Precision measurements of the astrophysical neutrino flux

by

Austin Schneider

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Physics)

at the

University of Wisconsin – Madison

2020

Defended on August 6th 2020

Dissertation approved by the following members of the Final Oral Committee:

Keith Bechtol · Assistant Professor of Physics

Jessi Cisewski-Kehe · Assistant Professor of Statistics

Francis L Halzen · Professor of Physics

Albrecht Karle · Professor of Physics

Abstract

The IceCube neutrino observatory has established the existence of a high-energy all-sky neutrino flux of astrophysical origin. This discovery was made using events interacting within a fiducial region of the detector surrounded by an active veto with reconstructed energy above 60 TeV, commonly known as the high-energy starting event sample or HESE. This work revisits the HESE sample analysis with an additional 4.5 years of data, newer glacial ice models, and improved systematics treatment. In addition to updating the data, this analysis seeks to validate previous measurements and explore all possible avenues that may have been overlooked. This work reports on the latest astrophysical neutrino flux measurements and presents a source search for astrophysical neutrinos. The compatibility of these observations with specific isotropic flux models proposed in the literature and generic power-law-like scenarios are explored. It is found that the astrophysical neutrino spectrum, when assumed equal for neutrinos and anti-neutrinos and among neutrino flavors, is compatible with an unbroken power law, with a preferred spectral index of $2.87^{+0.20}_{-0.19}$ for the 68.3 % confidence interval. The techniques developed here will prove useful in the next generation of analyses, as more data is incorporated. These include a new approach to accounting for simulation statistical uncertainties in counting experiments and an improved method of computing the veto suppression effects on atmospheric neutrinos.

Acknowledgements

The body of work that follows tells a story of scientific achievement, highlighting the incremental progress others and I have made. However, it neglects something far more important, the people and community behind it all. Unfortunately, I have neither the time nor space to thank everyone properly. I have had incredible support from friends, family, mentors, communities of fellow scientists, among many others throughout this journey. Without their help, support, and constant positive influence, I would not be where I am now, nor would I be the person I am today.

Thank you, Albrecht, for your guidance, light-hearted perspective, and prudent advice. Your persistent input has helped me become a better scientist and communicator, despite my own stubbornness at times. There is still much for me to learn, but I have you to thank for setting me on the right track and keeping me from getting lost in the metaphorical woods. Carlos, Tianlu, Juliana, and Nancy, we've shared a long and challenging road bringing our analyses to fruition. Still, I hope you all remember some of the fun we had along the way. You all pushed me to be better in different ways, and it has been a blast having you as collaborators. I'm fortunate to have been surrounded by the fantastic people at WIPAC. I couldn't ask for a better bunch as friends and colleagues to accompany me through these years. Your attitudes, eagerness, and excitement have been absolutely contagious, and I have been humbled by your generosity and selfless drive to improve the institutions that shape our lives. I hope to see you all in the future, even as we depart from our shared experiences at WIPAC.

To my parents, Margaret and James, thank you for your unconditional love and support through all these years. You've been there for me at every turn, and I wouldn't have made it without you.

Contents

Chapter 1

Introduction

During the last decades, observations of low-energy extraterrestrial neutrinos have pushed forward our understanding of astrophysical environments and elucidated properties of neutrinos. These advancements stem from the fact that neutrinos only interact via the weak force, allowing them to escape dense astrophysical environments where they are produced and travel long distances to us unperturbed. These properties have enabled measurements of MeV neutrinos from the Sun, the closest detected extraterrestrial neutrino source, improving our understanding of the Sun's inner workings [1, 2] and have been pivotal in resolving the neutrino flavor-changing puzzle [3]. Similarly, the observation of neutrinos from supernova 1987A [4, 5], approximately 168,000 light-years away in the Large Magellanic Cloud, has provided irreplaceable information for supernovae physics [6], constraining neutrino properties [7, 8, 9, 10], and fundamental physics [11, 12].

Enabled by these unique neutrino properties, the study of high-energy extraterrestrial neutrinos is opening many new opportunities for discovery. These neutrinos can travel distances that far exceed the distance traversed by neutrinos observed in 1987. They are expected to be produced in high-energy hadronic processes in our Universe either directly from hadronic decaying or decaying charged leptons produced in hadronic interactions [13]. Regions of charged-particle acceleration are then prime candidates as high-energy neutrino sources. The observation of EeV cosmic rays indicates that objects of large size or high magnetic field strength are accelerating charged particles to high energies, potentially narrowing the search for neutrino sources to a subclass of objects [14, 15]. Measurements of the cosmic ray, gamma ray, and neutrino fluxes show similar energy content despite their disparate energy regimes, as demonstrated by recent measurements in Fig. 1.1. Despite this information and a wealth of cosmic-ray observations, the sources of ultra-high-energy cosmic rays are an unresolved mystery [16]. Thus, much like solar neutrinos that can escape their birthplace, these neutrinos are an indispensable probe for cosmic-ray sources that provide insight into the long-standing problem of cosmic-ray origin as they can escape dense environments to reach us unperturbed. By studying

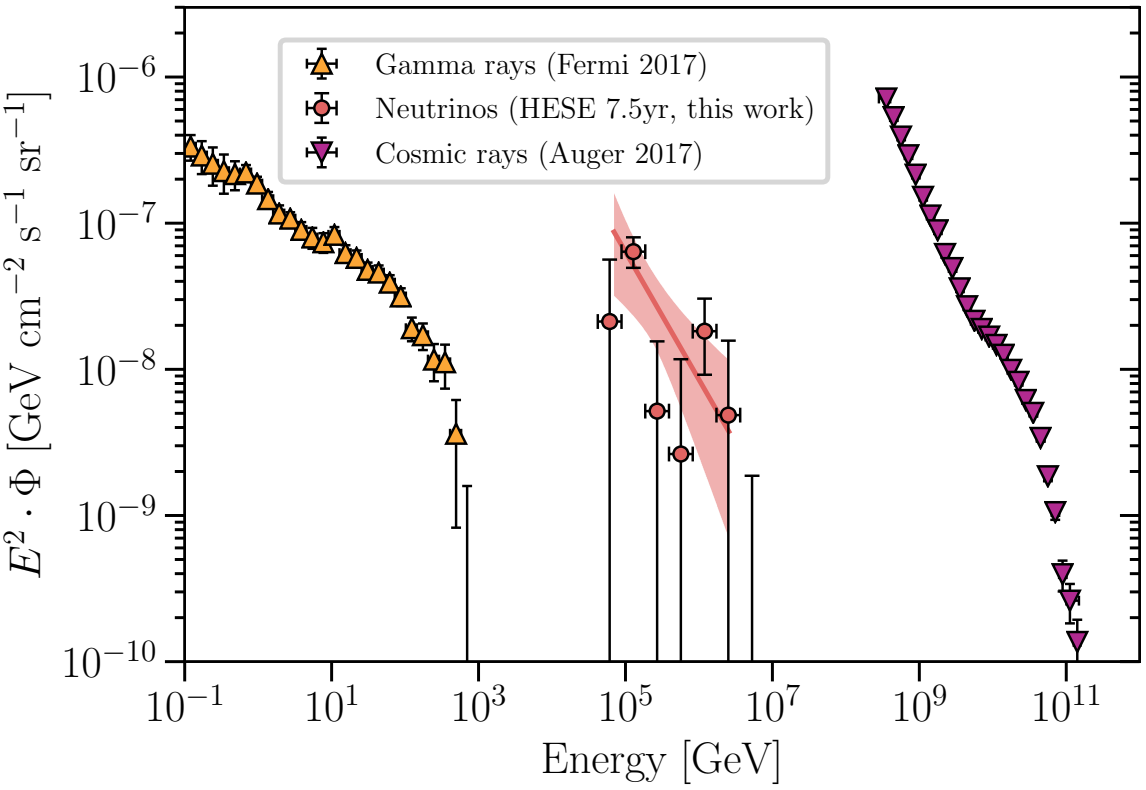

Figure 1.1: ***High-energy fluxes of gamma rays, neutrinos, and cosmic rays.*** The segmented power-law neutrino flux, described in Section 8.1.1.5, obtained in the analysis described in this work, is shown as red circles. The single power-law assumption, described in Section 8.1.1.1, is shown as the light red region. The high-energy gamma-ray measurements by Fermi [17] are shown in orange, while the very-high-energy cosmic-ray measurements by the Pierre Auger observatory [18] are shown as purple data points. The comparable energy content of these three fluxes is of particular interest in the investigation of cosmic-ray origin.

their flux and energy spectrum, we can constrain the acceleration environments that produce these neutrinos.

High-energy astrophysical neutrinos are also powerful probes of new physics [19]. This is in large part because neutrinos are charged under flavor [13, 20, 21], unlike other cosmic messengers. New flavor-non-trivial interactions can arise from breaking of space-time symmetries [22, 23], secret neutrino interactions with the cosmic-neutrino background [24, 25, 26, 27], flavored dark-matter neutrino interactions [28, 29, 30], or other non-standard interactions [31]. Beyond flavor, the very-long distances traversed by high-energy astrophysical neutrinos can be used for accurate time-of-flight [32] and neutrino-flux spectral distortion [33] measurements. Their high energy also provides a unique way to study very-heavy decaying and annihilating dark matter, whose other standard model products will not reach Earth [34]. Finally, these neutrinos can also probe the high-energy neutrino-nucleon cross section [35, 36, 37, 38, 39]. Such a measurement is of interest due to the

possibility of observing gluon screening [40] which could reduce the cross section at the highest energies [41, 42, 43], or new physics phenomena, *e.g.* low-scale quantum gravity [44], leptoquarks [45, 46, 47, 48, 49, 50, 51], sphalerons [52, 53], and micro black hole production [54, 55]; see [56] for a recent review.

The IceCube Neutrino Observatory has firmly established the existence of high-energy astrophysical neutrinos. Northern Sky measurements of through-going muon tracks [57, 58], all-sky measurements using events with interaction vertices contained in the detector fiducial volume [59, 60, 61, 62] such as High-Energy Starting Events (HESE), and additional studies extending to lower energies with contained cascades [63, 64] have all contributed to the characterization of the astrophysical neutrino flux. However, the energy spectrum, directional distribution, and composition of this neutrino flux are still too poorly constrained to differentiate between many astrophysical scenarios.

Figure 1.2 shows how the different components of the analysis are related. The following chapters cover these components except for Monte-Carlo generation and detector simulation. Chapter 2 provides some brief background on neutrino interactions, cosmic rays, and cosmic-ray air-shower products relevant for IceCube. These physical phenomena set the stage for what we must consider in the analysis, and information from this chapter is relevant for many of the components given in Fig. 1.2. Chapter 3 explains relevant details of the IceCube detector and discusses some inherent detector systematics relevant to the astrophysical neutrino flux measurement. Chapter 4 outlines the event selection used for these measurements, the calculation of relevant backgrounds, and a new technique for computing the background suppression effects of the selection. The new technique for computing this background suppression effect was developed in collaboration with Carlos Argüelles, Sergio Palomares-Ruiz, Logan Wille, and Tianlu Yuan. The content relating to this technique is reproduced from the corresponding publication [65]. Chapter 7 explores the statistical techniques needed to analyze the data and make inferences about the astrophysical flux parameters. This chapter also develops a new technique to account for the uncertainties inherent when using simulation with a limited sample size, a hurdle that must be overcome because of computationally limited background estimation. This technique was developed in collaboration with Carlos Argüelles and Tianlu Yuan; some content of this chapter is reproduced from the corresponding publication [66]. Finally, Chapter 8 provides the analysis description and discussion of results. Information from this chapter will also be reproduced in a forthcoming publication.

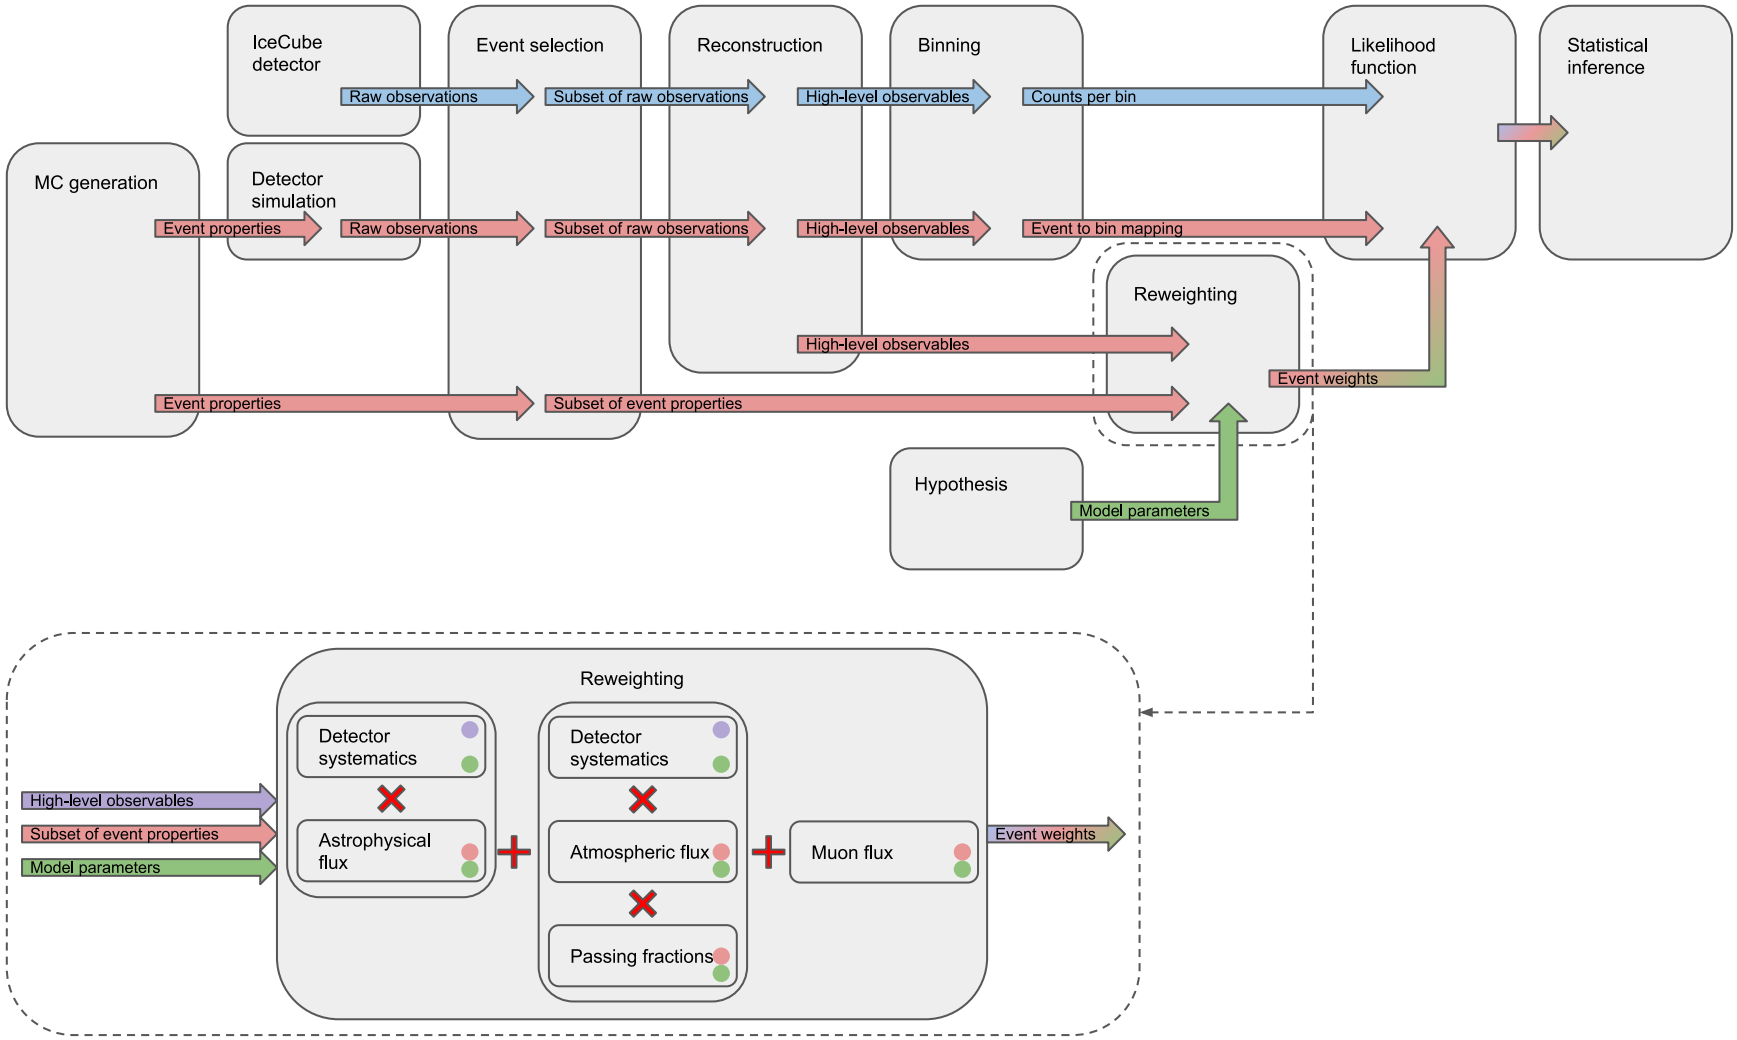

Figure 1.2: ***Anatomy of an analysis.*** This diagram shows the flow of information through the different components of the analysis. The upper portion begins with Monte-Carlo generation and the IceCube detector operation. These mark the beginning of data and simulation, which follow parallel paths until they are compared in the likelihood function. Along the way, model parameters are introduced in order to weight the simulation to a physical hypothesis. Reweighting the simulation involves several components whose relationship is described in the lower portion of the diagram. Contributions from three classes of fluxes are added together to get the total event weight. The atmospheric and astrophysical neutrino weighting both involve detector systematics, and the atmospheric weighting also includes the calculation of passing fractions.

Chapter 2

Background information

2.1 Cosmic rays and air showers

Cosmic rays are high-energy protons and atomic nuclei that propagate through space. Around 90 % of cosmic rays we observe are protons, with a smaller proportion of alpha particles, and less than 1 % which are heavy nuclei [67]. Most cosmic rays originate from outside our solar system but within our galaxy. Increased solar activity has been shown to inversely correlate with the majority of cosmic ray observations, which can be understood as an exclusion from our stellar neighborhood by the magnetized solar wind. The bulk of cosmic rays have gyro-radii small enough to be contained in the Milky-Way by galactic magnetic fields. Additionally, gamma-ray measurements at the energies relevant for neutral pion decay show a spatial distribution consistent with a galactic population of cosmic rays. If the bulk of cosmic rays were of extra-galactic origin, we would expect neutral pions to be produced in abundance within intergalactic space, and this is not supported by the gamma-ray measurements [68]. In fact, the peak of the gamma-ray spectrum from neutral-pion decay is observed with galactic gamma-ray data in the $\sim 70\,\mathrm{MeV} - 40\,\mathrm{GeV}$ energy range, and accurately predicts the measured cosmic ray spectrum in the $\sim 200\,\mathrm{MeV} - 100\,\mathrm{GeV}$ energy range [69].

Observations of the relative proportions of primary and secondary nuclei suggest that cosmic rays, on average, traverse distances thousands of times greater than the galactic disk's thickness before reaching Earth. This proportion of secondary nuclei decreases with energy, meaning that higher energy particles traverse less material on average before reaching us. We expect this trend if acceleration and propagation are separate processes that occur on different timescales. For this reason, we believe cosmic rays to be accelerated by dense astrophysical objects and then subsequently diffused by propagation through the interstellar medium. Simultaneous propagation and acceleration also fail because the timescale to accelerate cosmic rays to their observed energy is greater than the age of the galaxy in such a scenario. While the acceleration mechanisms

and particular sources of cosmic rays are unknown, we can narrow the candidates to astrophysical objects that are large or have strong magnetic fields.

A small proportion of cosmic rays have been observed in excess of 1×10^{19} eV. At these energies, cosmic rays cannot be sufficiently bound by the galactic magnetic field to produce an isotropic flux. Thus, cosmic rays at this energy scale must either point towards nearby accelerators or arrive roughly isotropically from extra-galactic sources. This highest energy population's spectrum and spatial distribution has been shown as distinct from the lower-energy population and likely to be of extragalactic origin. The sources and acceleration mechanisms of these extra-galactic cosmic rays are unknown to date. They remain a subject of significant study and interest, particularly because much more extreme environments are required for acceleration to these energies. A simple energetic consideration tells us that the maximum energy achievable through acceleration by an object with a size R and magnetic field strength B is $E < ecBR$. Tighter constraints on the maximum energy can be derived when considering diffusive shock acceleration specifically. This limit is adequately summarized by the Hillas diagram, which compares the size and magnetic field strength of various objects to the upper limit on cosmic ray energy. Figure 2.1 shows this limit for 1×10^{20} eV protons considering various acceleration scenarios and energy losses. Clearly, there are many candidates for cosmic-ray production, and narrowing the field may require another approach.

Cosmic rays can be observed interacting in the Earth's atmosphere, as they produce extensive particle air-showers whose constituents and products can be observed with a wide range of techniques. The charged particle products of these interactions have been observed as early as 1912 [71]. These air showers are of particular concern to neutrino detectors as they produce high-energy muons and neutrinos in relative abundance, both of which can be observed by neutrino detectors even with significant shielding and overburden.

