Goto. C++ Interpreter – CINT. CQ publishing.

[135] R. Carr. The CUfits Package, 2015. Internal Double Chooz communication 6097.

[136] D. van Heesch. Doxygen user manual, 2015. [Online; accessed 10-Jul-2015].

Part VI

Appendices

Appendix A

Covariance matrix procedures

A.1 Multisim methods

A Gaussian random variable x with mean μ and variance σ^2 can be simulated by the following basic Monte Carlo procedure:

I.1 Draw a (pseduo-)random number u from a standard normal distribution.

I.2 Take the square root of the variance σ^2 to find the standard deviation σ.

I.3 Calculate σu.

I.4 Add the result of Step I.3 to μ to generate $x = \mu + \sigma u$.

This procedure can be generalized to simulate a vector of b correlated variables $\mathbf{x} = (x_1, ...x_b)$ with means $\boldsymbol{\mu} = (\mu_1, ..., \mu_b)$ and covariances described by the matrix $\mathbf{V}$. The above steps generalize to:

II.1 Independently draw b (pseduo-)random numbers u_i from a standard normal distribution, forming the vector $\mathbf{u}$.

II.2 Perform a Cholesky decomposition of $\mathbf{V}$ to find the lower triangular matrix $\mathbf{L}$.

II.3 Calculate $\mathbf{Lu}$.

II.4 Add the result of Step II.3 to $\boldsymbol{\mu}$ to generate $\mathbf{x} = \boldsymbol{\mu} + \mathbf{Lu}$.

Repeating Steps II.1–II.4 a total of N times produce a set of N simulated $\mathbf{x}$. This set can be used in a variety of applications. The following sections describe two common uses. In Double Chooz and other neutrino oscillation experiments, methods like these are often called multisims.

A.1.1 Frequentist tests

Suppose $\mathbf{x} = \mathbf{n}$ describes the number of events in each bin of an observable spectrum, $\boldsymbol{\mu}$ describes the expected bin contents, $\mathbf{V} = \mathbf{M}$ is the covariance matrix for those expectations, and the analysis does not include pull parameters. In this case, performing Steps II.1–II.4 for a large N produces an ensemble of simulated datasets. (If the analysis includes pull parameters, their values must be simulated according to the same procedure, with the resulting values applied to each simulated spectrum.) This ensemble represents variations around $\boldsymbol{\mu}$ which correspond to the uncertainties encoded in $\mathbf{M}$. Such ensembles are useful in frequentist techniques such as the confidence interval derivation of Feldman and Cousins [117].

A.1.2 Monte Carlo–based error propagation

Suppose $\mathbf{x}$ is some general set of parameters. For example, $\mathbf{x}$ could describe the bin contents of the simulated E_ν spectrum discussed in Sec. 15.3.1, for which $\boldsymbol{\mu}$ represents the central values. In this instance, $\mathbf{V} = \mathbf{M}^{rxtr}(E_\nu)$, as defined in Eq. 15.19. Uncertainties on $\mathbf{x}$ can be propagated to a different set of parameters $\mathbf{y}(\mathbf{x})$ through N performances of Steps II.1–II.4 and the following additional steps.

For each simulated $\mathbf{x}$, the corresponding $\mathbf{y}(\mathbf{x})$ is calculated. In the example given above, $\mathbf{y}$ describes the E_{vis} spectrum, which is determined from $\mathbf{x}$, the E_ν spectrum, by reweighting individual events. In the m^{th} simulation, the weight of the k^{th} event is:

$$w_m^k = \frac{(x_i)_m^k}{(\mu_i)^k} \tag{A.1}$$

where i identifies the bin in which the k^{th} event occurs. A covariance matrix $\mathbf{U}$ for the parameters $\mathbf{y}$ can be constructed according to:

$$M_{ij} = \frac{1}{N}\sum_{m=0}^{N}\Big((y_i(\mathbf{x}_m) - y_i(\boldsymbol{\mu})\Big)\Big((y_j(\mathbf{x}_m) - y_j(\boldsymbol{\mu})\Big) \tag{A.2}$$

where N should be made sufficiently large to produce negligible statistical uncertainty on the M_{ij}. In the example used here, $\mathbf{U} = \mathbf{M}^{rxtr}(E_{vis})$.