When cosmic rays interact with the atmosphere, some nuclear fragments can be produced, but hadronization occurs due to the high energy of the incident cosmic ray and initiates a hadronic particle shower. As part of the hadronic shower, pions and kaons are produced in abundance. These subsequently decay or interact in the low-density atmosphere. Leptonic and semi-leptonic decays of charged pions and kaons that produce muons or neutrinos have a large decay branching fraction, resulting in an abundance of muons and neutrinos in any air shower that begins hadronically. Decays of charged pions ($\pi^+ \to \mu^+\nu_\mu$), K_L^0 ($K_L^0 \to \pi^\pm e^\mp \nu_e$, $K_L^0 \to \pi^\pm \mu^\mp \nu_\mu$), and K^+ ($K^+ \to \mu^+\nu_\mu$, $K^+ \to \pi^0 e^+ \nu_e$, $K^+ \to \pi^0 \mu^+ \nu_\mu$) all have direct contributions to the muon and neutrino fluxes; secondary contributions from pion production in kaon decay are also present, particularly for K_S^0. The yield of muons and neutrinos is highly zenith dependent. This dependence is a result of the interplay between the probability of meson interaction and meson decay. Unlike those that decay, pions and kaons that interact with the atmosphere or Earth will not produce neutrinos, and decay is much more likely in the low-density atmosphere. Between the zenith and horizon, the average path length of

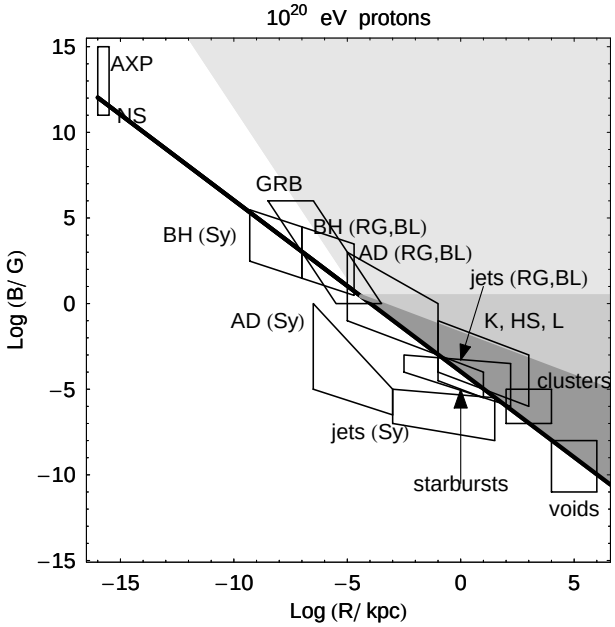

Figure 2.1: ***Hillas criterion for*** $\mathbf{1 \times 10^{20}\,eV}$ ***protons.*** This figure is reproduced here from [70]. The black line denotes the region above which objects meet the Hillas criterion. The shaded regions show regions allowed for different acceleration scenarios: light gray for one-shot acceleration, gray for one-shot acceleration with synchrotron losses, and dark gray for one-shot acceleration and diffusive shock acceleration. Boxes denote the parameters of various astrophysical objects: the central parsecs (AD) of active galaxies (low-power Seyfert galaxies (Sy) and powerful radio galaxies (RG) and blazars (BL)), relativistic jets, knots (K), hot spots (HS) and lobes (L) of powerful active galaxies (RG and BL); non-relativistic jets of low-power galaxies (Sy); starburst galaxies; gamma-ray bursts (GRB); galaxy clusters and inter-cluster voids; immediate the neighborhood of neutron stars (NS), anomalous X-ray pulsars and magnetars (AXP).

mesons through the atmosphere increases with zenith angle, thereby increasing the yield of neutrinos for the same flux of cosmic rays. For this reason, neutrino production from pions and kaons is peaked at the horizon, as shown in Fig. 2.2.

Neutrinos can also be produced by charmed hadrons present in air showers. These hadrons are much shorter-lived than the pion and kaon, and almost invariably decay before interaction. As the interaction probability for these charmed hadrons is so low, the production of neutrinos from them has practically no dependence on the zenith angle. This flux of neutrinos is still unobserved, although its spectrum has been predicted ($\sim E^{-2.7}$) and constraints placed on its normalization. As this atmospheric flux is effectively isotropic, some have suggested that it may be confused for the astrophysical flux. However, the astrophysical flux can be distinguished by its energy spectrum and the effect described in Section 6.1.2.

Muons produced in these air showers reach the ground where they can be detected, and often penetrate many kilometers into the Earth's surface. Underground neutrino detectors can, therefore, be sensitive to muon backgrounds despite large overburdens of rock, water, or ice. On the other hand, Neutrinos have a small interaction cross such that the effective area for neutrinos is approximately 10^6 times smaller than that

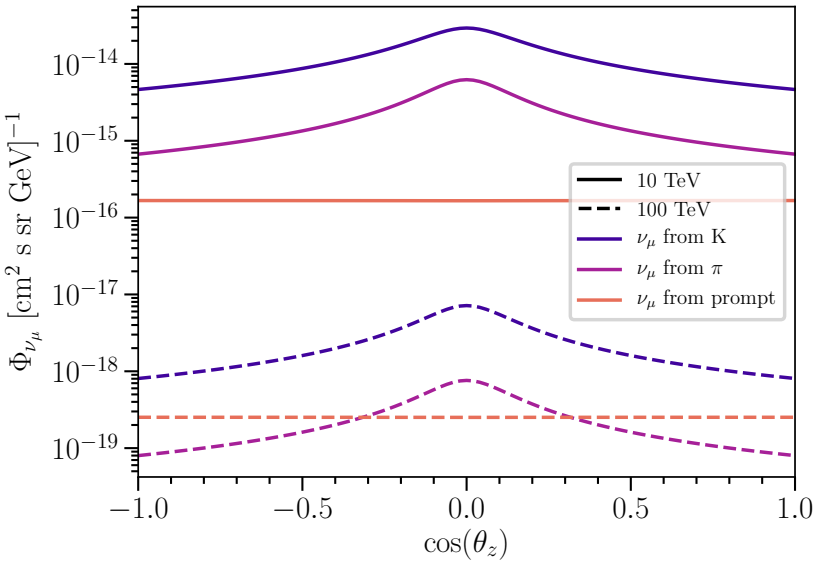

Figure 2.2: ***Production of atmospheric muon neutrinos.*** The atmospheric muon neutrino flux at the Earth's surface is plotted as a function of the zenith angle measured with respect to the IceCube detector center. At these high energies above several TeV, the production of muon neutrinos is dominated by kaon decays. The kaon and pion production is peaked at the horizon, where the path of these particles includes the longest stretch of atmosphere before reaching the Earth. This larger path length results in a higher decay probability for these mesons. Production from the prompt decay of charmed hadrons has no angular dependence as the decay length of these hadrons is short with respect to the atmosphere for all directions.

for similar energy muons in the few TeV energy regime. With this small interaction cross section, neutrinos are likely to pass through the Earth without interacting. Neutrino detectors are then able to observe the atmospheric neutrino flux coming from all directions. Chapter 6 explores in detail how these backgrounds are estimated for IceCube.

2.2 Neutrino interactions and detection

Neutrinos are the only neutral leptons in the standard model of particle physics. Although fundamentally, neutrinos only interact through gravity and the exchange of weak bosons, there exists a wide range of processes that dominate the relevant physical behavior of neutrinos at different energy scales. Such interactions include nuclear capture, inverse beta-decay, quasi-elastic scattering, resonant particle production, coherent elastic scattering, deep inelastic scattering (DIS), and ultra-high energy interactions [73, 74]. Above TeV neutrino energies, only two known processes remain relevant for detection: DIS and resonant W production. Deep inelastic scattering refers to processes that probe the fundamental components of hadrons. For neutrinos, this means the exchange of a weak boson with a quark. The momentum imparted to the quark will produce a hadronic cascade of secondary particles. The details of the lepton side of the interaction depend heavily on

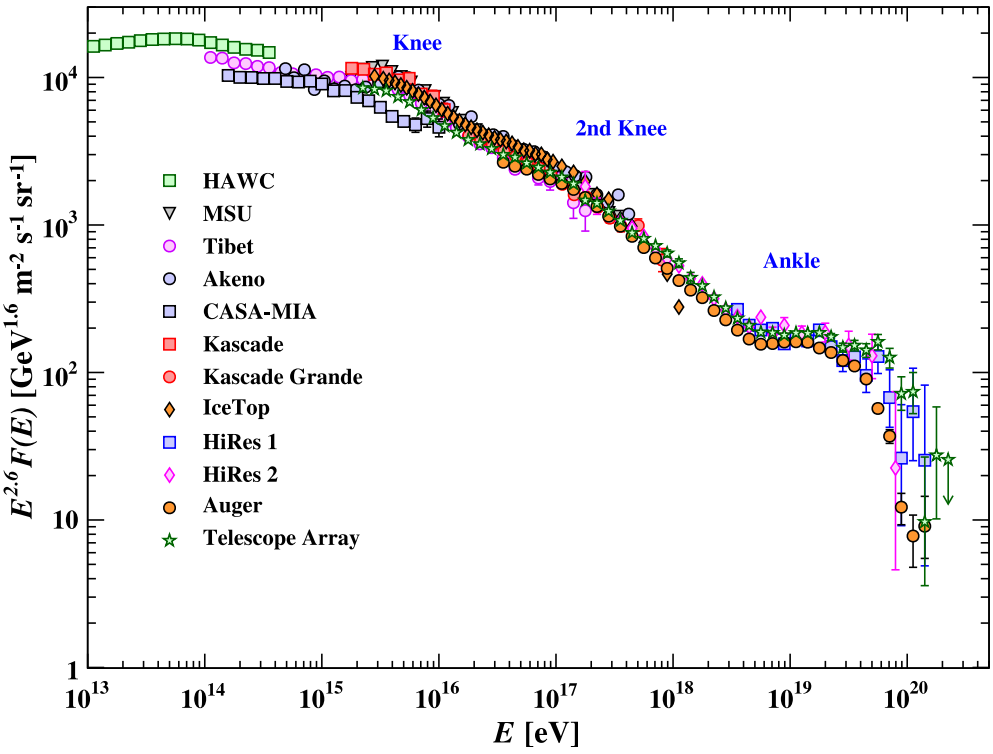

Figure 2.3: ***Cosmic Ray Spectrum.*** This figure is reproduced here from [72]. The measured flux of cosmic rays is plotted from a variety of cosmic ray observatories. The flux is multiplied by $E^{-2.6}$ to highlight the different spectral features. Above 4×10^{18} eV the harder spectrum of extra-galactic cosmic rays is apparent before the spectrum cuts off.

the species of incident neutrino and weak boson exchanged. We can divide these DIS interactions into two categories based on the weak boson exchanged. Interactions involving the exchange of a Z^0 are referred to as "neutral current" (NC), and those exchanging a W^+ or W^- are referred to as "charged current" (CC). Modern techniques for observing neutrino interactions rely on detecting the charged particle products of the initial interaction. As a consequence of this, the observable energy can be very different for NC and CC events. Both interactions produce a hadronic cascade, but on the leptonic side of the interaction, NC events have an outgoing neutrino (not observable), while CC events have an outgoing charged lepton (observable). These two interactions are shown in Fig. 2.4.

The third interaction relevant above TeV neutrino energies is the resonant production of a W boson, otherwise known as the Glashow resonance (GR) [75]. In matter on Earth, this process occurs when an anti-electron neutrino combines with an atomic electron to produce an on-shell W^+ as shown in Fig. 2.6. If we consider the rest frame of the electron, then this resonance occurs at a neutrino energy of 6.3 PeV. For atomic electrons we should consider the rest frame of the atom, and in this case there is a Doppler broadening of the resonance of $\sim 20\,\%$ due to the orbital motion of the electrons [76]. In practice, this broadening is

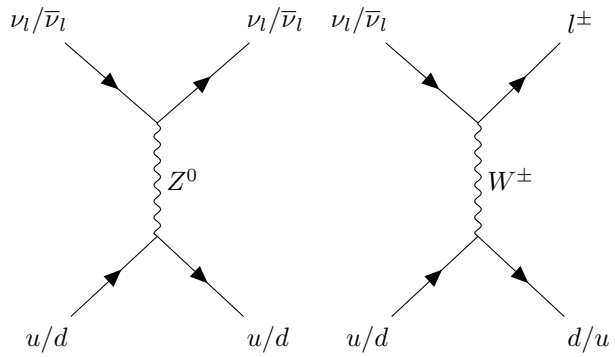

Figure 2.4: The neutrino deep inelastic scattering processes NC (left) and CC (right) in matter is shown in the figure above for interactions with nucleon component quarks. In both cases, significant momentum can be imparted to the outgoing quark, which will result in the production of a hadronic particle cascade. In NC interactions, only the hadronic cascade may be detectable as the interaction product is a neutrino, which is unlikely to undergo another interaction within the detection medium. Interactions of the CC variety, on the other hand, produce a charged lepton in addition to the hadronic cascade. This charged lepton can also be detected if it receives enough energy.

small compared to the energy resolution of modern neutrino detectors that have access to this energy scale, and any further broadening from thermal motion will be even smaller. The production of a W^+ and its subsequent decay can result in either a hadronic shower similar to a NC interaction, or a leptonic final state similar to a CC interaction. These two possibilities correspond to the hadronic and leptonic decay modes of the W, respectively.

The cross section of these processes is shown as a function of the neutrino energy in Fig. 2.5. The growth of the cross section with energy counteracts the falling neutrino spectrum to only a small degree. Despite the large cross section, only a few Glashow resonance events are expected in 10 years of IceCube detector operation due to the comparably low neutrino flux at these energies. Absorption of neutrinos in the Earth becomes significant above $\sim 1\,\text{TeV}$, and is notably peaked near the 6.3 PeV resonance energy.

Through either a NC interaction or the hadronic decay of a W boson, neutrinos can induce a hadronic shower. In a hadronic shower, both charged hadrons and leptons are produced, which can be detected through well-established methods. Charged current interactions produce a hadronic shower by imparting momentum to a quark that is then hadronized, although the charged lepton produced in the interaction is also detectable and can significantly alter the event's detection signature.

Through either a CC interaction or the leptonic decay of a W boson, neutrinos can produce a detectable charged lepton, although the detection signature differs depending on the charged lepton's flavor. Focusing on dense detection media like ice or water, the three flavors of charged leptons' detection signatures are as follows. High energy electrons and positrons immediately interact with the detection media to initiate an electromagnetic cascade where charged leptons and high energy photons are alternately produced by one

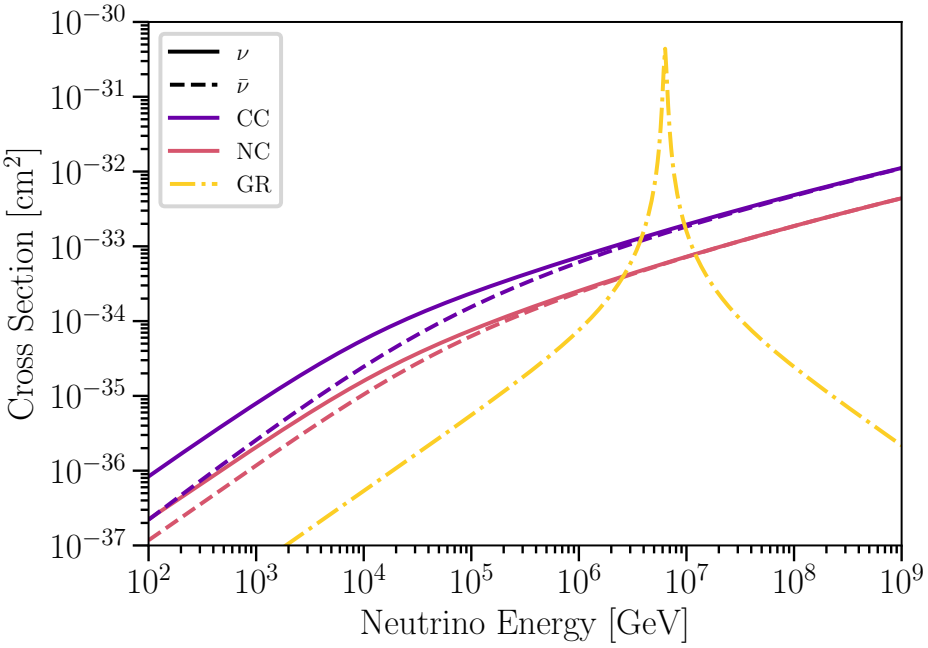

Figure 2.5: ***Neutrino interaction cross sections.*** Charged current (CC), neutral current (NC), and Glashow resonance (GR) cross sections are shown as a function of neutrino energy. The CC and NC cross sections come from [77], while the GR cross sections are computed using the forms in [78] and corrected for the Doppler shift of electrons in the molecular orbitals of water [76].

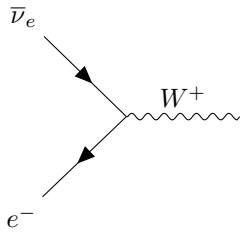

Figure 2.6: The production of an on-shell W^+ boson through the combination of an anti-electron neutrino and electron. This resonant process occurs at neutrino energies around 6.3 PeV.

another. This electromagnetic cascade develops over ~ 10 radiation lengths (about $\sim 5\,\mathrm{m}$ in total) at 10 TeV with a lateral extension of $\sim 20\,\mathrm{cm}$ (twice the Molière radius), expanding within the dense detection medium, and has an extreme directional bias because of the momentum of the first charged lepton.

Muons from high energy neutrino interactions do not interact as readily as electrons and positrons due to their larger mass. Instead, muons can travel several kilometers in dense media before losing enough energy to quickly decay. Along their entire path length, muons lose energy by interacting with the detection medium. These "energy loss" interactions include ionization, electron-positron pair production, bremsstrahlung, and photo-nuclear interactions. Although these processes are highly stochastic, the average energy loss of muons approximately follows $-dE/dx = a + bE$ where a is determined by the ionization energy losses and b is defined by the other processes. In general, a and b are both functions of muon energy E, but this linear approximation where a and b are constant holds locally as both are slowly varying as a function of E. For muon energies above 1 TeV, the so-called "radiative" term bE dominates the average energy losses. Figure 2.7 shows the energy loss rate for the different processes. Once the radiative losses have taken over above ~ 1 TeV, the losses grow exponentially with energy.

The strong dependence of the energy loss rate on muon energy means that the energy lost while traversing the detector can be used to estimate a muon's energy. For simple observables like the total energy lost over 1 km the stochastic energy losses introduce large variations between muons of the same energy, reducing their power as a proxy for the muon energy. Figure 2.8 shows the energy of muons lost within 1 km of ice; a distribution with very long tails. In practice the muon energy can currently be determined to within a factor of 2, however improved techniques that take advantage of more detailed information may achieve a resolution as small as 10 % in the future.

Taus produced through a CC interaction, or the leptonic decay of a W boson are also detectable. The short decay length of a tau, 50 m/PeV, means it is likely to decay very close to the neutrino interaction vertex. Taus decay hadronically to produce a hadronic shower with a branching ratio of 64.79 %. The tau's leptonic decay modes are decay to a charged lepton and corresponding neutrino of either electron or muon flavor. In the electron case, an electromagnetic shower results; whereas a far traveling muon is produced in the muon case. For taus produced via the decay of a W boson, the event is indistinguishable from taus produces via CC and NC interactions other than by the resonance energy at which this process occurs. Although a tau may traverse tens of meters, which is detectable by IceCube, the energy losses are still negligible over these distances.

From the wide variety of methods for detecting charged particles, water Cherenkov detectors are most common for detecting neutrinos in this high-energy regime above 1 TeV. Water Cherenkov detectors have the advantage that their detection medium is both abundant and inexpensive, a major motivator in the design of

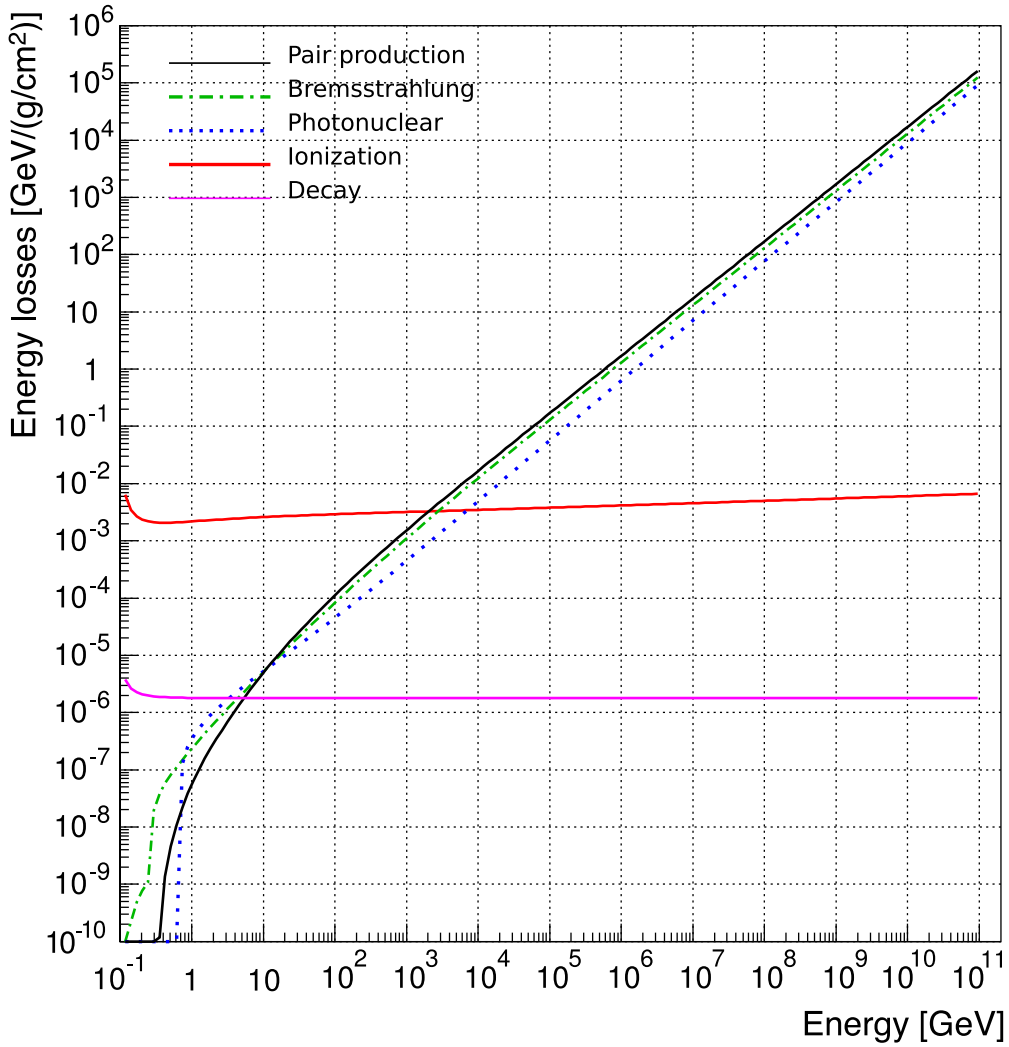

Figure 2.7: ***Muon Energy Losses*** The muon energy loss rate computed for water is plotted as a function of the muon energy for the five different loss processes. This figure is reproduced here from [79]. Radiative energy losses dominate above $\sim 1\,\text{TeV}$. Although it is a smaller contribution to the energy losses, photo-nuclear interactions are the main source of uncertainty above $\sim 1\,\text{TeV}$.

IceCube. Cherenkov radiation results from charged particles propagating through a dielectric medium faster than the phase velocity of light. As charged particles pass through the dielectric material, the ionization of the medium induces the emission of light. For particles faster than the phase velocity of emitted light, the emission forms a conical coherent wave front at a well-defined angle to the particle's path.

These detectors use Cherenkov radiation to detect particles, but the high-energy secondary charged leptons themselves contribute only a small fraction of the Cherenkov photons. The tertiary particle showers produce most of the Cherenkov photons that these leptons give rise to when they lose energy in the detection medium. This increased light-yield allows detectors to be sparsely instrumented while maintaining detection efficiency and reconstruction quality.

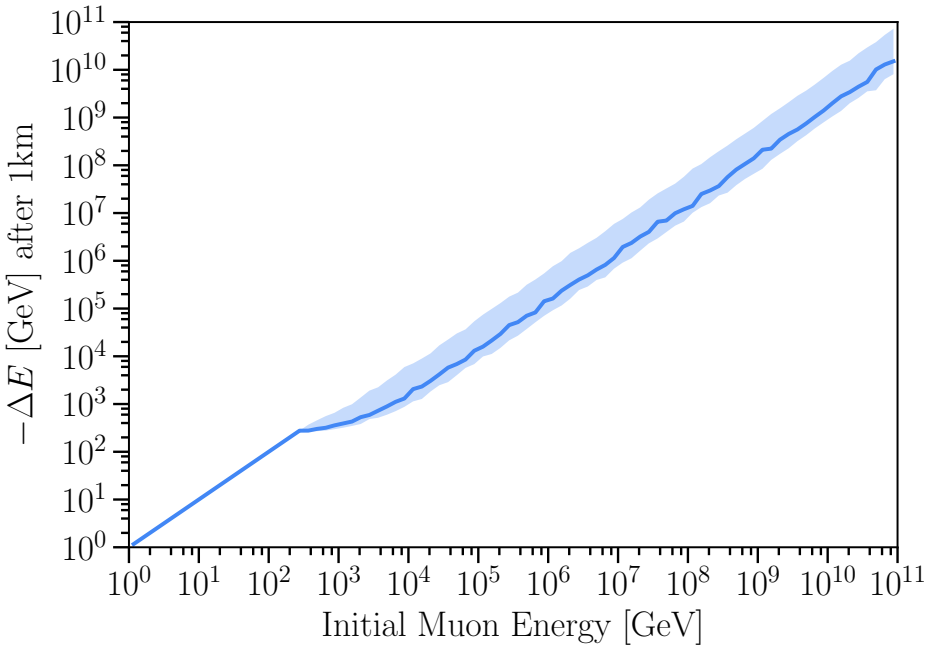

Figure 2.8: ***Muon energy loss in* 1 km *of ice.*** Muons are propagated through 1 km of ice, and their lost energy recorded. For each initial muon energy, the MAP and 90 % HPD region are plotted here. There is a strong correlation between the initial energy and energy lost. However, the distribution has very long tails, making a precise determination of the muon energy based on this observable impossible. Below a few hundred GeV muons lose all of their energy within 1 km of ice, narrowing the distribution of ΔE.

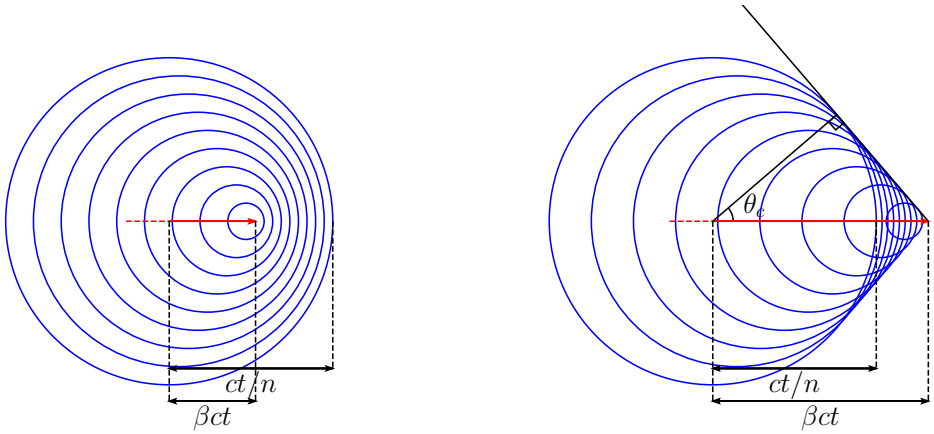

Figure 2.9: ***Cherenkov radiation.*** This diagram shows the photon wavefronts as a charged particle moves through a dielectric medium. On the left for $\beta < 1/n$ and on the right for $\beta \geq 1/n$. The Cherenkov angle θ_c is shown on the right and describes a ray perpendicular to the coherent wavefront. In the Antarctic ice the Cherenkov angle $\theta_c = 40.72°$.

Water Cherenkov detectors often make use of photo-multiplier-tubes (PMTs) to detect Cherenkov photons. PMTs have a thin photo-cathode held at a high-voltage differential with respect to an anode; this allows the production and acceleration of a photo-electron when photons pass through the photo-cathode. Accelerated photo-electrons hurtle towards an amplification stage composed of many dynodes, each held at a large voltage difference to the adjacent dynodes. This setup allows a single photo-electron to produce many secondary electrons upon interaction with a dynode, starting an exponential cascade of electrons across the dynode stages. The resulting cascade of electrons is amplified with respect to the original signal enough to be detected electronically as a voltage change. In this way, PMTs can be sensitive to single photons, provided one can reach the photo-cathode and produce a photo-electron.

Chapter 3

The IceCube South Pole Neutrino Observatory

3.1 The IceCube detector

IceCube is a gigaton-scale Cherenkov detector embedded in the Antarctic ice [80] at the geographical South Pole [81]. The detector consists of photo-multiplier tubes (PMTs) and digitization electronics contained within glass pressure housings; this single unit of the detector is referred to as a "Digital Optical Module" (DOM) [82]. The basic operation of the IceCube PMTs remains the same as in other experiments, but special care has been taken to design the DOMs to support their operation in the extreme environment of the Antarctic ice-sheet. The PMT and associated electronics are held inside the DOM's glass pressure housing, designed to withstand a sustained 250 bar and up to 690 bar temporarily. A mu-metal cage surrounds the PMT to reduce the ambient magnetic field's effect on the PMT. Without this shielding, we would expect $5-10\,\%$ lower collection efficiency, poorer charge resolution, and gain variations of up to $20\,\%$. The "digitization" step performed within the sensor modules was an essential advancement over IceCube's predecessors, as all sensors are located more than 1400 m below the glacier surface, preventing accurate analog readout of the detector from the surface. The high voltage supply is also located within the DOM for similar reasons. Figure 3.1 shows a DOM and its components.

DOMs are arranged into vertical columns connected by cabling, called "strings." Each string consists of 60 DOMs spaced ~ 17 m apart. There are 86 strings in total, 78 of which are configured as described previously and arranged horizontally in a hexagonal grid with ~ 125 m spacing. The remaining eight strings comprise the DeepCore sub-array, which has DOMs arranged with smaller vertical and horizontal spacing within a hexagonal cell of the main array, and has PMTs with higher quantum efficiency [83]. The DOMs record discretized charge and timing information if a PMT readout voltage corresponding to at least 0.25 photo-electrons (PE) is observed. The discretized information is referred to as "hits", where each hit has a

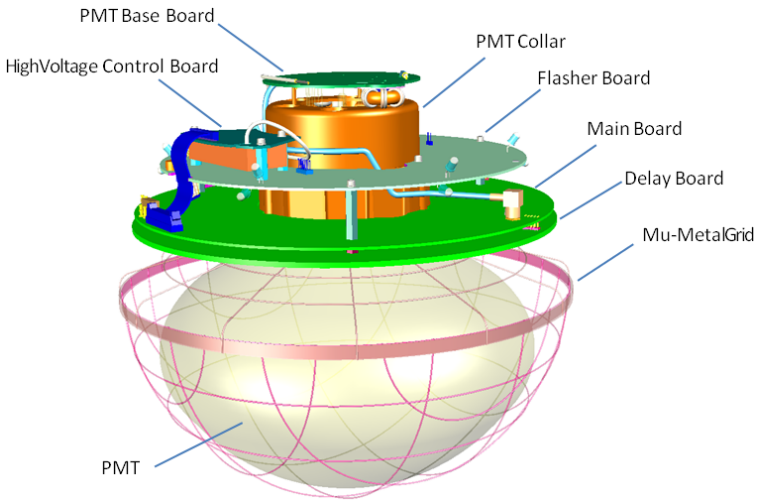

Figure 3.1: ***IceCube Digital Optical Module.*** This figure is reproduced from [81].

recorded charge and time.

Although IceCube was designed and constructed before discovering astrophysical neutrinos, elements of its design are optimized for their detection. Previous attempts to observe the astrophysical neutrino flux proved elusive and suggested that a kilometer-scale detector would be necessary to observe the flux [84, 85, 13], and motivating IceCube's scale. The spacing of strings is optimal for detecting events of several TeV. The denser vertical spacing of DOMs provides improved zenith angle resolution so that down-going background muons may be better identified. The detector's timing resolution is at the nanosecond scale to make reconstruction over hundreds of meters possible. Finally, DOMs face downwards to increase the amount of detectable light from up-going neutrino events. Figure 3.2 shows the layout of IceCube from the side, with each line indicating a string and the circles denoting each DOM.

In IceCube, neutrinos are detected by observing the Cherenkov light emitted by particles produced due to the neutrino interacting with the ice or the bedrock below the detector. Neutrino neutral-current (NC) interactions initiate a hadronic shower that appears as a "cascade"-like morphology in the detector, and an out-going neutrino that is not observable. Here, "cascade"-like refers to highly localized energy deposition and roughly spherical light emission. Charged-current (CC) interactions produce a hadronic shower at the site of the neutrino interaction and an out-going charged lepton. When the CC interaction is triggered by an electron neutrino (ν_e), an electron (e) is produced, and its subsequent interaction starts an electromagnetic shower. This interaction is observable as a cascade-like morphology, which is indistinguishable from a NC cascade. Suppose the incident particle is instead a muon neutrino (ν_μ). In that case, a muon (μ) is produced in the interaction which, at the energies we are concerned with, will generally traverse several kilometers and exit the kilometer-scale detection volume while depositing energy stochastically [72]. IceCube observes

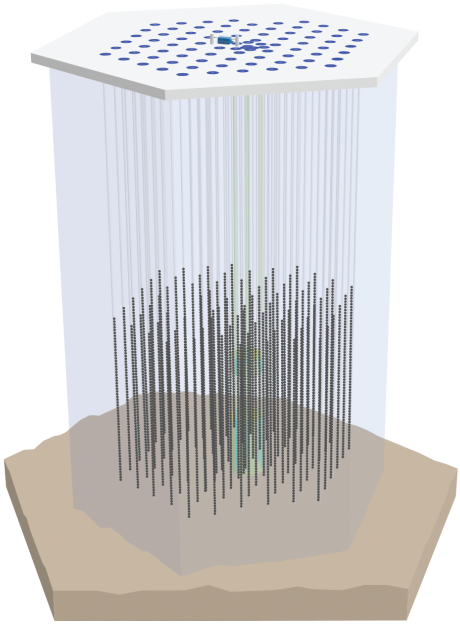

Figure 3.2: ***IceCube detector configuration.*** The 86 strings are shown as vertical lines, with each DOM denoted by a small circle. Blue circles indicate IceTop water Cherenkov tanks at the top. The green region highlights the DeepCore sub-array. Extending down to a depth 50 m is the firn layer, represented by the thin white section at the top of the diagram. Below that is 1400 m of ice before the detector begins at a depth of 1450 m. The detector itself is ~ 1 km in height, terminating at a depth of 2450 m. The bedrock begins at a depth of ~ 2820 m.

this muon as a "track"-like morphology, where "track"-like refers roughly to a long and narrow cylinder of energy depositions and light emission pointed in the same direction. Finally, if the interaction is due to a tau-neutrino (ν_τ) the initial hadronic interaction is accompanied by an out-going τ, which subsequently decays with a mean decay length of $\sim 50\,\mathrm{m/PeV}$. When the initial interaction and subsequent decay of the τ can be distinguished, such events are classified as "double cascades" [86, 87]. The morphologies induced by a tau-neutrino interaction are often trackless, but the τ leptonic decay to $\nu_\tau\mu\bar{\nu}_\mu$ is not negligible ($\sim 17\,\%$ [88]) and in this case, a track can also be observed. By distinguishing between these three "morphologies" (cascades, tracks, and double-cascades), IceCube is sensitive to the neutrino flux's flavor composition.

Astrophysical neutrinos are expected to arrive in roughly equal amounts of the three neutrino flavors [89, 90, 23, 91, 92, 93], be distributed isotropically across the sky, and dominate the observed neutrino flux above $\sim 100\,\mathrm{TeV}$ [58]. At energies above $\sim 100\,\mathrm{TeV}$ IceCube has obtained results consistent with this expectation so far [94, 95, 96, 97, 98].

3.2 Detector systematic uncertainties

To imagine the possible systematic uncertainties for an analysis it is helpful to consider the chain of events that leads to our observations. As this section is concerned with detector systematic uncertainties, we start immediately after the neutrino interaction as charged particles produce Cherenkov photons. Once a photon is produced, it propagates through the ice, influenced by the varying optical properties. A photon will undergo many scatterings and may potentially be absorbed before reaching an optical module. Near the optical modules, the ice properties are entirely altered by the melting and refreezing that occurred during drilling and deployment. After passing through this local ice, which may involve additional scatterings, a photon may reach the PMT surface or be obstructed by many other components. From any incident photons, the PMT itself facilitates photo-electrons' production with a quantum efficiency of $\sim 25\,\%$. After producing a photo-electron, the signal is amplified with a gain of 10^7 before delay and digitization of the signal through a variety of channels.

The behavior of the entire process after producing a photo-electron is well understood to the degree that it does not significantly affect high-energy measurements that rely on many observed photo-electrons. Section 4.2 discusses one of the last modifications to the calibration of processes after photo-electron production that may affect high-energy analyses.

Before the production of photo-electrons, there remain many possible systematic uncertainties in the photon propagation and detection. These systematic uncertainties can be organized as arising from either incomplete knowledge of the ice properties or other details of the detector that affect the response. Ice properties are separated into global ice effects that were present before deployment – such as anisotropy, scattering length, layering of dust concentration, absorption of photons in the bulk ice, and depth dependence of these properties – and local ice properties that were introduced during deployment, *i.e.* effects of the refrozen ice surrounding the DOMs [99].

Below, we examine these systematics' effect on the observables described in Chapter 5 after the event selection in Section 4.1 is applied.

As far as local effects go, the air bubbles introduced in the drilling process and concentrated in the center of the hole during refreezing increase the scattering of light, particularly in the vertical direction [81]. Some additional effects may be present due to the relative orientation of the module, hole-ice, and cable. Uncertainties in the optical module light acceptance and local ice effects can be modeled with three parameters

- DOM efficiency ($\epsilon_{\texttt{DOM}}$),
- head-on efficiency ($\epsilon_{\texttt{head-on}}$), and

- lateral efficiency ($\epsilon_{\texttt{lateral}}$).

The first parameter is an overall change in the efficiency of all the DOMs in the detector, with respect to the individual baseline of each DOM, and encapsulates many potential systematic effects. The last two parameters are part of a parameterization of the efficiency's angular dependence, which depends most strongly on local effects [81, 100, 101]. In this parameterization the relative efficiency is

$$A(\eta) = 0.34(1 + 1.5\eta^3/2) + \epsilon_{\texttt{lateral}}\eta(\eta^2 - 1)^3 + \epsilon_{\texttt{head-on}}e^{10(\eta-1.2)}, \tag{3.1}$$

where η is the photon angle of incidence with respect to the photo-multiplier tube; see [81, 100, 101] for a detailed discussion of this parameterization. The $\epsilon_{\texttt{head-on}}$ parameter modifies the photon efficiency in the vertical direction, while the $\epsilon_{\texttt{lateral}}$ parameter modifies the lateral direction's efficiency. Of these three parameters, only $\epsilon_{\texttt{DOM}}$ and $\epsilon_{\texttt{head-on}}$ have a significant effect on the observable distributions in this analysis, and so $\epsilon_{\texttt{lateral}}$ is fixed to a nominal value obtained from calibration data in the simulation used for this analysis. Dedicated simulations are run for different values of relative $\epsilon_{\texttt{DOM}}$ to incorporate uncertainties that stem from these parameters into the analysis. From these simulations interpolating b-splines are constructed with `PHOTOSPLINE` [102, 103] to describe the ratio between expected event distributions. Then, for each simulation, we construct two-dimensional histograms in the observed quantities for each neutrino flux – *i.e.* conventional, prompt, and astrophysical – and for each expected morphology – *i.e.* cascade, track, and double cascade. It is apparent from these observable distributions that $\epsilon_{\texttt{DOM}}$ primarily changes the overall normalization of the event rates. We then smooth these histograms by constructing a spline using `PHOTOSPLINE` [102, 103], resulting in a set of interpolating b-splines, one for each $\epsilon_{\texttt{DOM}}$. Finally, we linearly interpolate between the splines in the $\epsilon_{\texttt{DOM}}$ dimension. This systematic correction is applied multiplicatively to the expectation of the sample. Fig. 3.3 shows the change in event rate, with respect to the analysis nominal simulation, in the deposited energy and zenith distributions when shifting the $\epsilon_{\texttt{DOM}}$ by one sigma (the prior width). This parameter's primary effect is to increase or decrease the rate in the sample, having minor deposited energy and zenith dependence. Thus, this systematic parameter has the largest impact on the absolute normalization of the reported fluxes.

A similar procedure is performed to include the effect of changing the head-on efficiency. Again, a dedicated simulation is run for several values of $\epsilon_{\texttt{head-on}}$ to compute the systematic correction, in this case: -3, -1, 0, and 1; where 0 corresponds to the nominal value [80]. This correction is applied multiplicatively to the expectation resulting in the distributions in Fig. 3.4, where this parameter is varied within one standard deviation. The $\epsilon_{\texttt{head-on}}$ parameter primarily modifies the relative rate of observed up-going and down-going events.

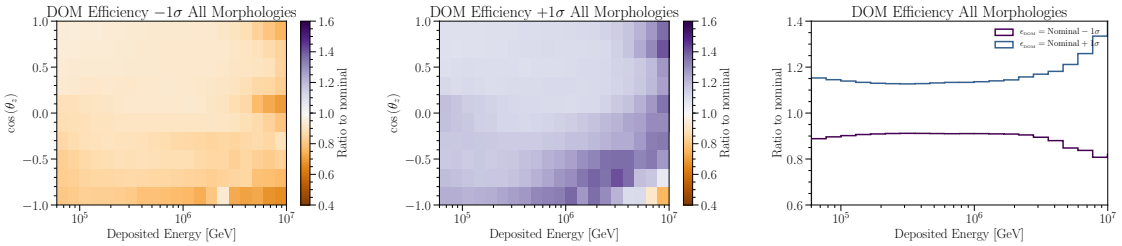

Figure 3.3: ***Effect of changing the DOM efficiency on the sample.*** These figures show the change relative to the nominal detector efficiency. Left and center panels show this ratio as a function of the deposited energy and the cosine of the zenith angle for decreasing and increasing the efficiency by one sigma, respectively. The color scale is proportional to the change, and less saturated colors correspond to lesser change. The rightmost panel shows the ratio of energy distributions, when decreasing and increasing the efficiency, as dark and light lines as a function of the deposited energy. In these three panels all morphologies are considered.

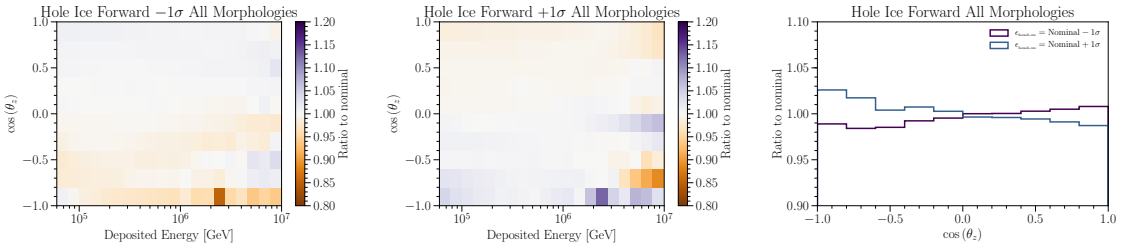

Figure 3.4: ***Effect of changing the hole ice on the sample.*** The layout and colors have the same meaning as in Fig. 3.3, and show the effect of decreasing and increasing the hole ice forward efficiency. Since this effect predominantly affects the angular distribution, the rightmost panel shows the ratio to nominal as a function of the cosine of the zenith angle. In these three panels all morphologies are considered.