A.2 Separation of normalization, shape, and mixed components

Suppose the covariance matrix $\mathbf{M}$ represents uncertainty on a spectrum with predicted bin contents $\boldsymbol{\mu}$. This matrix can be decomposed into components representing normalization uncertainty, spectrum shape uncertainty, and mixed normalization-shape uncertainty. Normalization uncertainties are associated with linear scaling factors, as represented in Eq. 15.9, which have the same relative effect on all μ_i. Shape uncertainties are associated with parameters that change individual μ_i but conserve the total spectrum integral, $\mu_{tot} = \sum_i \mu_i$. Mixed uncertainties are associated with parameters that change μ_{tot} and have different effects on individual μ_i.

Following the derivation in [131], the elements of $\mathbf{M} = \mathbf{M}^{norm} + \mathbf{M}^{shape} + \mathbf{M}^{mixed}$ are given by:

$$
\begin{aligned}
M_{ij}^{norm} &= \frac{\mu_i \mu_j}{\mu_{tot}^2} \sum_{k,l} M_{kl} \\
M_{ij}^{shape} &= M_{ij} - \frac{\mu_j}{\mu_{tot}} \sum_k M_{ik} - \frac{\mu_i}{\mu_{tot}} \sum_k M_{kj} + \frac{\mu_i \mu_j}{\mu_{tot}^2} \sum_{k,l} M_{kl} \\
M_{ij}^{mixed} &= \frac{\mu_j}{\mu_{tot}} \sum_k M_{ik} + \frac{\mu_i}{\mu_{tot}} M_{kj} - 2\frac{\mu_i \mu_j}{\mu_{tot}^2} \sum_{k,l} M_{kl}
\end{aligned}
\tag{A.3}
$$

Appendix B

Oscillation fit software

B.1 Overview

All Rate+Shape analyses in this thesis were performed with CUfits, a C++ package designed to execute χ^2-based oscillation fits to Double Chooz data. A variety of other studies, including cross checks of the RRM method, have also been performed with CUfits. The name of the package stems from its origin in Cluster United, one of the analysis groups organized in preparation for the first Double Chooz publication. Since that analysis, CUfits has been expanded and revised continously, including a major overhaul in late 2014. This appendix outlines the basic structure and function of CUfits, as of the July 2015. The package supports many features beyond this discussed here, including a reactors-off term, analytic spectrum shapes, lattice-based minimization, and other details involved in the more complex fits of Ch. 15.

B.2 Basic code information

B.2.1 Dependencies

The CUfits package requires ROOT [132] installed with RooFit [133] and CINT [134].

B.2.2 Source code

The CUfits code is maintained on the Cluster United SVN repository. The code can be checked out by any user with access to that system. Installation instructions can be found in the Double Chooz

document database [135].

B.2.3 Directory structure

The CUfits directory structure is illustrated in Fig. B.1. The top-level `CUfits` directory contains a `Makefile` and `Linkdef.h` file which specify rules for making the CUfits package. The subdirectories have these contents:

Subdirectory	Contents
dat	Input files for fitter programs, generally in .root format, and some programs and other files used to generate them
docs	Auto-generated documentation, from Doxygen [136]
include	Links to header files in `CUfits/src`
lib	Shared library containing classes, namespaces, etc. defined in `CUfits/src`
mac	Executable programs, such as fitters in `fitter` directory, and associated text files listing inputs for these programs
src	Files defining common classes, namespaces, etc. used throughout the CUfits package

B.3 Basic fit example

B.3.1 Program description

This section describes a minimal example of a Rate+Shape fit program in CUfits. Users with source code access can find corresponding code in `CUfits/mac/fitter/secondPub/secondPub_fitter.cc`. The basic strategy of this program, like all fitter programs in CUfits, is to construct a predicted IBD spectrum based on a set of inputs, construct a χ^2 statistic from that prediction and a dataset, and minimize that statistic by adjusting the predicted spectrum.