The global properties of the ice are taken into account in different ways for different effects. The scattering and absorption of photons in the ice is azimuthally anisotropic because of the ice flow [80]. This azimuthal anisotropy is modeled as the effect is described in [104], although newer modeling of the anisotropy through the birefringent properties of ice [105] will supersede this in the future and is expected to produce some changes to the reconstructed event directions. For the double cascade morphology, changes to the scattering and absorption lengths can alter the event's apparent length because changes to scattering can alter the arrival time of photons with respect to the axis connecting the two cascades. Therefore, this effect can bias the double cascade length reconstruction if the orientation of the anisotropy axis and the strength of the anisotropy are not well modeled. Calibration measurements well constrain the anisotropy axis, however, the strength of this effect is more uncertain. Uncertainties of this effect are incorporated by parameterizing the bias in the length reconstruction with an analytic function. This parameterization is given by the piece-wise

function

$$\begin{aligned}\Delta l(\theta_z, \phi, a_{\mathtt{s}}) &= a_{\mathtt{s}} \cdot \Delta l(\theta_z, \phi, a_{\mathtt{s}} = 1) \\ \Delta l(\theta_z, \phi, a_{\mathtt{s}} = 1) &= \left(b_1 \cos^4\left(\theta_z - \pi/2\right) + b_2\right) - \left(a_1 \cos^4\left(\theta_z - \pi/2\right) + a_2\right) \\ &\quad \cdot \sin\left(2\left(\phi - \phi_0\right)\right) \cdot \begin{cases} 1 & \texttt{if} \;\; \sin\left(2\left(\phi - \phi_0\right)\right) < 0 \\ \frac{1}{2} & \texttt{if} \;\; \sin\left(2\left(\phi - \phi_0\right)\right) > 0 \end{cases},\end{aligned} \tag{3.2}$$

where ϕ_0 is the angle of the anisotropy axis, a_1 and a_2 are amplitudes, and b_1 and b_2 are offsets. The parameterization is chosen based on the sinusoidal behavior of the bias as a function of the azimuth. The conditional term defines two regions in line with the sinusoid, where the bias is positive in one and negative in the other. The bias strength in these two regions is off by a factor of 2, giving rise to the terms in the conditional. Both the amplitude and the average bias vary with the zenith angle approximately as $\cos^4$. The amplitude and offset terms define the bias at two extremes where the axis is entirely aligned or misaligned with the azimuthal plane. Parameters a_1, a_2, b_1, b_2 are determined by a fit to simulation after performing reconstructions with variable anisotropy strength assumed by the reconstruction algorithm. Different values of these parameters are determined for each class of interaction. Implicitly this treatment assumes that the bias effect is linear with the anisotropy scale and that swapping the anisotropy scale of simulation and reconstruction produces the inverse of the length bias.

Although we have obtained an analytic parameterization of the observable bias, this cannot be used directly in our analysis without violating the statistical model's assumptions, as events would move from bin to bin. The effect of this bias on the distribution of event observables is instead parameterized with splines in the same way as the previously mentioned ice and detector systematics, using discrete values of the anisotropy scale. The result of applying this change in anisotropy to this sample is shown in appendix Fig. 3.5. Other observables used in this analysis are not strongly affected by this systematic uncertainty, so the effect is neglected.

The bulk ice scattering and absorption uncertainty is sub-leading, and the impact is evaluated by repeating the analysis with three different ice variants. The three ice variants used are a 10 % increase in overall light scattering, a 10 % increase in overall light absorption, and a simultaneous 7.5 % reduction of both absorption and scattering. We found that the inclusion of the effects of bulk ice scattering, bulk ice absorption, and the discrete atmospheric flux choices previously mentioned in Section 6.2 increases the reported uncertainty of the astrophysical parameters by at most 20 % with respect to errors computed without these effects. For this reason, these effects are not included in the analysis or reflected in the reported errors of model parameters.

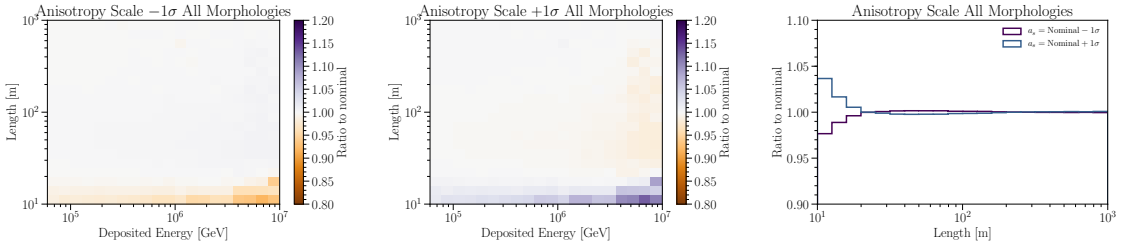

Figure 3.5: ***Effect of changing ice anisotropy on length distribution in the sample.*** The layout and colors have the same meaning as in Fig. 3.3, and show the change relative to the baseline to ice model anisotropy. The rightmost panel shows the ratio of length distributions when decreasing and increasing the anisotropy, as dark and light lines as a function of the reconstructed length. In these three panels all morphologies are considered.

Chapter 4

Searching for astrophysical neutrinos: High energy starting events

4.1 Event selection

The high-energy starting-event (HESE) sample was originally developed as a part of the work that discovered the astrophysical neutrino flux [59] and remains unchanged in this analysis. It aims to isolate astrophysical neutrinos by reducing the background of not only atmospheric muons but also atmospheric neutrinos. In order to do so, the outer parts of the detector are used as a veto layer, which aims to select only events with a contained interaction vertex, here referred to as "starting events".

The fiducial volume approximately excludes the top-most 90 m of the detector, 90 m from the outer layer of DOMs on the detector sides, the bottom-most 10 m of DOMs, and a 60 m thick horizontal layer directly below the region of ice with the largest dust content. The veto then consists of the DOMs excluded from the fiducial volume. Figure 4.1 shows a schematic of the veto. The veto's top and side regions are thick to increase the probability that atmospheric muons are caught entering the detector. The bottom region is thinner, as there is no up-going background from atmospheric muons and exists only to make the response of the selection to neutrinos uniform in angle. The horizontal region below the dust layer serves to veto muons entering the detector from the side within the dust layer.

We first define the approximate trigger time (t_0) and vertex position ($\vec{x}_0$). The event trigger time is defined as when the integrated charge deposition in the detector reaches 250 PE, excluding the DOMs in DeepCore, and considering only charges with hard-local-coincidence (HLC) triggers [106, 82, 81]. On each DOM, at least one "hit" is recorded if the module detects a voltage corresponding to at least 0.25 PE. HLC triggered hits are DOM hits that are in coincidence ($\pm 1\,\mu$s) with another hit on the four closest DOMs of

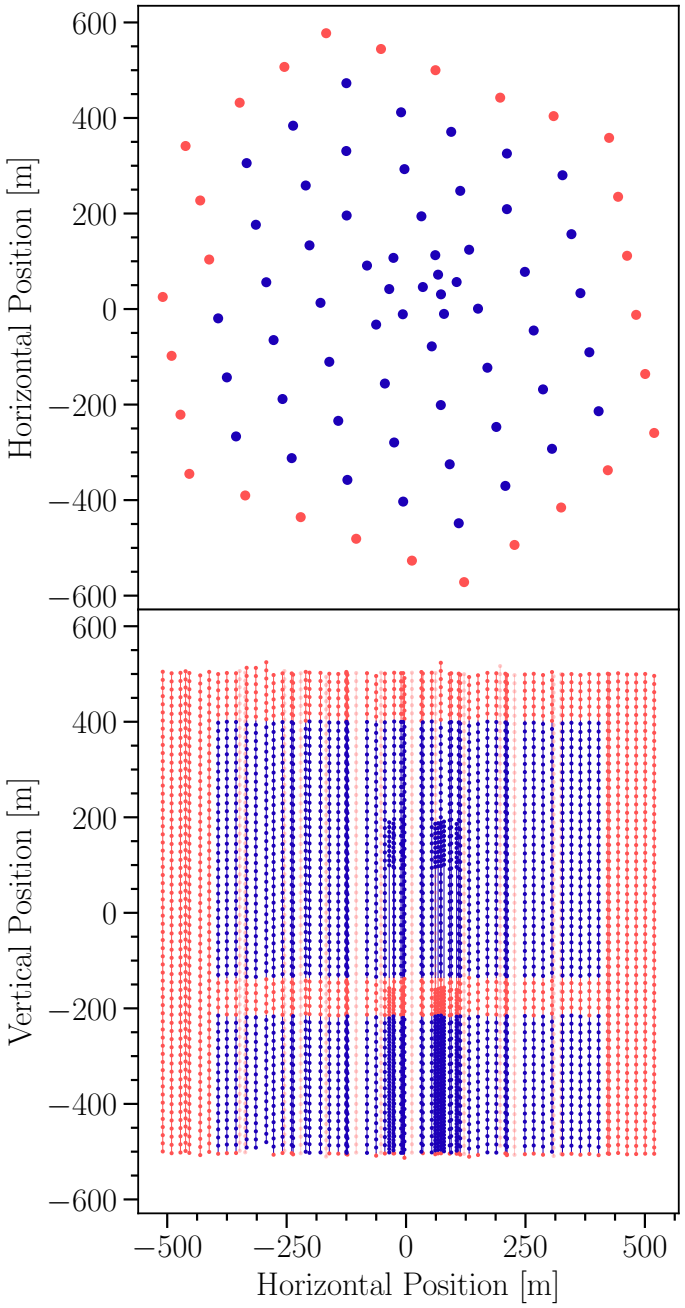

Figure 4.1: ***HESE veto.*** Diagram of the IceCube detector with the veto DOMs indicated. Top panel: an overhead view of the IceCube detector. Positions of strings with only veto DOMs are shown in red, while those with at least one non-veto DOM are shown in blue. Bottom panel: a side view of the IceCube detector strings and DOMs. Veto DOMs are indicated with red circles and non-veto DOMs with blue circles. Strings in front of or behind the region without veto-DOMs are semi-transparent.

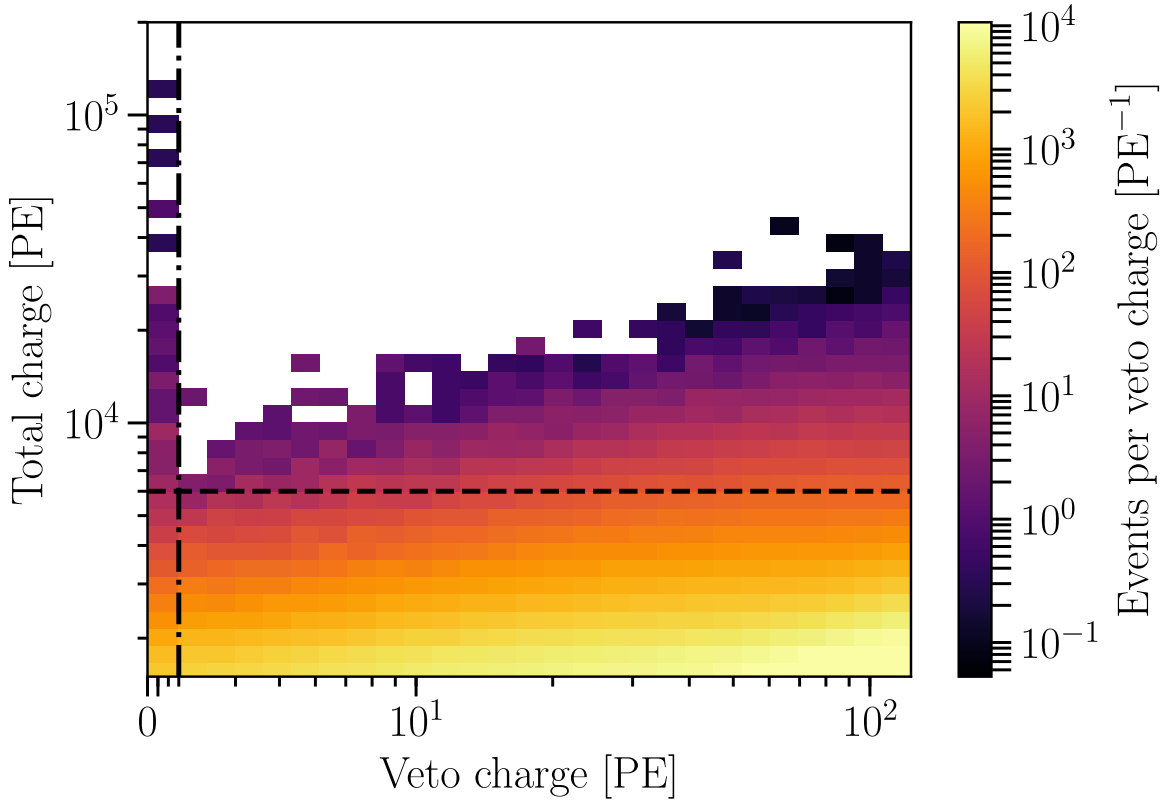

Figure 4.2: ***Veto charge vs. total charge.*** The color scale shows the event density with respect to the veto charge, prior to the veto cuts and charge cuts, where darker color implies higher event density and lighter implies lower event density. The vertical axis is the total charge, in photo-electrons, deposited in the detector, while the horizontal axis is the veto charge as defined in Section 4.1. The horizontal axis is plotted in linear-scale from $0-3$ PE and in log-scale from 3 PE upwards. The high-charge population of low-veto-charge events (between 0 and 3 PE) can be clearly seen against the background of higher-veto-charge events. The vertical dashed line indicates the cut on total charge of events at 6000 PE, and the horizontal dashed-dotted line indicates the cut on veto charge of events at 3 PE

the same string. Choosing only HLC triggered hits reduces the noise, and excluding DeepCore makes the detector response more uniform. To further prevent random noise adding to the threshold over long time periods, a three-microsecond time window is considered when computing the event start time. This time window encompasses the time a high-energy muon takes to traverse the detector. The event's interaction vertex is approximated by the charge-averaged position of the event's first 250 PE from HLC hits. To reduce the background of muons entering from outside the detector, we select only events with three PEs or less from veto hits and with veto hits on fewer than three DOMs. A veto hit with time t_1 and position $\vec{x}_1$ is defined as an HLC triggered hit that meets the conditions summarized in Table 4.1a. As the first condition, the hit must be on a DOM within the veto region. Hits are required to arrive before $t_0 + 50$ ns to select only

hits from light originating outside the fiducial volume. This cut is sufficient as the minimum distance from the fiducial volume to a veto DOM is 17 m. The time of veto hits is required to be within 3 µs of the trigger time to reduce the contribution of noise. Finally, to select only hits that may be related to the event vertex, veto hits are required to be on DOMs within $3\,\mu\text{s} \cdot c$ (~ 899 m) of the event vertex position. The 3 PE veto cut removes atmospheric events that are likely to deposit charge in the veto region but is not efficient at removing lower energy atmospheric events. The distribution of data events with respect to the veto charge and total charge is shown in Fig. 4.2. Atmospheric muon events comprise the bulk of the distribution above a veto charge of 3 PE, whereas, below this threshold, the distribution is dominated by a population of neutrinos at higher total charge. In the Northern sky, both atmospheric and astrophysical neutrinos can contribute significantly to this low-veto-charge population. However, in the Southern sky, atmospheric neutrinos are accompanied by muons from the same air-shower, whereas astrophysical neutrinos are not. Thus, in the Southern sky, we expect astrophysical neutrinos to be the dominant component in this population.

The veto allows astrophysical neutrino events in the Southern sky to be separated from the vast majority of atmospheric muons and from a portion of atmospheric neutrinos in a manner that provides some advantages with respect to non-veto methods [64, 107]. The physical separation of the veto and non-veto regions allows us to separately characterize the veto's response to incoming atmospheric muons independently of the non-veto portions of the detector. This characterization is a key ingredient in calculating the event selection response to atmospheric neutrino events with accompanying muons. Atmospheric events with accompanying muons are modeled as lone neutrinos that have an additional suppression effect. This calculation of this suppression is discussed at length in Section 6.1.2, and subverts the need for more expensive air-shower simulations. Improved background rejection can also be achieved through other methods [64], but without the same physical separation. The clear separation of the detector into fiducial volume and veto region also allows for a specific type of muon background estimation. A data-driven estimate of the atmospheric muon background is obtained using the outermost layer to identify muons, and performing the event selection in a reduced volume. Chapter 6 describes this method in more detail. The event selection effective area is also approximately isotropic for astrophysical neutrinos (without accounting for absorption in the Earth), but the background has a highly zenith-dependent acceptance. This property enables up/down asymmetry tests of the astrophysical flux in data that can be performed without detailed knowledge of the detector response. Although it is possible to design more sensitive event selections, these sanity checks are robust against certain classes of unknown systematic effects and have provided the community with confidence in the original discovery of astrophysical neutrinos with this sample.

To minimize the number of atmospheric events in the selection, we select only events with at least 6000 PE deposited in the detector. This charge cut removes background events with downward fluctuations

Veto-hit conditions
1) Hit on DOM within veto region
2) $t_1 \leq t_0 + 50\,\text{ns}$
3) $t_1 \geq t_0 - 3\,\mu\text{s}$
4) $\lvert\vec{x}_0 - \vec{x}_1\rvert \leq 3\,\mu\text{s} \cdot c$

(a) ***Veto-hit conditions.***

Parameter	Value
Event start time (t_0) charge threshold	250 PE
Maximum veto charge	3.0 PE
Maximum DOMs with veto hits	2
Minimum total charge	6000 PE
Trigger time window	3 µs

(b) ***HESE cuts and definitions.***

Table 4.1: ***HESE cuts and veto hit definitions.*** This table contains the criteria a hit must satisfy to be considered a veto hit (left) and the definitions of cut values and related quantities (right). Here, t_0 is the approximate trigger time (when 250 PE have accumulated) and $\vec{x}_0$ is the approximate vertex position (the charge average position of the first 250 PE). The DOM position and hit time are given by $\vec{x}_1$ and t_1, respectively. Only events with 6000 PE or more in total charge are considered as part of the sample. Veto hits must be on DOMs within the veto region. Only hits that are at most 3 µs before the trigger time and less than 50 ns after the trigger time are considered for the veto. Hits that are outside a ~ 899 m radius from the approximate vertex position are also excluded from the veto. If the veto hits constitute more than three photo-electrons or are distributed over more than two DOMs, the event is rejected from the sample.

in muon stochastic energy losses, keeping only events guaranteed to be high energy. A charge cut is preferred over a reconstructed energy cut as it is more closely related to the observed event light yield and thus is a more robust estimator of expected veto charge and is less reliant on simulation details. This cut was determined using burn-sample data, equivalent to 10 % of two years of detector operation, by requiring that no identified muons pass the charge cut. For the isotropic flux analyses in Section 8.1 a reconstructed deposited energy cut is placed at 60 TeV to reduce muon contamination further and limit the impact of normalization uncertainties and unknown shape uncertainties for this background component. As shown in Section 6, the muon component does not significantly contribute to the sample above 60 TeV. This 60 TeV cut is not imposed for the source search analyses in Section A.1 as they have a conservative analysis design that uses scrambled data to model the background test-statistic distribution.

Table 4.1b summarizes the event selection criteria; these are limited to cuts on the total charge, veto charge, and veto hit multiplicity. Combining these cuts and the definition of the approximate event interaction vertex fully specifies the event selection. Figure 4.3 shows the effective area of the sample for astrophysical neutrinos of three flavors, averaged between neutrinos and antineutrinos. A data taking period of approximately seven and a half years is used for this analysis. The chosen data taking period corresponds to a detector livetime of approximately 2635 days once offline periods are accounted for. Table 4.2 summarizes the number of observed events that pass the selection using reconstructed quantities described in 5. A total of 102 events were observed. Of these events, 60 have deposited energies above 60 TeV: 41 cascades, 17 tracks, and 2 double cascades. Even though the event selection has not changed with respect to previously reported results [59, 60, 61, 62] the event properties and the selected events themselves have changed due to a re-calibration of the single photo-electron (SPE) charge distributions of each digitizer in the detector. This re-calibration is

described in more detail in Section 4.2. This re-calibration's net effect is a decrease in the total observed charge of $\sim 4\,\%$ on average. This results in some events dropping below the total charge cut, which was not changed in a corresponding manner. Eight events were dropped from the sample because of this re-calibration. Seven events were removed because their total charge is now less than 6000 PE, with event numbers: 5, 6, 42, 53, 63, 69, and 73. Of these, only three – one cascade and two tracks – have deposited energy above 60 TeV where the astrophysical component measurements are performed. Removal of the events that failed the charge cut was not entirely necessary but is more an artifact of the unblinding procedure since we did not consider correspondingly shifting the charge cut. An additional track event, event 61, was also removed as it now fails the veto criterion: the time to accumulate 250 PE increased, changing the vertex position, and allowing more hits to meet the veto hit criteria. Finally, a ninth event, event 62, is not included in the sample due to a loss of low-level data required for the re-calibration.

Of the two double cascades above 60 TeV, one has a high probability of originating from a ν_τ interaction [108], while for the other event the ν_τ origin is simply favored with respect to a ν_e or ν_μ origin. Below 60 TeV flavor discrimination is generally poorer, but in particular, the identification of tau neutrinos is no longer robust. In this lower energy region, we reconstruct 41 cascades, 10 tracks, and 2 double cascades. Although it is tempting to apply the same degree of interest to all double cascade events, the two below 60 TeV are in a region of reconstructed parameter space with higher contamination from other neutrino flavors and larger background uncertainties. These events are likely neutrinos as indicated by their up-going direction, and have a similar likelihood of astrophysical or atmospheric origin as their flavor is unknown and energy is lower.

Category	$E < 60$ TeV	$E > 60$ TeV	Total
Total Events	42	60	102
Up	19	21	40
Down	23	39	62
Cascade	30	41	71
Track	10	17	27
Double Cascade	2	2	4

Table 4.2: ***Observed events by category.*** The left-most column indicates the event category, which may correspond to a particular choice of morphology or direction. The right-most column shows the total number of data events observed in a given category. Intermediate columns split events into those with less than or greater than 60 TeV reconstructed deposited energy.

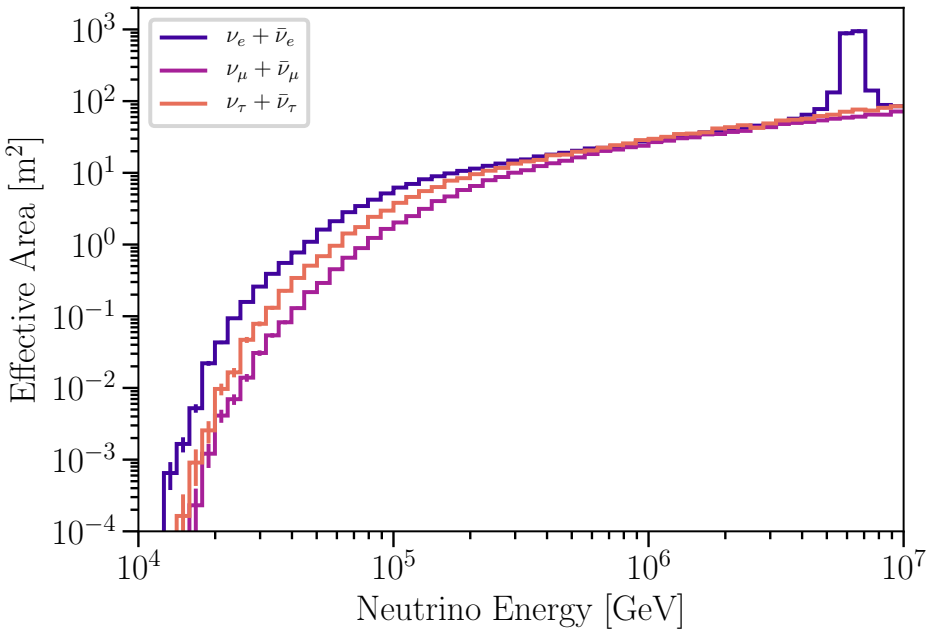

Figure 4.3: ***All-sky energy dependent astrophysical neutrino effective area.*** The direction averaged effective area for the three neutrino flavors is shown as a function of the neutrino energy incident to Earth. This effective area takes into account the effects of absorption in the Earth. This effective area can be multiplied by the all-sky astrophysical neutrino flux to obtain the expected number of astrophysical neutrinos. On the other hand, to obtain the expected number of atmospheric events, the atmospheric passing fraction must be included in the calculation of the effective area.

4.2 Single photo-electron charge distribution calibration

Each of IceCube's digital optical modules uses a discriminator circuit to trigger on the PMT voltage level and a selection of digitizers with different time windows, digitization rates, and dynamic range to record the PMT output after a trigger. To perform accurate physics measurements, we must convert the digitized waveforms into deposited charge time stamps. This conversion is achieved with *in situ* measurements of various calibration constants for each DOM. The voltage and timing calibration is described in detail by [81]. The PMTs in the DOMs operate at a gain of 10^7. In the PMT operation that results in this gain, the process of producing knock-off electrons is highly stochastic, meaning that a single photo-electron can give rise to a broad distribution of charges at the final stage with a FWHM of 0.75 times the mode and a positive tail extending up to 5 times the mode. The PMT gain as a function of high voltage is calibrated using the response to background "dark-noise" photons to achieve the target gain of 10^7. For single photo-electrons (SPEs), the charge before amplification e is known. By integrating the voltage samples (as a function of time, and dividing by the load resistance) of many single photo-electron waveforms, a histogram of response charges is formed for each high voltage setting. The histogram of response charges above a particular charge is fit

with the sum of an exponential and Gaussian distribution. The Gaussian component peak (SPE peak) is used to set the PMT high voltage such that the gain at this point in the distribution is 10^7. The discriminator threshold is then set to be 0.25 times the voltage corresponding to the Gaussian peak. This calibration procedure is repeated at the start of each IceCube season. The stability of the DOM operating conditions makes this calibration frequency sufficient.

After the digitizer readout and gain of the PMT have been calibrated, the interpretation of the digitized waveform readout as deposited-charge then relies on the calibration of the single photo-electron charge distribution (SPE curve). Previous IceCube analyses used the same sum of exponential and Gaussian distributions that were fit during the calibration procedure to interpret the readout waveforms. However, in 2015 further investigation revealed that a more accurate procedure than that used for initial calibration resulted in SPE peaks that were, on average, 4.3 % higher than expected. This realization prompted a detector wide re-calibration effort for all IceCube data seasons, internally referred to as "Pass2." The offset in the initial calibration procedure is in part due to contamination from multi-photo-electron pulses and failures of the fitting procedure to match the data well. The Pass2 re-calibration procedure uses raw waveform information collected via an unbiased random filter designed to capture detector noise and background muon events. Like the calibration procedure, the integrated charges from each digitizer are first deconvolved from electronics effects and then binned so that a function can be fit to the distribution. Depending on the available information and fit quality, an exponential plus Gaussian distribution or the sum of an exponential and two Gaussian distributions is fit to each histogram above the 5th percentile of the data. This procedure is performed on a per-season basis for each digitizer on every DOM. Once the new SPE curves have been obtained, the data is reinterpreted accordingly and processed with the filters used for the 2017 data taking season to provide a uniform set of filters for all the re-processed years of data.

The changes introduced by the new filters are small for most IceCube analyses. However, the re-calibration of the SPE curves has resulted in a change to inferred deposited charge, and therefore energy, of approximately -4 % on average. This systematic bias was partially accounted for in other IceCube analyses that consider the uncertainty in the absolute photon efficiency of the DOMs. This analysis and others forthcoming use the new Pass2 calibration to obtain more accurate estimations of the deposited charge and energy of events in the detector. Notably, in this analysis, the reconstructed deposited energy has changed compared to previous analyses of this sample due to the charge re-calibration. As the selection uses cuts on the charge of events which have not been altered, some events have been removed from the sample after re-calibration. Further calibration efforts of the DOM response are ongoing and have since yielded more accurate measurements that will be used in future analyses [109].

Chapter 5

Event reconstruction and simulation

Reconstruction of the neutrino events involves determining the interaction vertex, the incident direction, and the energy depositions – positions and magnitudes – in the detector. The interaction vertex determined as discussed in Section 4.1 is used only in the event selection, while the more sophisticated reconstructions described in this section introduce the interaction vertex as a free parameter. We separately consider hypotheses formed according to the three morphologies: track, cascade, and double cascade. For each of these hypotheses, we determine the expected light arrival time distribution on all DOMs given an event hypothesis and maximize the likelihood of these light distributions given the data with respect to the direction, vertex, and energy depositions. This process is described in sections 6 and 8.2 of [110].

To incorporate information about neutrino flavor, we assign each event a reconstructed morphology according to the classification algorithm described in [111, 108]. This algorithm can be broken into five distinct steps.

1. A series of quality cuts are made on the double cascade reconstruction output, namely, the double cascade reconstruction must not fail, and the opening angle between the track and double cascade reconstructions must be less than or equal to 30°.

2. If the quality cuts fail, the likelihood of the reconstructed cascade and track directions are compared to assign the morphology as cascade or track.

3. If the quality cuts pass, events with reconstructed length $L < 10\,\mathrm{m}$ are assigned the cascade morphology.

4. Remaining events with energy confinement $E_C < 0.99$ are assigned the track morphology. To define the energy confinement we first define the energy of each cascade within 40 m as $E_{1,C}$ and $E_{2,C}$. Then the confinement is given by $E_C = (E_{1,C} + E_{2,C})/E_{\mathtt{tot}}$, where $E_{\mathtt{tot}}$ is the total reconstructed-energy of the event.

5. Remaining events with energy asymmetry $-0.98 \leq E_A \leq 0.3$ are assigned the double cascade morphology, and other remaining events are assigned the cascade morphology.

Here, the energy asymmetry is the ratio of the difference between the two double cascade reconstructed energy depositions and the sum of the two energy depositions, $E_A = (E_1 - E_2)/(E_1 + E_2)$, where E_1 and E_2 are the reconstructed energy of each cascade.

This method effectively separates "cascade-like" and "track-like" morphologies at high energies, but more importantly, it produces a majority sample of tau neutrinos in the double cascade category above 60 TeV. Despite this categorization method, there remain contributions from all flavors in all three morphological categories, which are outlined in Table 5.1. The non-negligible rate of misidentification is not an issue for inference, though, and the asymmetry in the contributions can be used to constrain the flavor composition of the neutrino events. A more detailed analysis of the astrophysical neutrino flavor content is presented in [108].

Morphology	Cascade	Track	Double Cascade
Total	72.7 %	23.4 %	3.9 %

Morphology	Cascade	Track	Double Cascade
ν_e	56.7 %	9.8 %	21.1 %
ν_μ	15.7 %	72.8 %	14.2 %
ν_τ	27.6 %	10.5 %	64.7 %
ν_e CC	51.9 %	8.8 %	18.2 %
ν_μ CC	8.7 %	71.6 %	10.9 %
ν_τ CC	23.6 %	9.8 %	62.9 %
ν CC	84.3 %	90.2 %	92.0 %
ν NC	14.8 %	2.6 %	6.9 %
ν GR	0.9 %	0.3 %	1.2 %
μ	0.0 %	6.9 %	0.0 %

Table 5.1: ***Expected events by category for best-fit parameters above* 60 TeV.** Each column specifies the morphology of reconstructed events. Each row specifies a particle type, interaction type, or combination thereof. The top table provides the percentage of events expected in each morphology with respect to the total number of events. The bottom table provides the percentage of events in each category for a particular morphology, where percentages are computed with respect to the total number of expected events of the specified morphology. Here, CC stands for deep inelastic charged-current scattering, NC for its neutral-current counterpart, and GR for Glashow resonance. The percentages have been rounded to one decimal point.

The magnitude of energy depositions can be reconstructed to $\sim 10\,\%$ accuracy if they are contained within the detector [110]. In this sample, the median deposited energy resolution is $\sim 7.9\,\%$, $\sim 11\,\%$, and $\sim 7.8\,\%$ for cascades, tracks, and double cascades, respectively. Fig. 5.1 shows the median resolution of the reconstructed electromagnetic-equivalent deposited energy as a function of the simulated true electromagnetic-

equivalent deposited energy within the detector for the three reconstructed morphologies. Some reconstruction uncertainty stems from hadronic cascades, as they have more variability in Cherenkov light yield than their electromagnetic counterpart [110]. The deposited energy is correlated with the neutrino energy, which can be used to constrain the neutrino energy spectrum. However, in neutral current interactions, the outgoing neutrino can take away a large fraction of the initial energy making the deposited energy much smaller than the neutrino energy. The left panel of Fig. 5.2 shows this correlation. The deposited energy is peaked close to the neutrino energy and the long tails of the reconstructed distribution. To visualize this more clearly, Fig. 5.3 shows the distribution of reconstructed deposited energy for slices in true neutrino energy, where the selection truncates the tail of the distribution for lower neutrino energies. If we use the deposited energy as a proxy for the neutrino energy, then we obtain a median neutrino energy resolution of $\sim 11\,\%$, $\sim 30\,\%$, and $\sim 18\,\%$ for reconstructed cascades, tracks, and double cascades respectively. We can compare these to the resolutions of the deposited energy to see the impact of other effects. Cascades have additional uncertainty that stems from the neutrino interactions' kinematics and from differences in light yield for electromagnetic and hadronic showers as IceCube is unable to differentiate between these types of showers. Tracks also suffer from the same kinematics issues but lack complete information in ν_μ CC events where the resulting muon exits the detector. Finally, double cascades have more uncertainty than cascades because of the additional degrees of freedom in the reconstruction hypothesis associated with the production and decay of a tau.

The angular reconstruction is more straightforward by comparison, as the average angle between the primary neutrino and secondary particles of the interaction is smaller than a quarter degree above 10 TeV. This deviation is negligible compared to reconstruction uncertainties and is even smaller at the energy scale we are concerned with. Cascades, tracks, and, double cascades in this sample have a median zenith resolution of $\sim 6.3^\circ$, $\sim 1.5^\circ$, and $\sim 5.0^\circ$ respectively. The analysis of the astrophysical flux in Section 8.1 does not use the azimuthal directional information. Azimuthal resolution of the different event categories is larger than the zenith resolution by $0.3 - 0.6^\circ$ because the inter-DOM spacing is smaller along the vertical axis than the horizontal. The track angular resolution in this sample is worse than the resolution in ν_μ dominated samples [107] due to the non-negligible contamination from ν_e and ν_τ given in Table 5.1, the shorter length of tracks that start within the detector, and the presence of an initial hadronic cascade for ν_μ CC events. The large separation between the DOMs limits the angular resolution of cascades, a limited number of unscattered photons that can be detected, and our modeling of photon propagation in ice [80, 81]. We can explain the better angular resolution of tracks and double cascades by considering the longer path length of the highly-boosted secondary lepton. The taus and muons travel farther in the detector, allowing better reconstruction of their direction. To summarize the angular reconstruction's behavior, the right panel of Fig. 5.2 shows the distribution of reconstructed zenith angles as a function of the true neutrino zenith angle.

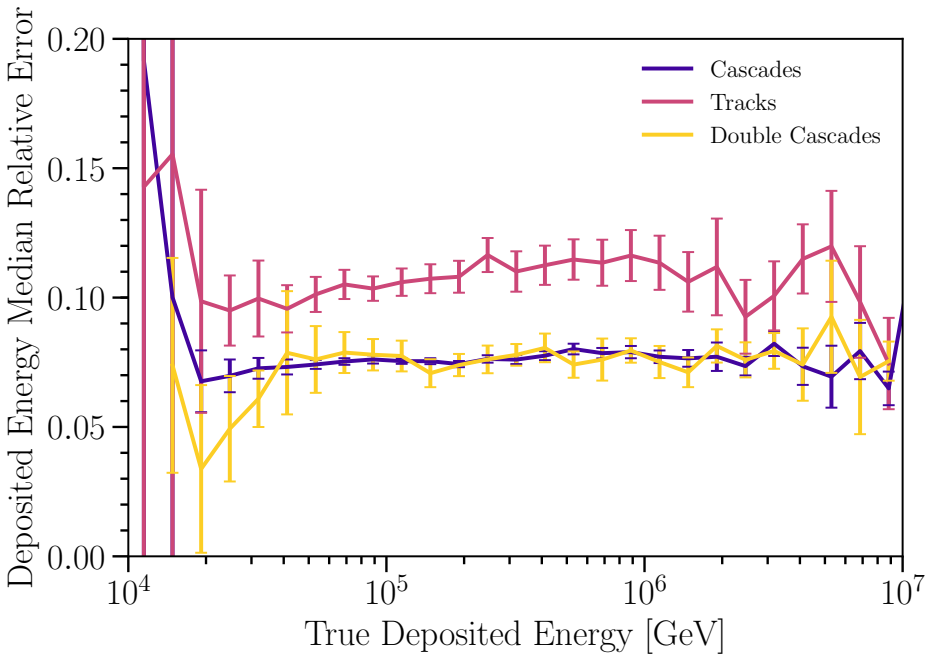

Figure 5.1: ***Deposited energy resolution.*** Each line shows the median energy resolution for a reconstructed morphology plotted as a function of the true deposited energy in the detector, with error bars indicating the median's statistical uncertainty. At these energies, cascade deposited energy uncertainty is affected by the spatial extension of the showers not modeled by the reconstruction, whereas the stochasticity of the losses dominates the uncertainty in track deposited energy.

The large smearing in this matrix arises from the cascades that dominate the data sample.

The reconstruction method used in this analysis has been changed compared to the previous iterations [59, 60, 61, 62] to enable better treatment of reconstruction uncertainties, which improves the accuracy of the analysis. Previously, progressively narrower brute-force scans of the neutrino direction were used for data. We now use a minimizer to determine the best-fit neutrino direction of data events, which offers a significant computational speed improvement. Additionally, the morphology determination is now performed algorithmically, whereas previous analyses performed morphology identification by hand. Although these changes may worsen the reconstruction's accuracy for individual events, they also make it feasible to run the reconstruction and classification on simulation events. By using the same algorithmic procedure for simulated events, we now account for reconstruction and classification uncertainties on an event-by-event basis, as opposed to using average uncertainties. Finally, the third morphological category (double cascades) is now included in the classification scheme instead of only tracks and cascades, which adds additional flavor information to the fit.

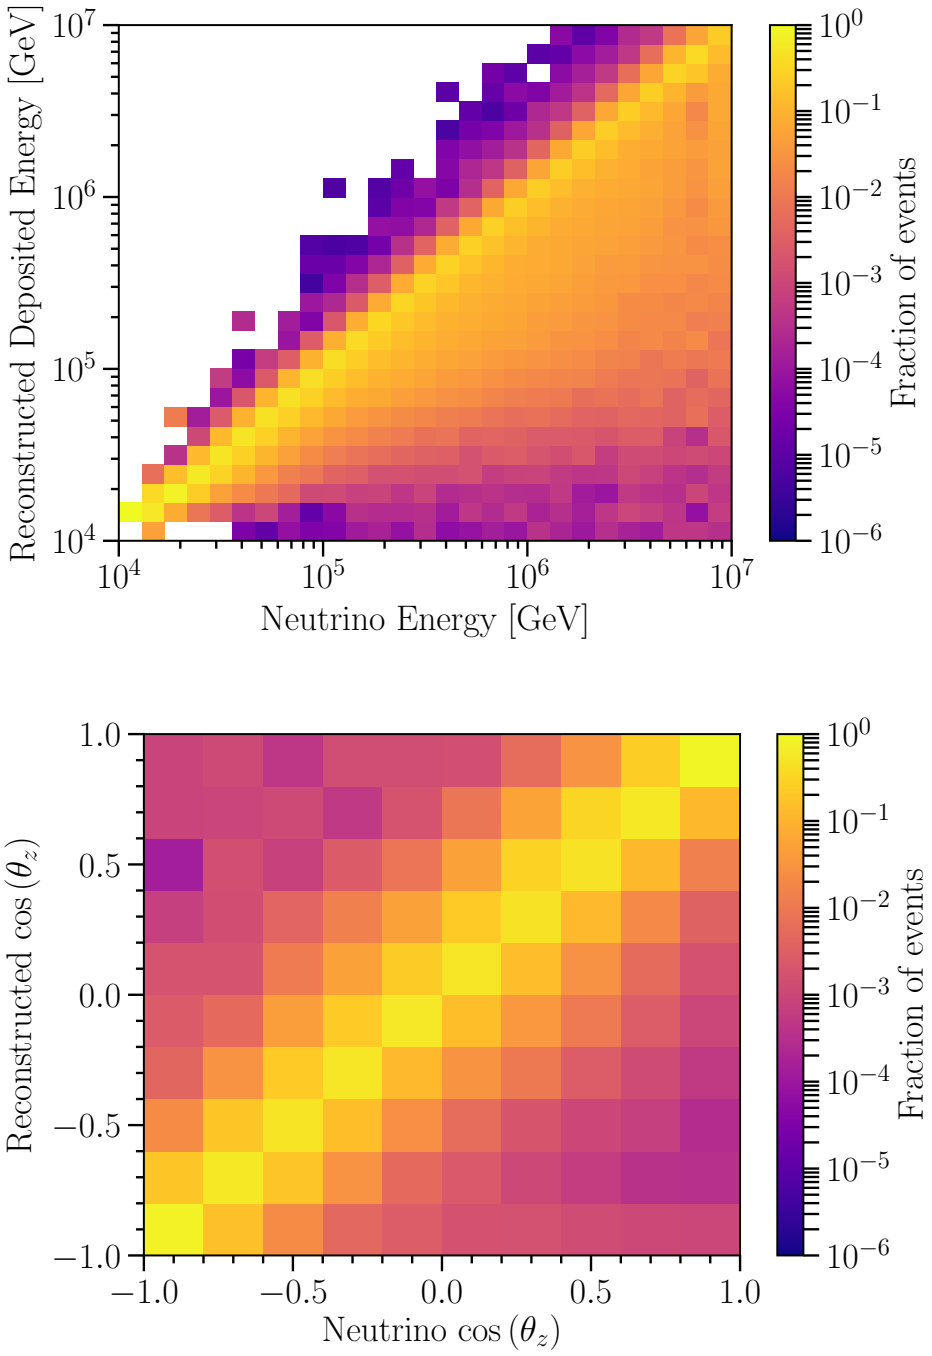

Figure 5.2: ***Distribution of expected reconstruction quantities as a function of true parameters.*** Transfer matrices, evaluated using simulation weighted to the single power-law best-fit parameters, are shown for all morphologies combined. The probability of a reconstructed deposited energy for a given neutrino energy (left) and the probability of a reconstructed cosine of the zenith angle for a particular cosine of the neutrino zenith angle (right) are shown. The matrices are column normalized. The asymmetry of the energy transfer matrix (left) is due to energy conservation, preventing large over fluctuations in reconstructed energy, and the wide range of visible energies possible for NC events which can lose large fractions of energy to the outgoing neutrino.

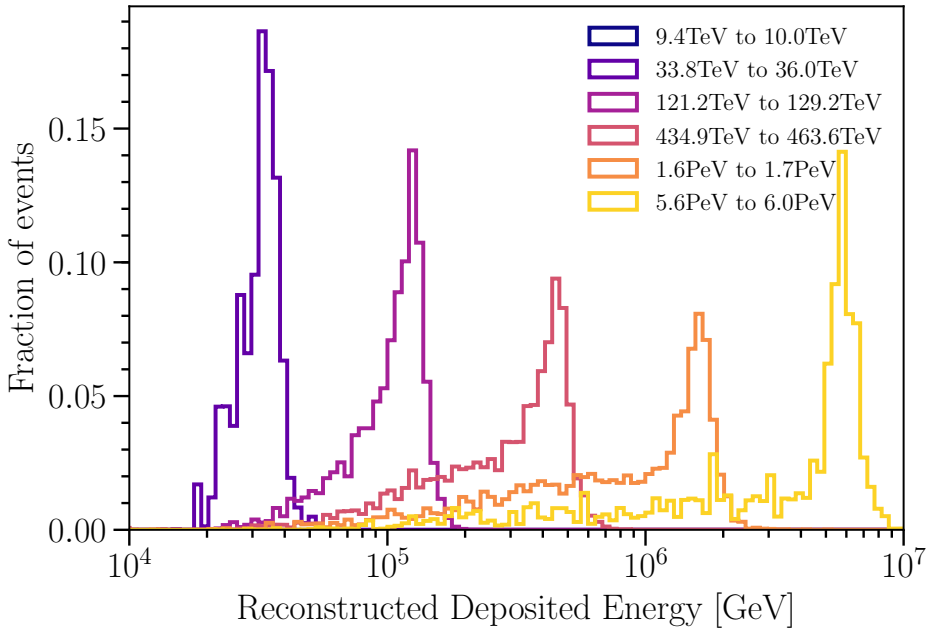

Figure 5.3: ***Deposited energy probability distributions.*** Probability distribution of reconstructed deposited energies for slices of true neutrino energy for all morphologies weighted to the best-fit parameters. The lowest energy bin in the legend does not contribute to the sample.