The program requires a single command-line argument, the name of a configuration file. For this example, that file is `CUfits/mac/fitter/secondPub/GdII_singleIP_nonLiRed.txt`. It contains constants and file paths (relative to the `CUfits/dat` directory) used in this example program. The format of this file is described in Section B.4.1. In this example, the file specifies inputs corresponding

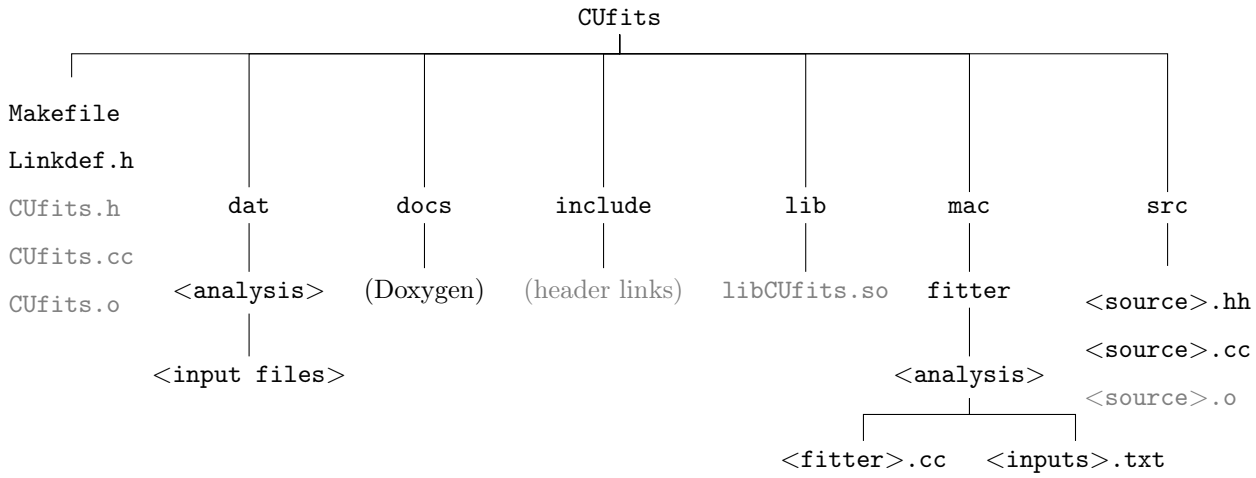

Figure B.1: The directory structure of the CUfits package. Items in gray are created during compilation.

to the previous Gd-based analysis, without the ^{9}Li veto. Because the official Rate+Shape fit used a ^{9}Li veto, included two integration periods, and treated some uncertainties via pull terms, the results of this example program differ slightly from the official, published results.

B.3.2 Outline of the program

This basic fit program can be divided into six sections (as indicated by major comment lines in `secondPub_fitter.cc`), as outlined below. Objects referenced in this outline are defined in the following section.

1. Set up the basics.
   (a) Define a `CUreader` object which will read parameters from the configuration file.
   (b) Enumerate the spectra to be fit. This example fits a single IBD spectrum: the Gd selection in the far detector.
2. Create the predicted spectrum.
   (a) Create a `CUspectrum` object, which will hold the prediction for the spectrum to be fit.

(b) Add predictions for signal and background event types from, based on `TTree` objects in .root files referenced in the configuration file.

(c) For diagnostic purposes, output the predicted number of events for each type.

3. Input the data, or to perform a sensitivity study, set the "data" equal to the prediction.

   (a) Load data from a .root file, or make a copy of the predicted spectrum and apply oscillation if desired.

   (b) For diagnostic purposes, output the number of data events.

4. Create (relative) covariance matrices.

   (a) Create (relative) normalization and uncorrelated error matrices, using CUfits functions.

   (b) Load (relative) shape matrices from .root files specified in the configuration file.

5. Create the χ^2 function (called the FCN, in MINUIT terminology).

   (a) Create a `CUpullInput` object, which in this example is just a null pointer.

   (b) Create a `CUchiSquare` object, which contains all of the information needed to calculate the χ^2 for this fit.

   (c) Create a `RooFormulaVar` object, which will pass that information to RooFit, where it will become the FCN which MINUIT will minimize.