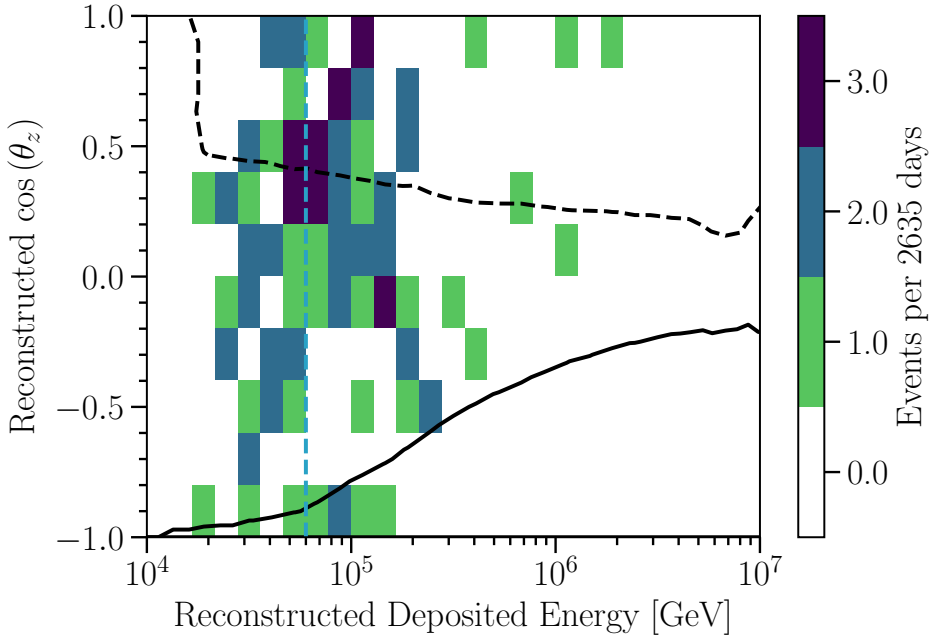

Figure 5.4: ***HESE events observed in* 7.5 years.** Histogram of the observed events as a function of their inferred deposited energy and cosine of the reconstructed zenith angle. The vertical dashed line indicates the low-energy threshold of 60 TeV. Up-going events are in the lower portion of the plot towards $\cos(\theta_z) = -1.0$, while down-going events are in the upper portion towards $\cos(\theta_z) = 1.0$. The deficit of up-going high-energy events is due to the absorption of neutrinos in the Earth. The solid line indicates the contour of 50 % probability for neutrinos to be absorbed in the Earth, below which the probability of absorption is higher. The black dashed line indicates the contour of 50 % probability for atmospheric neutrinos to be rejected by an accompanying muon, above which this rejection probability is higher. However, this contour's left vertical edge is an artifact of the simulation for this selection, as no simulated events are present to the left of this line.

Chapter 6

Estimation of backgrounds

In the search for astrophysical neutrinos, the predominant backgrounds are from atmospheric neutrinos and atmospheric muons. Single atmospheric neutrinos can produce the same event signatures as astrophysical neutrinos because they are not fundamentally different from astrophysical neutrinos. Atmospheric neutrinos can only be distinguished from their astrophysical counterparts by examining the population properties or the fact that muons produced in the same air shower can also simultaneously reach the detector. Muons from cosmic ray air showers differ from neutrinos in their event signature as they always produce light while entering the detector, while neutrinos can interact within the detector volume before any light is produced near the edge of the detector. Muons can also only be observed up to a certain amount of overburden, beyond which they will lose their energy and decay before reaching the detector. The up-going observation region is then free from atmospheric muons. The task in this section is to model these backgrounds and estimate the contribution of them in the final event selection. For atmospheric neutrinos, the initial flux is modeled with existing calculations. However, the estimation of this background in the sample must also account for the fraction of atmospheric neutrinos rejected because of muons present in the air shower. Section 6.1.1 details how the atmospheric neutrino flux is modeled, while Section 6.1.2 explores a new calculation of the atmospheric neutrino passing fractions and explores the many factors that may affect this calculation, and Section 6.1.2.5 applies this calculation to the HESE selection. The passing fraction calculation described in this chapter was developed in collaboration with Carlos A. Argüelles, Sergio Palomares-Ruiz, Logan Wille, and Tianlu Yuan [65]. Finally, Section 6.2 explores the contribution of lone muons to the sample.

6.1 Neutrino background estimation

6.1.1 Neutrino fluxes

Atmospheric neutrinos are predominantly produced by the decay of pions and kaons, which we shall call the "conventional" component. At energies above 1 TeV, the conventional component spectrum is softer than the incident cosmic-ray spectrum by one unit in the spectral index due to the interactions of these mesons in the atmosphere. The conventional neutrino flux is largest at the horizon, $\cos\theta_z = 0$ [112, 113, 114, 115] because the larger path length in the thin atmosphere increases the proportion of pions that can decay before interacting. To model the conventional neutrino flux, we use a parameterization of the Honda *et al.* 2006 flux calculation [114] given in [116] which at the highest energies uses the analytic parameterization of the neutrino flux in [112]. This calculation is tuned to match observations of atmospheric muons, which remain difficult to predict from first principles.

A sub-leading – yet unobserved – contribution due to charmed hadron decays is expected to be important above ~ 100 TeV [117]. Since the charmed hadrons decay promptly and do not interact in the atmosphere at the energies relevant for this analysis, we call this the "prompt" component. Thus, the prompt component has a spectral index close to the incident cosmic-ray spectrum at these energies. The prompt flux is also constant with respect to the cosine of the zenith angle, as the short decay time of the charmed hadrons removes the effect of different path lengths through the atmosphere. To model the prompt neutrino flux at Earth's surface, we use the flux computed in [117].

The angular and energy distribution of the initial atmospheric neutrino flux is modified with respect to the flux at Earth's surface because of the absorption of high-energy neutrinos in the Earth. We account for this using a dedicated Monte Carlo, similar to the one described in [118], which simulates the propagation of neutrinos in the Earth. In this Monte Carlo, we use the isoscalar neutrino cross sections given in [77] for the neutrino-nucleon interactions and the Earth density model described in [119]. Neutrino-electron scattering can be safely neglected except for resonant-W production [75], which we include. We ignore the uncertainties on the Earth opacity as they are known to be sub-leading in this energy range [78, 77, 120].

Modeling the atmospheric neutrino flux with these two components does not account for the contribution of K_s [121], which is $\sim 10\,\%$ at 100 TeV and well-within our uncertainties. In order to account for uncertainties in the cosmic-ray flux [122] and hadronic interactions [123] we parameterize the atmospheric neutrino flux as

$$\phi_\nu^{\mathrm{atm}} = \Phi_{\mathrm{conv}}\left(\phi_\nu^\pi + R_{K/\pi}\phi_\nu^K\right)\left(\frac{E_\nu}{E_0^c}\right)^{-\Delta\gamma_{CR}} + \Phi_{\mathrm{prompt}}\phi_\nu^p\left(\frac{E_\nu}{E_0^p}\right)^{-\Delta\gamma_{CR}}, \tag{6.1}$$

where ϕ_ν^π, ϕ_ν^K, and ϕ_ν^p are the conventional pion, kaon, and prompt atmospheric neutrino fluxes at a neutrino energy E_ν respectively.

The parameters $\Phi_{\texttt{conv}}$ and $\Phi_{\texttt{prompt}}$ are normalizations for the conventional and prompt normalizations respectively; $R_{K/\pi}$ allows us to modify the relative kaon to pion contributions; the $\Delta\gamma_{\texttt{CR}}$ parameter allows for hardening or softening of the atmospheric neutrino component to account for uncertainties in the cosmic-ray flux slope. These parameters are incorporated into the analysis as nuisance parameters with priors as summarized in Table 7.1. This analysis refrains from directly using prior information from other IceCube neutrino studies to provide results independent of other neutrino samples.

The conventional normalization prior is motivated by studies of the total uncertainty due to cosmic-ray and high-energy hadronic processes [123]; it is chosen to be Gaussian for simplicity with an expectation centered on the baseline flux model and appropriate variance. The standard deviation of the Gaussian prior on the cosmic-ray slope parameter has been chosen to be 0.05 in order to accommodate values measured at intermediate [124] and high [125, 126, 127] energies. The uncertainty on corrections to the ratios of neutrinos-to-anti-neutrinos ($2\nu/\left(\nu+\bar{\nu}\right)_{\texttt{atmo}}$) and kaon-pion yields ($R_{K/\pi}$) were estimated by comparing the expectation of different atmospheric neutrino calculations, and a Gaussian prior width was chosen that encompasses their predictions. Details regarding these corrections and the different atmospheric neutrino flux calculations can be found in [128, 129]. Finally, the parameters $E_0^c = 2020\,\text{GeV}$ and $E_0^p = 7887\,\text{GeV}$ are points of fixed differential flux for the conventional and prompt components.

This simple parameterization of the atmospheric neutrino flux and its uncertainties is chosen because it models the primary systematic effects of physical uncertainties that are observable with this sample. Of course, this neglects other physical uncertainties, such as those regarding the hadronic interactions in cosmic ray air showers and the composition of the cosmic ray particles. These effects do not produce modifications to the observations that can be statistically discerned with the amount of data available. Later it will be shown that even the modeled systematic uncertainties have a small effect on astrophysical measurements compared to the statistical uncertainties.

The above description models the flux of atmospheric neutrinos, but does not fully encapsulate the detector response to atmospheric neutrinos. We have so far neglected the effect of accompanying muons on the detector response. This is discussed at length in Section 6.2.

6.1.2 Atmospheric neutrino passing fractions

This discussion and study of atmospheric passing fractions are directly adapted from collaborative work with Carlos A. Argüelles, Sergio Palomares-Ruiz, Logan Wille, and Tianlu Yuan [65]. Figures and equations are reproduced from [65], and the text follows the same thread of discussion.

As noted in [130], muons produced in the same air-shower may trigger the detector veto in coincidence

with the neutrino interaction. Dedicated simulations of cosmic ray air showers where neutrino interactions are forced can provide observational predictions for the distribution of these combined muon+neutrino events. However, available simulation techniques to create these estimates remain prohibitively expensive for the desired accuracy. Instead, we look to model this effect by computing an average efficiency with respect to the case where no muons reach the detector. To account for this when weighting the neutrino-only simulation, each atmospheric neutrino flux component, i, is multiplied by an efficiency. This efficiency is referred to as the "atmospheric neutrino passing fraction", and denoted by $\mathcal{P}^{i,\alpha}_{passing}$, for each neutrino flavor α.

The passing fraction depends on air shower development details, which includes modeling of the hadronic interaction and the energy losses of muons in the material between the shower and the detector. Air-shower properties are averaged over in this calculation, but the passing fraction is still a probability that is conditional on the neutrino properties. Of particular interest are the passing fraction's potential systematic dependencies, which include the cosmic ray spectrum and composition, the hadronic interaction model, muon energy losses, and the atmospheric density model. The passing fraction calculation is designed so that these systematic dependencies may be changed between different alternatives. Additionally, the detector response enters only in one place in the formalism as a probability of muon rejection. This fact allows for simple variation of the detector response modeling in a generic way. A framework was developed to perform this calculation that leverages the Matrix Cascade Equation (MCEq) package in key portions of the calculation. Validation of the computed passing fractions is performed against detailed air shower simulations with COsmic Ray SImulations for KAscade (CORSIKA), and excellent agreement is found.

Departure from the lone-neutrino case is based on the coincidence of a muon and neutrino from the same air-shower that occurs in both time and direction. Thus,

$$\mathcal{P}_{\text{reach}}\left(E^{\text{f}}_{\mu} \,|\, E^{\text{i}}_{\mu},\, X_{\mu},\right) \tag{6.2}$$

is naturally a key component of the calculation, which describes the probability of a muon with initial energy E^{i}_{μ} and slant depth X_{μ} to reach the detector with energy E^{f}_{μ}. As the IceCube detector is deep underground, the contribution to slant depth from the atmosphere is very small compared to that of the Earth, and so can be neglected. With only the slant depth of the Earth contributing, X_{μ} is fully specified by the zenith angle θ_z in detector coordinates and the depth d. Stochasticity of the muon energy losses is accounted for by the distribution $\mathcal{P}_{\text{reach}}$, which is computed by tabulating simulations of muons propagated in ice by the software `MMC`. Using the conditional probability distribution $\mathcal{P}_{\text{reach}}$ instead of average muon behavior ensures that contributions from the distribution tails are accounted for. The contributions from the tails become important for very large slant depths, where muons on average do not reach the detector. Departure from the neutrino-only case requires both the muon to reach the detector and the muon to trigger a response, in

this case for the muon to *light* the veto. The quantity

$$\mathcal{P}_{\text{light}}(E_\mu^{\text{f}}, d) \tag{6.3}$$

is the probability that a muon of energy E_μ^{f} and depth d at the detector boundary *lights* the veto, such that the event is rejected. In practice, this $\mathcal{P}_{\text{light}}$ is computed by tabulating detector simulations of atmospheric muons and incorporating the specifics of the event selection. With both $\mathcal{P}_{\text{light}}$ and $\mathcal{P}_{\text{reach}}$ in hand, the probability of detecting an incident muon $\mathcal{P}_{\text{det}}$ can be computed as

$$\mathcal{P}_{\text{det}}\left(E_\mu^{\text{i}},\, X_\mu(\theta_z, d), d\right) \equiv \int dE_\mu^{\text{f}}\, \mathcal{P}_{\text{reach}}\left(E_\mu^{\text{f}} \,|\, E_\mu^{\text{i}},\, X_\mu(\theta_z, d)\right)\, \mathcal{P}_{\text{light}}(E_\mu^{\text{f}}, d)\ . \tag{6.4}$$

Note that losses also depend on the medium and not only on the slant depth, whose dependence on θ_z we write explicitly. If the medium surrounding the detector is homogeneous, the dependence on X_μ can be exchanged for distance without loss of generality. To simplify the notation from here on, the dependence on X_μ will be replaced by a dependence on θ_z and the dependence on d will be neglected. These dependencies can be added back in at the end of the calculation without any change to the result.

For an atmospheric neutrino of known energy and direction, the passing fraction is defined as the probability that muons from the same air shower do not trigger the detector veto. This probability is denoted $\mathcal{P}_{\text{pass}}$ [130, 131] and can be written as the ratio

$$\mathcal{P}_{\text{pass}}(E_\nu, \theta_z) = \frac{\phi_\nu^{\text{pass}}(E_\nu, \theta_z)}{\phi_\nu(E_\nu, \theta_z)}\ , \tag{6.5}$$

where E_ν is the neutrino energy, θ_z is the zenith angle, ϕ_ν^{pass} is the differential flux of atmospheric neutrinos accompanied by muons that are detected, and ϕ_ν is the total differential atmospheric neutrino flux. In the next sections, passing fractions for electron neutrinos and muon neutrinos are derived. As tau neutrinos by comparison, make up a tiny portion of the atmospheric flux, there is no dedicated discussion for this case. However, the concerns for tau neutrinos are very similar to those of electron neutrinos minus the differences in their production.

6.1.2.1 ν_e passing fraction

Electron neutrinos (or antineutrinos) are produced alongside a positron (or electron) in the decay of their parent particle, for example $K^+ \to \pi^0 + e^+ + \nu_e$. Rarely does an electron neutrino have a sibling muon, as this would involve a lepton flavor violating process. Instead, muons that accompany electron neutrinos are produced in other branches of the air shower. Because the different shower branches are uncorrelated, to first order the average properties of muons in a prototypical air shower can be used for the passing fraction calculation. This section aims to track the production and propagation of these muons from uncorrelated

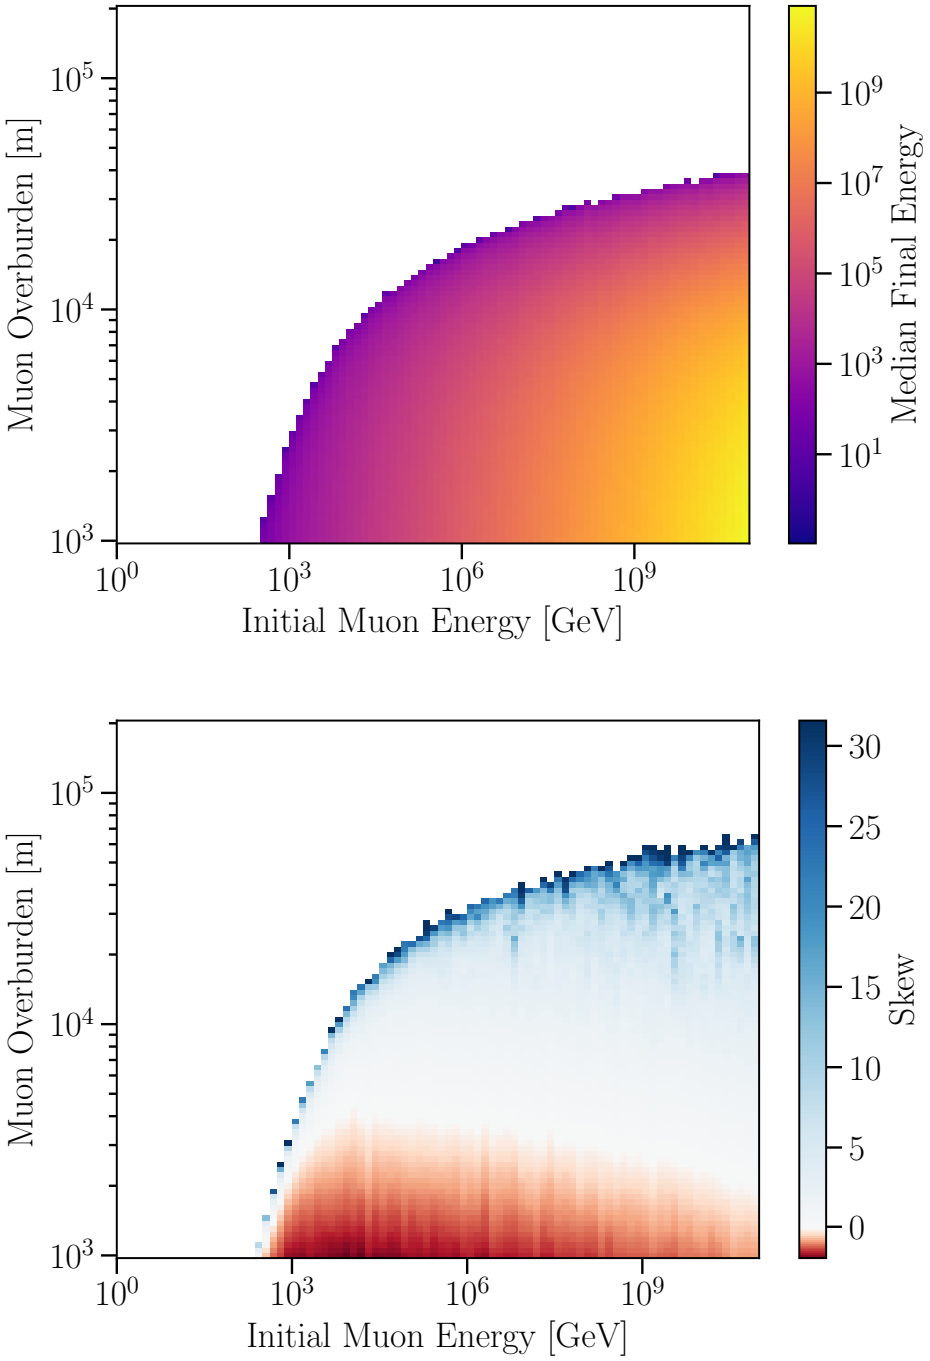

Figure 6.1: ***Muon Final Energy Statistics.*** Muons propagating in ice were simulated for a wide range of initial energies, and at a set of predefined distances their energy was recorded. This results in a distribution of final muon energies for each combination of initial muon energy and overburden. These two plots show statistics of these distributions, with the upper plot displaying the median and the lower plot displaying the skew. Previous calculations only used the median of these distributions, creating a one-to-one mapping from initial to final muon energy for each overburden or angle, as shown in the upper panel. However, these distributions' large skew show us that the tails become important for large overburdens as the lower panel shows.

branches of the shower and compute the flux of neutrinos not vetoed by these uncorrelated muons. Consider a shower from a single cosmic-ray primary particle of type A with energy $E_{\rm CR}$. The average neutrino yield $\frac{dN_{A,\nu}}{dE_\nu}(E_{\rm CR}, E_\nu, \theta_z)$ from such a shower can be computed with `MCEq`. So the atmospheric neutrino flux can be written as

$$\phi_\nu(E_\nu, \theta_z) = \sum_A \int dE_{\rm CR}\, \frac{dN_{A,\nu}}{dE_\nu}(E_{\rm CR}, E_\nu, \theta_z)\, \phi_A(E_{\rm CR}) \ , \tag{6.6}$$

where $\phi_A(E_{\rm CR})$ is the flux of primary cosmic rays of type A. the passing fraction can be obtained by modifying the integrand of Eq. 6.6 so that it is weighted by the Poisson probability *Pzmproto* of detecting an accompanying muon. In this way the uncorrelated passing fraction can be written as

$$\mathcal{P}^{\rm uncor,GJKvS}_{\rm pass}(E_\nu, \theta_z) = \frac{1}{\phi_\nu(E_\nu, \theta_z)} \sum_A \int dE_{\rm CR}\, \frac{dN_{A,\nu}}{dE_\nu}(E_{\rm CR}, E_\nu, \theta_z)\, \phi_A(E_{\rm CR})\, \mathcal{P}^{\rm shower}_{0\mu}\left(N_\mu = 0; \bar{N}^{\rm GJKvS}_{A,\mu}(E_{\rm CR}, \theta_z)\right) \ , \tag{6.7}$$

where $\mathcal{P}^{\rm shower}_{0\mu}\left(N_\mu = 0; \bar{N}^{\rm GJKvS}_{A,\mu}(E_{\rm CR}, \theta_z)\right) = \exp -\bar{N}^{\rm GJKvS}_{A,\mu}(E_{\rm CR}, \theta_z)$, and $\bar{N}^{\rm GJKvS}_{A,\mu}(E_{\rm CR}, \theta_z)$ is the average number of muons that are detected from an air shower. This average number of muons is computed using the yield of muons from the air shower $dN_{A,\mu}/dE^{\rm i}_\mu(E_{\rm CR}, E^{\rm i}_\mu)$ and the probability of detecting those muons, such that

$$\bar{N}^{\rm GJKvS}_{A,\mu}(E_{\rm CR}, \theta_z) = \int dE^{\rm i}_\mu\, \frac{dN_{A,\mu}}{dE^{\rm i}_\mu}(E_{\rm CR}, E^{\rm i}_\mu, \theta_z)\, \mathcal{P}_{\rm det}\left(E^{\rm i}_\mu, \theta_z\right) \ , \tag{6.8}$$

Some terms in the above equations are labeled with GJKvS because these quantities are calculated in the same manner as was done in [131] under certain choices for $\mathcal{P}_{\rm det}$. Namely, when $\mathcal{P}_{\rm light}$ is defined as a Heaviside function with a boundary at $E^{\rm f}_\mu = 1\,{\rm TeV}$, and $\mathcal{P}_{\rm reach}$ is defined as a delta function that matches the initial muon energy to the median final muon energy.

Fig. 6.2 shows the effect of the two muon treatments on the passing fractions as defined in Eq. 6.10. The median approximation overestimates the passing fraction at the horizon, while near the vertical direction, it is underestimated. In the horizontal region, muons with smaller initial energy such that their median range is less than the distance to the detector represent a small but non-negligible fraction of muons that may trigger the detector veto. Thus, in the median muon treatment, the passing fraction is overestimated because of too few muons. With less overburden in the vertical region, the probability for muon detection rapidly increases to approximately one with initial muon energy. However, this turnover is faster with the median treatment than the full treatment. Therefore, more muons contribute to the median case calculation, and the passing fraction is underestimated.

In Eq. 6.8 the energy available to other branches of the shower to produce uncorrelated muons is overestimated as some energy must be reserved for producing the electron neutrino or rather the branch of the shower that produces the electron neutrino. Complete modeling of the connection between the three relevant particles (cosmic ray primary, the electron neutrino, and muons) is cumbersome. It would involve

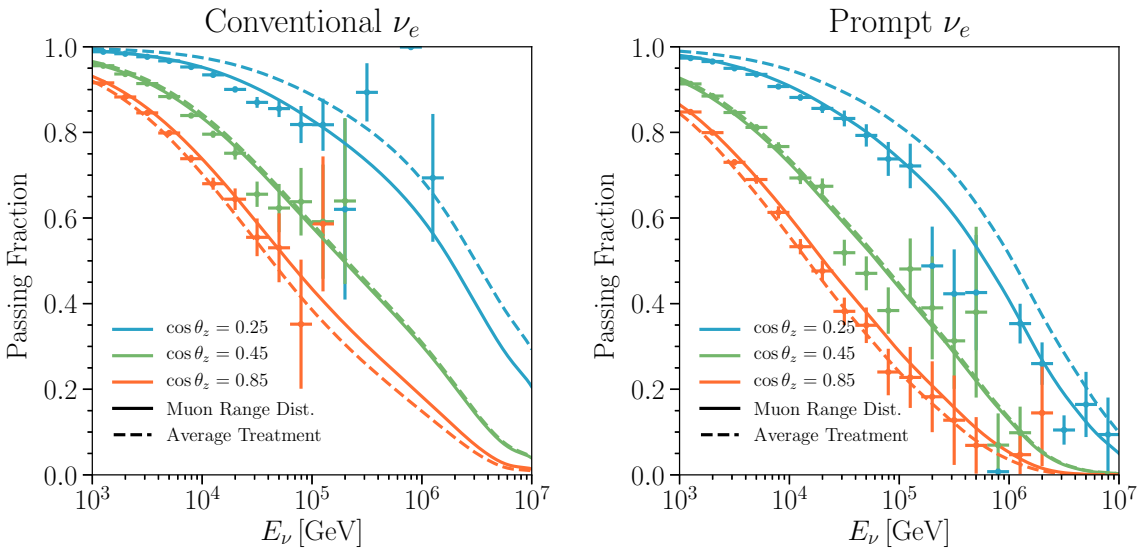

Figure 6.2: ***Passing fractions: effect of the treatment of muon losses in ice.*** Results are shown for three values of $\cos\theta_z$ (from top to bottom): 0.25 (blue), 0.45 (green), and 0.85 (orange); using the full muon range distribution (solid) or the median muon range (dashed). Results from the `CORSIKA` simulation are shown as crosses, with statistical error bars only. In all cases, the H3a primary cosmic-ray spectrum [132], the SIBYLL 2.3 hadronic-interaction model [133, 134] and the MSIS-90-E atmosphere-density model at the South Pole on July 1, 1997 [135, 136] are used. A depth in ice of $d_{\rm det} = 1.95\,\text{km}$ (like the center of IceCube) and a Heaviside $\mathcal{P}_{\rm light}(E_\mu^{\rm f}) = \Theta(E_\mu^{\rm f} - 1\,\text{TeV})$ are assumed. *Left panel:* Conventional ν_e passing fraction. *Right panel:* Prompt ν_e passing fraction.

modeling the shower branch producing the electron neutrino in its entirety, and then separately modeling the uncorrelated branch. However, a simple approximation can be made that at least accounts for the energy necessary to produce the electron neutrino. The electron neutrino must be produced by a parent particle p that has a kinematically allowed energy E_p, so the remaining energy is $E_{\rm CR} - E_p$. The muon yield is then modeled using the average behavior of a shower that begins with energy $E_{\rm CR} - E_p$. Because we are now concerned with the parent particle of the electron neutrino, the slant depth through which this parent particle travels must be considered, as this will affect the probability that it decays to a neutrino instead of interacting. Now considering the parent particle energy and slant depth, the yield of neutrinos from the air shower can be expanded as

$$\frac{dN_{A,\nu}}{dE_\nu}(E_{\rm CR}, E_\nu, \theta_z) = \int dE_p \int \frac{dX}{\lambda_p(E_p, X)} \frac{dN_{p,\nu}}{dE_\nu}(E_p, E_\nu) \frac{dN_{A,p}}{dE_p}(E_{\rm CR}, E_p, X) \,, \tag{6.9}$$

where $dN_{p,\nu}/dE_\nu(E_p, E_\nu)$ is the spectrum of neutrinos from a parent particle p with energy E_p, $dN_{A,p}/dE_p(E_{\rm CR}, E_p, X)$ is the spectrum of parent particles from the air shower at slant depth X, and $dX/\lambda_p(E_p, X)$ is the probability of the parent to decay between the slant depth X and $X + dX$. $\lambda_p(E_p, X)$ is the product of the local density and the parent decay length such that $\lambda_p(E_p, X) = \rho(X)\tau_p E_p/m_p$, where $\rho(X)$ is the local density, τ_p is the lifetime of the parent, E_p is its energy, and m_p is its mass. Now Eq. 6.7 can be rewritten with the above

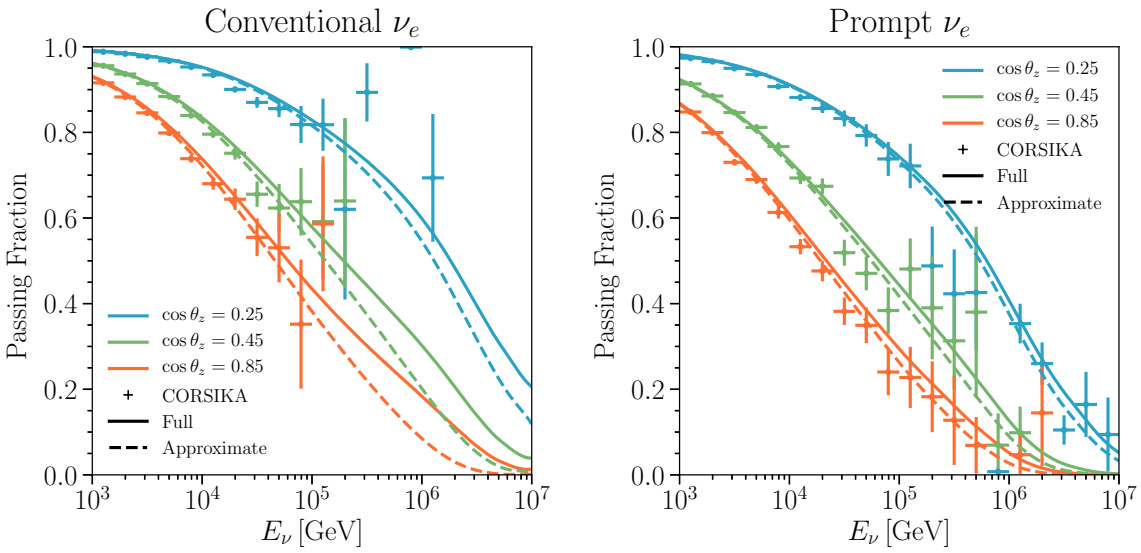

Figure 6.3: ***Passing fractions: effect of approximations on the energy of the shower giving rise to uncorrelated muons.*** Results are shown for three values of $\cos\theta_z$ (from top to bottom): 0.25 (blue), 0.45 (green), and 0.85 (orange); with the approach of this work (solid), Eq. (6.10), where the energy carried by the neutrino parent is subtracted from the rest of the muon producing shower (i.e., $E_{\rm CR} - E_p$), and without this subtraction (dashed), Eq. (6.7), considering the cumulative muon yield from a shower with energy $E_{\rm CR}$. Results from the `CORSIKA` simulation are shown as crosses, with statistical error bars only. *Left panel:* Conventional ν_e passing fraction. *Right panel:* Prompt ν_e passing fraction.

expansion and the simple approximation that accounts for the available energy,

$$\mathcal{P}^{\rm uncor}_{\rm pass}(E_\nu,\theta_z) = \frac{1}{\phi_\nu(E_\nu,\theta_z)} \sum_A \sum_p \int dE_p \int \frac{dX}{\lambda_p(E_p,X)} \int dE_{\rm CR} \frac{dN_{p,\nu}}{dE_\nu}(E_p,E_\nu)\, \frac{dN_{A,p}}{dE_p}(E_{\rm CR},E_p,X)\, \phi_A(E_{\rm CR})\, \mathcal{P}^{\rm shower}_{0\mu}\left(N_\mu = 0; \bar{N}_{A,\mu}(E_{\rm CR}-E_p,\theta_z)\right) \ . \quad (6.10)$$

This treatment remains consistent with the calculation of the neutrino flux mentioned earlier, as if the Poisson probability of muon non-detection is removed (i.e. $\mathcal{P}^{\rm shower}_{0\mu} = 1$) then the passing fraction is 1 by construction.

The naïve approximation of the muon yield using $E_{\rm CR}$ without the subtraction of E_p tends to underestimate the passing fraction, as shown in Fig. 6.3. This underestimation occurs because the muon yield is larger for showers of greater energy. Below 100 TeV, the absolute difference between the two methods is below 0.05; however, at higher energies and especially more vertical directions, the relative difference becomes more important.

Fig. 6.4 shows a comparison between this calculation and results from the CORSIKA simulation in addition to the differences between electron neutrino and antineutrino fluxes. The agreement between simulation and this calculation modeling is excellent for both neutrinos and antineutrinos. The conventional electron antineutrino passing fractions are lower than conventional electron neutrinos, primarily because positively charged mesons are preferentially produced, leading to a harder spectrum than the negatively

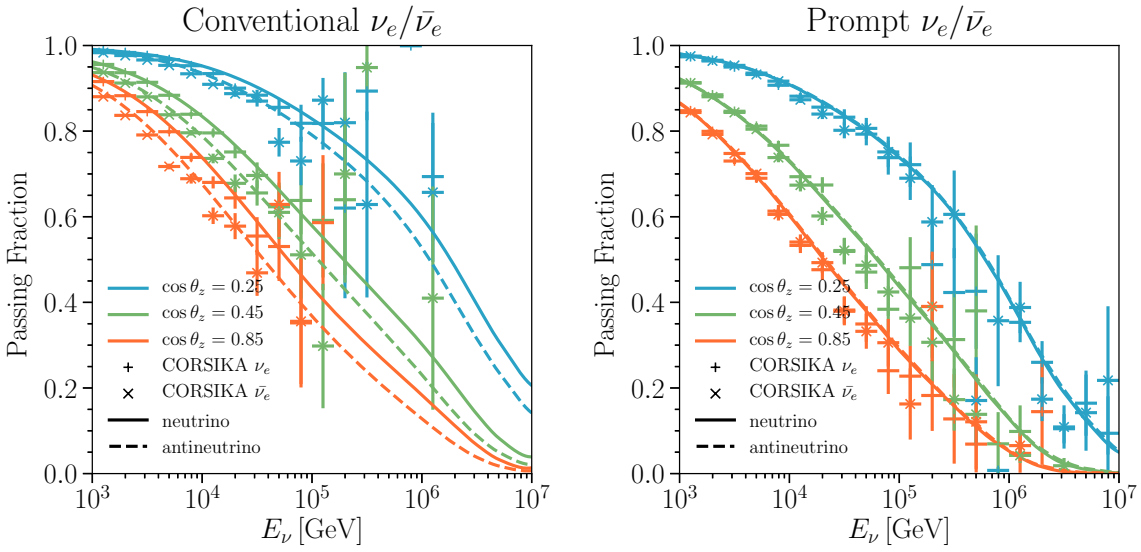

Figure 6.4: ***Passing fractions: neutrinos versus antineutrinos.*** Results are shown for three values of $\cos\theta_z$ (from top to bottom): 0.25 (blue), 0.45 (green), and 0.85 (orange); for neutrinos (solid) and antineutrinos (dashed). Results from the `CORSIKA` simulation for neutrinos (+) and antineutrinos (×) are also shown, with statistical error bars only. In all cases, the H3a primary cosmic-ray spectrum [132], the SIBYLL 2.3 hadronic-interaction model [133, 134] and the MSIS-90-E atmosphere-density model at the South Pole on July 1, 1997 [135, 136] are used, and $d_{\text{det}} = 1.95$ km in ice and $\mathcal{P}_{\text{light}}(E_\mu^{\text{f}}) = \Theta(E_\mu^{\text{f}} - 1\,\text{TeV})$ are assumed. *Left panel:* Conventional ν_e passing fraction. *Right panel:* Prompt ν_e passing fraction. Note that for prompts there are not differences between ν_e and $\bar{\nu}_e$.

charged mesons. The passing fractions for the prompt flux are almost identical as they are mainly produced from gluon fusion [43].

6.1.2.2 ν_μ passing fraction

Muon neutrinos or antineutrinos produced in hadron decays will always have a sibling muon due to flavor conservation. Meaning the passing fraction for muon neutrinos has a correlated component in addition to the uncorrelated component described for electron neutrinos. This correlated suppression depends entirely on the sibling muon's behavior, which itself can be described as a function of the muon direction and initial energy. This section's task is now to track the sibling muons produced along with the neutrino and compute the flux of neutrinos not vetoed by this muon. Additionally, the treatment of correlated and uncorrelated muons can be unified, which we also explore. There is a direct relation between the energy of the parent particle, muon, and neutrino for two-body decays: $E_\mu^{\text{i}} = E_p - E_\nu$. This relation gives a delta function for the muon energy spectrum $dN_{p,\mu}^{2-\text{body}}/dE_\mu^{\text{i}}(E_p, E_\mu, E_\mu^{\text{i}}) = \delta(E_\mu^{\text{i}} - E_p + E_\nu)$. However, the two-body approximation is not appropriate for the calculation of the passing fractions for the prompt fluxes, as neutrinos are mainly produced by the decays of $D^\pm$, D^0, $\bar{D}^0$, Λ_c^+, and $D_s^\pm$ [137]. Treatment of n-body decays requires evaluating the muon distributions from these particles' decays and use of Eq. (6.11). In this work, $dN_{p,\mu}/dE_\mu^{\text{i}}$ was

generated for K_L^0, D^+, D^0, and D_s^+.[1] For general n-body decays the probability of not detecting this sibling muon is a function of the energy spectrum and can be written as

$$\mathcal{P}_{0\mu}^{\text{sib}}\left(\theta_z|E_p,E_\nu\right) = 1 - \int dE_\mu^{\text{i}}\,\mathcal{P}_{\text{det}}\left(E_\mu^{\text{i}},\theta_z\right)\frac{dN_{p,\mu}}{dE_\mu^{\text{i}}}(E_p,E_\nu,E_\mu^{\text{i}})\ . \tag{6.11}$$

Neglecting the uncorrelated muons for a moment, we can think about the correlated passing fraction as a separate calculation. To obtain the correlated passing fraction we can start with the calculation of the neutrino flux, expanding to show the integration over the parent energy and slant depth as was done in Eq. 6.9.

$$\phi_\nu(E_\nu,\theta_z) = \sum_A\sum_p\int dE_p\int\frac{dX}{\lambda_p(E_p,X)}\int dE_{\text{CR}}\,\frac{dN_{p,\nu}}{dE_\nu}(E_p,E_\nu)\,\frac{dN_{A,p}}{dE_p}(E_{\text{CR}},E_p,X)\,\phi_{\text{A}}(E_{\text{CR}}). \tag{6.12}$$

To simplify notation, the flux of parent particles in the air shower at slant depth X defined as

$$\phi_{A,p}(E_p,X) = \int dE_{\text{CR}}\,\frac{dN_{A,p}}{dE_p}(E_{\text{CR}},E_p,X)\,\phi_{\text{A}}(E_{\text{CR}})\ , \tag{6.13}$$

can be separated to obtain a more compact form for the neutrino flux

$$\phi_\nu(E_\nu,\theta_z) = \sum_A\sum_p\int dE_p\int\frac{dX}{\lambda_p(E_p,X)}\,\frac{dN_{p,\nu}}{dE_\nu}(E_p,E_\nu)\,\phi_{A,p}(E_p,X). \tag{6.14}$$

Again, weighting the integrand by the probability of detecting the muon, in this case the sibling muon, the correlated passing fraction can be written as

$$\mathcal{P}_{\text{pass}}^{\text{cor}}(E_\nu,\theta_z) = \frac{1}{\phi_\nu(E_\nu,\theta_z)}\sum_A\sum_p\int dE_p\int\frac{dX}{\lambda_p(E_p,X)}\,\frac{dN_{p,\nu}}{dE_\nu}(E_p,E_\nu)\,\phi_{A,p}(E_p,X)\,\mathcal{P}_{0\mu}^{\text{sib}}\left(\theta_z|E_p,E_\nu\right). \tag{6.15}$$

In the same way that we compared different muon treatments for the uncorrelated passing fraction, Fig. 6.5 shows the effect for the correlated passing fraction. As the correlated case is only concerned with a single muon where energy is correlated with the neutrino, an energy-dependent effect emerges. The muon detection probability is underestimated with the median treatment for lower energies, resulting in a passing fraction that is too high while the converse is true for higher energies. This effect is less pronounces in the horizontal region as the energy correlation becomes more washed out with larger overburden.

As an approximation to the passing fraction, the correlated and uncorrelated passing fractions can be multiplied together like in [131]. Using the median muon behavior, the 2-body approximation, and multiplying the two passing fractions we can obtain the same approximation to the passing fraction as in [131]

$$\mathcal{P}_{\text{pass}}^{\text{GJKvS}}(E_\nu,\theta_z) \equiv \mathcal{P}_{\text{pass}}^{\text{cor,SGRS}}(E_\nu,\theta_z)\,\mathcal{P}_{\text{pass}}^{\text{uncor,GJKvS}}(E_\nu,\theta_z)\ , \tag{6.16}$$

[1] The decay distributions for their antiparticles are identical.