6. Perform the fit.

   (a) Create a `RooMinuit` object, which will oversee the χ^2 minimization.

   (b) Call the MINUIT algorithms MIGRAD and HESSE to find the minimum and evaluate errors.

B.3.3 Program output

As the program runs, it will produce a series of outputs on the command line. Some of these lines come directly from CUfits functions, including the names of files from which inputs are being taken, predicted event counts, intermediate steps in χ^2 calculations, and information about fit settings. Interspersed within these lines are blocks of code produced by MINUIT [123], which is called through RooFit [133]. The final MINUIT block contains χ^2_{min}, fit parameter values at that point, and uncertainty estimates.

B.4 Classes and namespaces

The following classes and namespace are defined in `CUfits/src`.

B.4.1 `CUreader`

`CUreader` is a simple text-parsing class which reads input parameters from a configuration file. Programs generally require a single `CUreader` object.

B.4.2 `CUspectrum`

The `CUspectrum` class organizes all objects and functions needed to construct a predicted IBD spectrum and to dynamically update it during χ^2 minimization. One `CUspectrum` object is required for each IBD spectrum used in a fit. For example, a fit to Gd and H spectrum in each of the Near and Far Detectors would require four `CUspectrum` instantiations.

Qualitatively, a `CUspectrum` object contains the following components:

1. Nominal spectra for each signal and background event type (organized in to `CUeventSet` objects, described below), plus the total spectrum
2. Pull parameters for each event type, energy scale parameter, and other parameters

When in use in a fit, the class performs the following tasks:

1. Updates the predicted IBD spectrum as oscillation parameters and pull parameters change during a fit, via reweighting or the application of bin modulations as described in Sec. 15.4.2, through the function `CUspectrum::updateHists`
2. Returns event histograms, parameters, etc. for use in χ^2 calculations conducted in `CUchiSquare`, or for other purposes

B.4.3 `CUeventSet`

The `CUeventSet` organizes the objects associated with a single event type in a predicted spectrum. Qualitatively, it contains:

1. A `TTree` of events

2. An overall normalization factor for these events

3. (Optional) Sets of normalization factors which are applied depending on the prompt or delayed vertex of each event

B.4.4 `CUpullInput`

The `CUpullInput` class organizes central values, uncertainties (with support for asymmetric uncertainties), and correlations for pull parameters. The purpose of this class is to simplify the user-end input of these parameters, particularly correlations. Only one `CUpullInput` object is required for a fit, even with multiple spectra. If pull parameters are not used in a fit, as in the example of Section B.3, the `CUpullInput` passed to the `CUchiSquare` object should be a null pointer.

B.4.5 `CUchiSquare`

The `CUchiSquare` class organizes objects and functions needed to calculate a χ^2 statistic. Each fit requires a single `CUchiSquare` object; multiple fits to the same data require their own `CUchiSquare` objects. For example, a program which performs a Rate+Shape fit to a dataset, followed by a Rate-Only fit, would require two `CUchiSquare` instantiations. Qualitatively, a `CUchiSquare` object contains:

1. Vector of the data events

2. Copy of the predicted spectrum or spectra

3. Covariance matrices used in the fit

4. (Optional) Pull parameters used in the fit, along with their errors and correlations

5. Settings for how to construct the χ^2:

   (a) χ^2 type: Rate+Shape, Rate-Only, or Shape-Only (not currently supported)

   (b) Statistical error type: Pearson, Neyman, or James (average of previous two)

   (c) Option to use a fixed covariance matrix, rather than dynamically updating it (not currently supported)

(d) Option to include a reactor-off term

6. (Optional) Parameters used to calculate a reactor-off term

B.4.6 `CUphysics`

`CUphysics` is a namespace containing oscillation parameters and other global parameters used in the fit, the function which calculates oscillation probability, including optional sterile neutrino oscillations, and a variety of `enums` which organize parameters throughout the CUfits package.

B.4.7 `CUutils`

`CUutils` is a namespace which holds a variety of utilities used throughout CUfits, including functions for creating binning schemes, histograms, and covariance matrices.