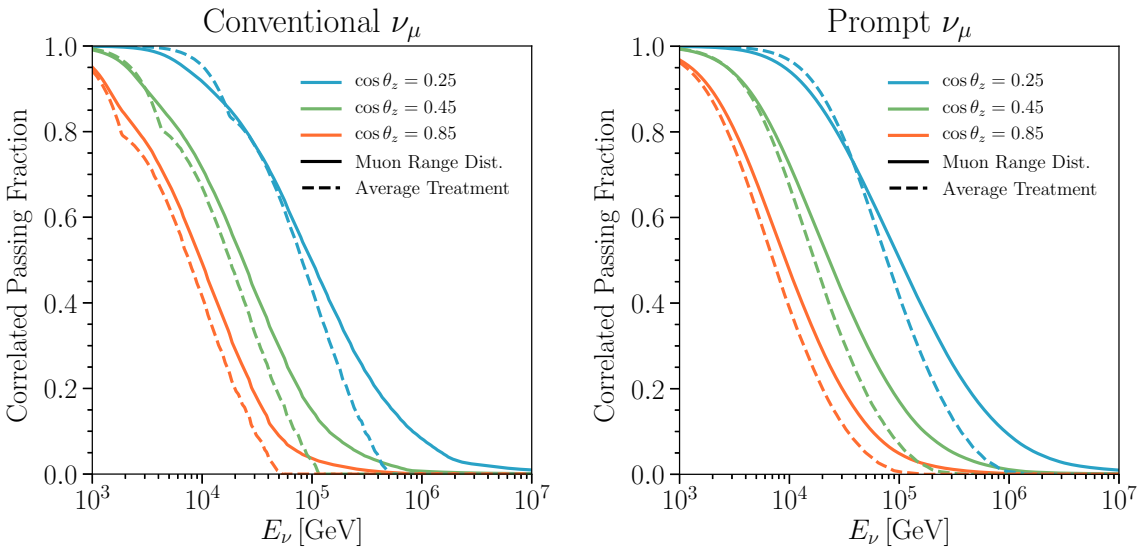

Figure 6.5: ***Correlated passing fractions: effect of the treatment of muon losses in ice.*** Same as Fig. 6.2 but only for the correlated part of the passing fraction for ν_μ.

which contains the correlated passing fraction derived in [130]

$$\mathcal{P}^{\text{cor,SGRS}}_{\text{pass}}(E_\nu,\theta_z) = \frac{1}{\phi_\nu(E_\nu,\theta_z)} \sum_A \sum_p \int dE_p \int \frac{dX}{\lambda_p(E_p,X)} \frac{dN_{p,\nu}}{dE_\nu}(E_p,E_\nu)\, \phi_{A,p}(E_p,X)\, \left[1 - \mathcal{P}^{\text{SGRS}}_{\text{det}}\left(E_p - E_\nu,\theta_z\right)\right] . \tag{6.17}$$

Instead of taking the "factorized" approach in Eq. 6.16, the correlated and uncorrelated equations can be combined at the level of muon detection. In this case, the veto can be thought of as a logical OR between the detection of an uncorrelated muon and the detection of a sibling muon. This logical OR can be achieved in the calculation by the multiplication of the non-detection probabilities $\mathcal{P}_{0\mu} = \mathcal{P}^{\text{sib}}_{0\mu}\mathcal{P}^{\text{shower}}_{0\mu}$. Similarly, weighting the integrand in the neutrino flux calculation by this quantity and maintaining the energy correction described previously, the complete passing fraction is obtained

$$\mathcal{P}_{\text{pass}}(E_\nu,\theta_z) = \frac{1}{\phi_\nu(E_\nu,\theta_z)} \sum_A \sum_p \int dE_p \int \frac{dX}{\lambda_p(E_p,X)} \int dE_{\text{CR}}\, \frac{dN_{p,\nu}}{dE_\nu}(E_p,E_\nu)\, \frac{dN_{A,p}}{dE_p}(E_{\text{CR}},E_p,X)\, \phi_A(E_{\text{CR}})$$
$$\times\, \mathcal{P}^{\text{sib}}_{0\mu}\left(\theta_z|E_p,E_\nu\right)\, \mathcal{P}^{\text{shower}}_{0\mu}\left(N_\mu = 0; \bar{N}_{\mu,A}(E_{\text{CR}} - E_p,\theta_z)\right) . \tag{6.18}$$

The numerator in the above equation corresponds to the passing flux $\phi^{\text{pass}}_\nu(E_\nu,\theta_z)$. This represents the final equation to be used in the passing fraction calculation as all the relevant effects are accounted for, save for the approximations described earlier. Eq. 6.18 naturally applies to muon neutrinos, but also can be directly applied to electron neutrinos for which no sibling muons are present meaning $\mathcal{P}^{\text{sib}}_{0\mu}\left(\theta_z|E_p,E_\nu\right) = 1$. Setting the sibling non-detection probability to 1 recovers the result for electron neutrinos in Eq. (6.10). Expanding

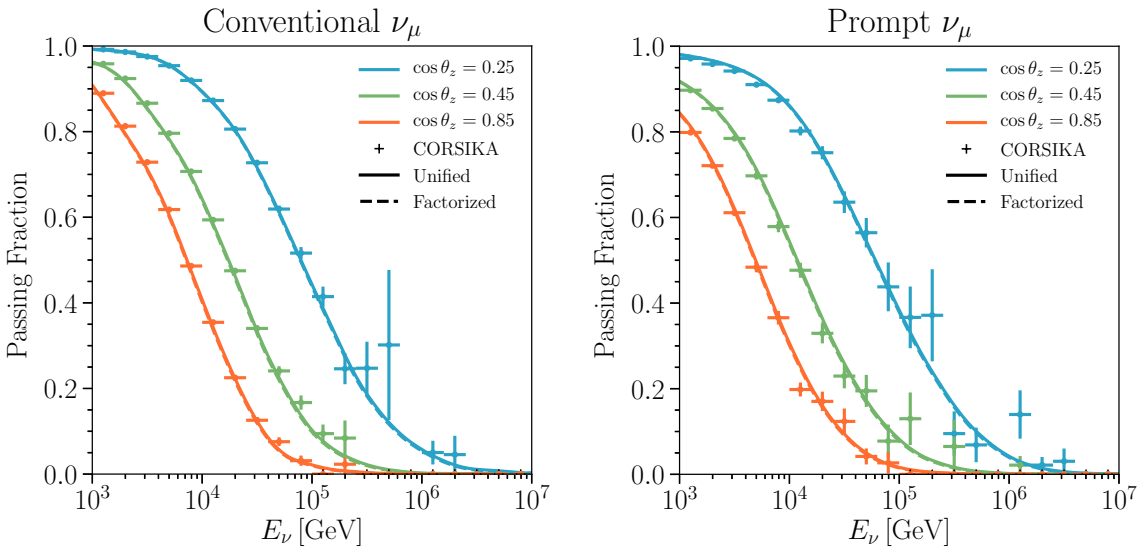

Figure 6.6: ***Passing fractions: differences between the unified and the factorized treatments for ν_μ.*** Same as Fig. 6.3 regarding the approximations on the energy of the shower which gives rise to uncorrelated muons. Comparison of the unified treatment (solid), Eq. (6.18), and the approximate treatment (dashed), Eq. (6.16), which factorizes the correlated and uncorrelated passing fractions. The result is driven by the correlated part.

the probabilities in the above equation, the explicit result is

$$\mathcal{P}_{\rm pass}(E_\nu,\theta_z) = \frac{1}{\phi_\nu(E_\nu,\theta_z)} \sum_A \sum_p \int dE_p \int \frac{dX}{\lambda_p(E_p,X)} \int dE_{\rm CR}\, \frac{dN_{p,\nu}}{dE_\nu}(E_p,E_\nu)\, \frac{dN_{A,p}}{dE_p}(E_{\rm CR},E_p,X)\, \phi_A(E_{\rm CR}) \times \left[1 - \int dE_\mu^{\rm i}\, \frac{dN_{p,\mu}}{dE_\mu^{\rm i}}(E_p,E_\nu,E_\mu^{\rm i}) \int dE_\mu^{\rm f}\, \mathcal{P}_{\rm light}(E_\mu^{\rm f})\, \mathcal{P}_{\rm reach}\left(E_\mu^{\rm f}|E_\mu^{\rm i},\theta_z\right)\right] e^{-N_{A,\mu}(E_{\rm CR}-E_p,\theta_z)}\,, \tag{6.19}$$

The factorized and unified approaches are compared in Fig. 6.6, where one can see that the differences are negligible compared to the other effects discussed in this section. The effect of the energy correction $E_{\rm CR} - E_p$ is smaller for muon neutrinos than in the for electron neutrinos, shown in Fig. 6.3, because for muon neutrinos $\mathcal{P}_{0\mu}^{\rm sib}$ is a much more dominant factor than $\mathcal{P}_{0\mu}^{\rm shower}$. Note, for the uncorrelated part of the passing fraction, this subtraction is more important at higher energies, where the correlated contribution is very small. Thus, the relative effect on the muon neutrino passing fraction is much smaller. Finally, comparisons of the passing fractions calculated for muon neutrinos and muon antineutrinos are shown in Fig. 6.7 and exhibit excellent agreement with the Monte Carlo results, as in the case of electron neutrinos (Fig. 6.4).

A comparison of both the neutrino and antineutrino calculations to results from the CORSIKA simulation is shown in Fig. 6.7. As with the calculation for electron neutrinos, the muon neutrino calculation shows excellent agreement with the air-shower simulations.

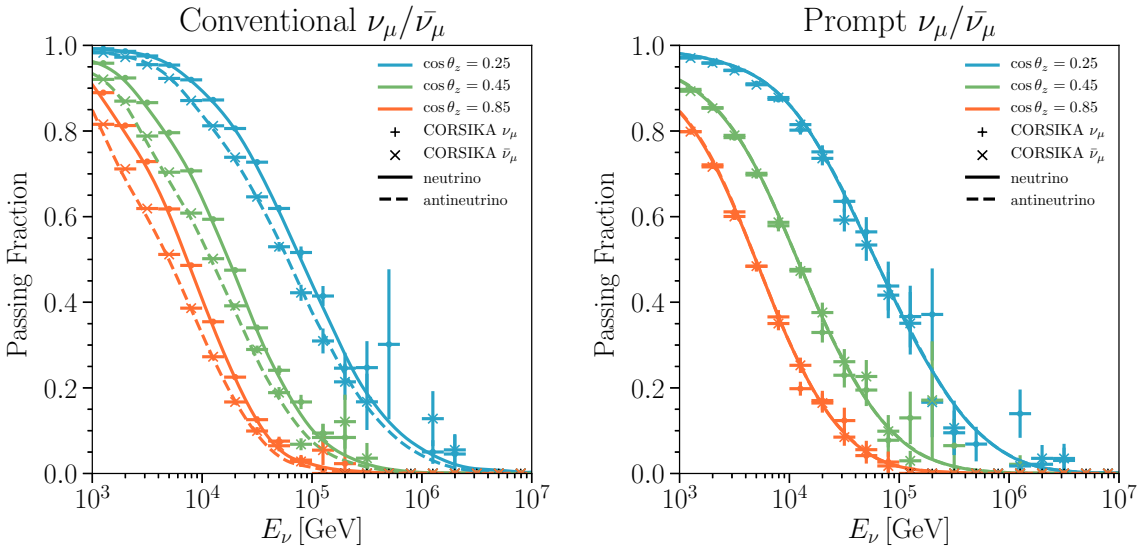

Figure 6.7: ***Passing fractions: neutrinos versus antineutrinos.*** Same as Fig. 6.4 but for ν_μ and $\bar{\nu}_\mu$. Note that for prompt neutrinos there are not differences between ν_μ and $\bar{\nu}_\mu$.

6.1.2.3 Calculation improvements

The calculation described in previous sections has many improvements with respect to earlier treatments. A review of some earlier calculations will help us to examine the differences.

The first proposed calculation of the doing-going atmospheric neutrino suppression due to muons accounted only for sibling muons produced by the same parent meson as the neutrino. The original calculation could then only be applied to atmospheric muon neutrinos [130]. This original calculation used several analytic approximations, applicable to neutrinos produced in pion and kaon decays under the assumption of a power-law cosmic-ray primary spectrum. As described previously, the muon energy loss behavior for this calculation was treated using the median approximation, and the triggering probability modeled as a Heaviside function. These choices allowed analytic passing fractions to be derived.

In [131], the analytic treatment was generalized to include muons from other branches of the shower, not associated with the neutrino production. This generalization allowed the suppression to be computed for atmospheric electron neutrinos. Analytic approximations were used to describe the average properties of muons from air showers. In fact, the calculation in [131] fit a parameterized $\bar{N}_{A,\mu}(E_{\rm CR}, \theta_z)$ to the results of CORSIKA simulation which used SIBYLL 2.1 for conventional neutrinos and DMPJET-2.5 for prompt neutrinos. The combination of these correlated and uncorrelated contributions led to Eq. 6.16. For these two calculations, the components were factorized, and the $E_{\rm CR} - E_p$ subtraction was not performed. This led to an overestimation of the shower energy that produces visible muons in the electron neutrino case.

The new approach explained in the previous sections and developed in [65] includes a full description of

the cosmic ray primary spectrum and composition as well as the hadronic interactions that give rise to the parent particles and eventually muons and neutrinos. This leads to the more accurate Eq. 6.18, which is fully consistent between electron and muon neutrinos, although it must be calculated numerically. Muon energy loss distributions are also fully accounted for in this treatment. Critically, this approach uses numerical techniques instead of analytic approximations, opening the door for a wide array of modifications to the calculation. These modifications can include: different detector configurations and responses, cosmic ray spectra, hadronic interaction models, atmospheric density models, and muon-energy-loss cross sections. With these modifications, the passing fraction calculation can be made consistent with the modern atmospheric neutrino flux calculations of `MCEq`.

Fig. 6.8 shows a direct comparison between the new calculation and its predecessor [131] for ν_e and ν_μ from the conventional and prompt fluxes in two different directions. In the more vertical direction, ν_e and ν_μ passing fractions for the conventional flux show a large difference, where the newer calculation provides a higher passing fraction. Two effects are important here, the muon range treatment and the energy of the uncorrelated shower branch. Near the horizon, these effects partially cancel, minimizing the difference between the two calculations, but that is not the case in the vertical direction. The shoulder present in the older calculation shows the transition between the pion and kaon dominated production of neutrinos. However, this is washed out by the more detailed muon range treatment of the new calculation. The differences in the prompt ν_e curves are more difficult to interpret. Comparisons of the muon range treatment and the energy subtraction for prompt ν_e would lead us to expect some small differences, but not quite of this magnitude. One additional factor that could explain this is the use of DPMJET-2.55 in the old calculation for prompt neutrinos; we see that DPMJET-2.55 gives larger passing fractions than SIBYLL 2.3c. Differences in the prompt ν_μ passing fractions are explained by the fact that [131] applies the conventional ν_μ passing fractions to prompt ν_μ.

6.1.2.4 Calculation systematics

The calculation outlined in these sections relies on a host of external information such as the cosmic ray flux, the physics governing decays and hadronic interactions, calculations of muon-energy-loss cross sections, and the detector response. In this section, a few of these potential sources of uncertainty are examined.

6.1.2.4.1 Muon Energy Losses Muons lose energy in a medium through three processes: ionization losses, e^+e^- pair production, bremsstrahlung, and photo-nuclear interactions. Below $\sim 1\,\text{TeV}$ losses are dominated by ionization, whereas, at higher energies, the radiative processes dominate. In order of importance, these are pair production, bremsstrahlung, and photo-nuclear interactions. Above $\sim 10\,\text{PeV}$ photo-nuclear

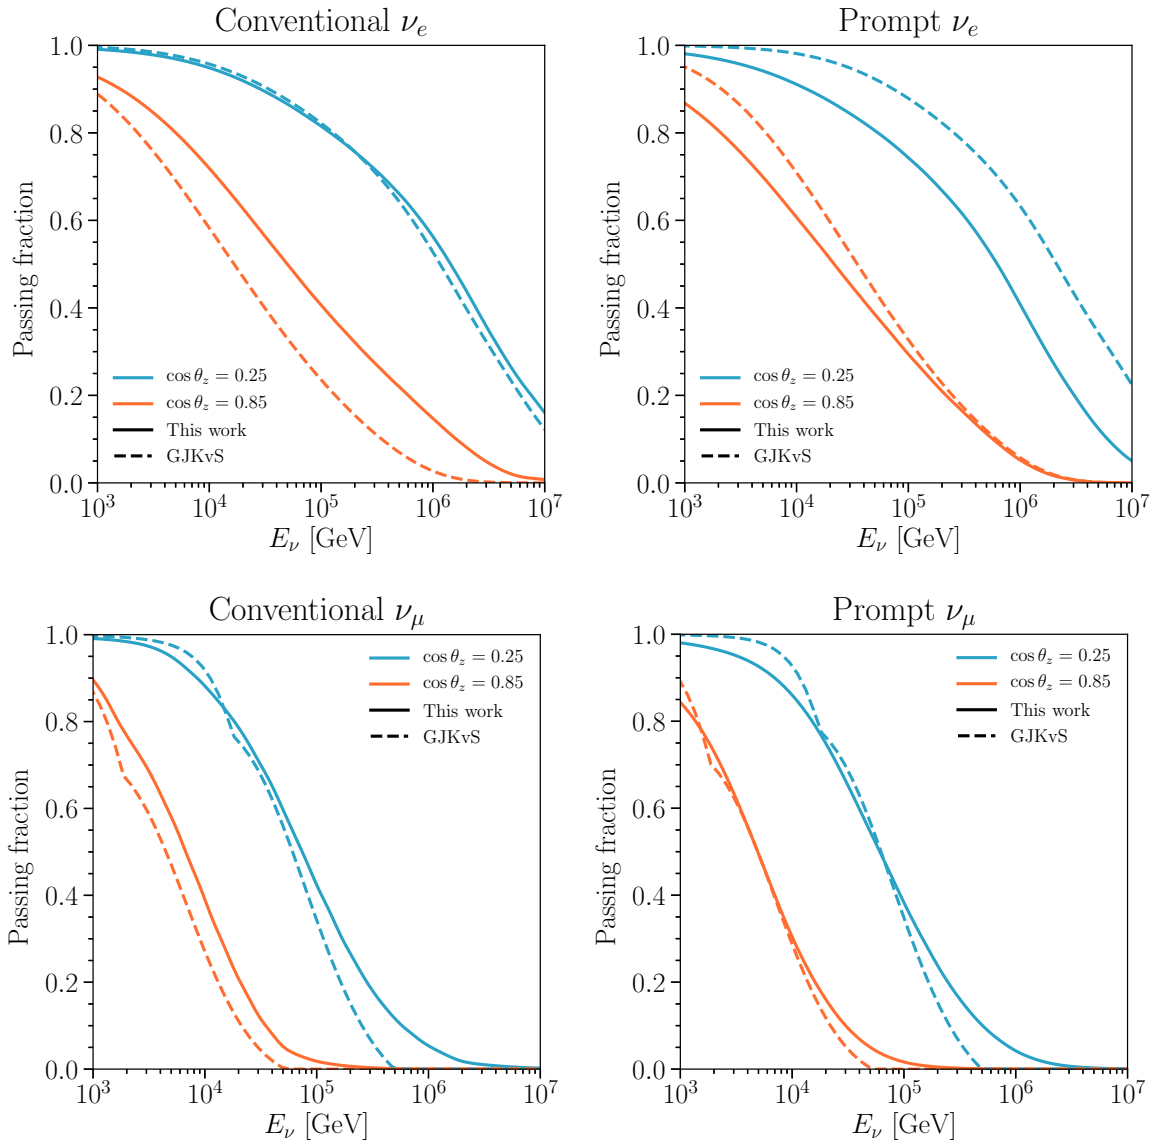

Figure 6.8: ***Passing fractions: comparison with previous work***. Results are shown for two values of $\cos\theta_z =$ (from top to bottom): 0.25 (blue) and 0.85 (orange); with the calculation in this work (solid) and with that in Ref. [131] (dashed). Our results are obtained with the H3a primary cosmic-ray spectrum [132], the SIBYLL 2.3c hadronic-interaction model [138], the MSIS-90-E atmosphere-density model at the South Pole on July 1, 1997 [135, 136], and assuming $d_{\rm det} = 1.95$ km in ice and $\mathcal{P}_{\rm light}(E_\mu^{\rm f}) = \Theta(E_\mu^{\rm f} - 1\,{\rm TeV})$. They include all the effects discussed in previous sections, Eq. (6.18). *Top-left panel:* Conventional ν_e passing fraction. *Top-right panel:* Prompt ν_e passing fraction. *Bottom-left panel:* Conventional ν_μ passing fraction. *Bottom-right panel:* Prompt ν_μ passing fraction.

interactions are comparable to bremsstrahlung [139]. These processes are well understood up to $\sim 10\,\text{TeV}$ in muon energy, above which the photo-nuclear interactions dominate the uncertainty in the cross section.

The photo-nuclear cross section used in MMC/PROPOSAL by default uses a data-driven parameterization of the proton structure function in a deep-inelastic scattering (DIS) formalism [140, 141]. This formalism includes contributions from both soft (non-perturbative) and hard (perturbative) physics [142]. Another approach describes the cross section with a generalized vector dominance model for the soft component [143] and a framework of the color dipole moment for the hard component [144, 145].

Without appropriate measurements and the corresponding data-derived uncertainties the cross section uncertainty can be examined by comparing the two calculation approaches. The default cross sections are slightly lower at the energies relevant for the passing fraction calculation, resulting in smaller passing fractions as more muons can reach the detector with higher energy. These differences in the cross section result in a maximum difference of 0.01 for the passing fractions.

6.1.2.4.2 Primary Cosmic Ray Spectra Previously shown comparisons all assume the Hillas-Gaisser three population model (H3a) [132]. This model considers five different nuclei mass groups with three populations. These cut off at a characteristic rigidity. Other models for the cosmic ray primary flux have been proposed, and current measurements leave a considerable amount of uncertainty in the energy regime relevant for the passing fractions.

A second model, GST-4gen [146], has a lower cutoff for the first two populations, which also have correspondingly harder spectra, while the third population is iron-only above the spectrum ankle. A fourth proton-only component is included to obtain better agreement with shower depth data around 1 EeV. However, this proton-iron composition is somewhat in tension with data from Auger [147, 148].

A third model, ZS [149], has three components derived from specific astrophysical sources. The lowest energy population comes from the expanding shells of supernovae remnants. The middle energy component comes from isolated supernovae and their interaction with the interstellar medium. Finally, the high energy component comes from massive star supernovae explosions and their interaction with their own stellar wind. This last process produces a very heavy composition.

Fig. 6.9 compares the effect these three models have on the passing fractions. The largest differences are for the ν_e passing fractions, which differ between models by less than 0.05, except when comparing the ZS model above $10^6\,\text{GeV}$. This can be understood as a consequence of the ZS model being designed only to describe the cosmic ray flux up to $10^8\,\text{GeV}$, meaning the ZS neutrinos flux predictions may not be accurate above $10^6\,\text{GeV}$. The muon neutrinos passing fractions have much smaller variations between models. This is likely a result of the sister muons dominating the passing fraction calculation, meaning that the population

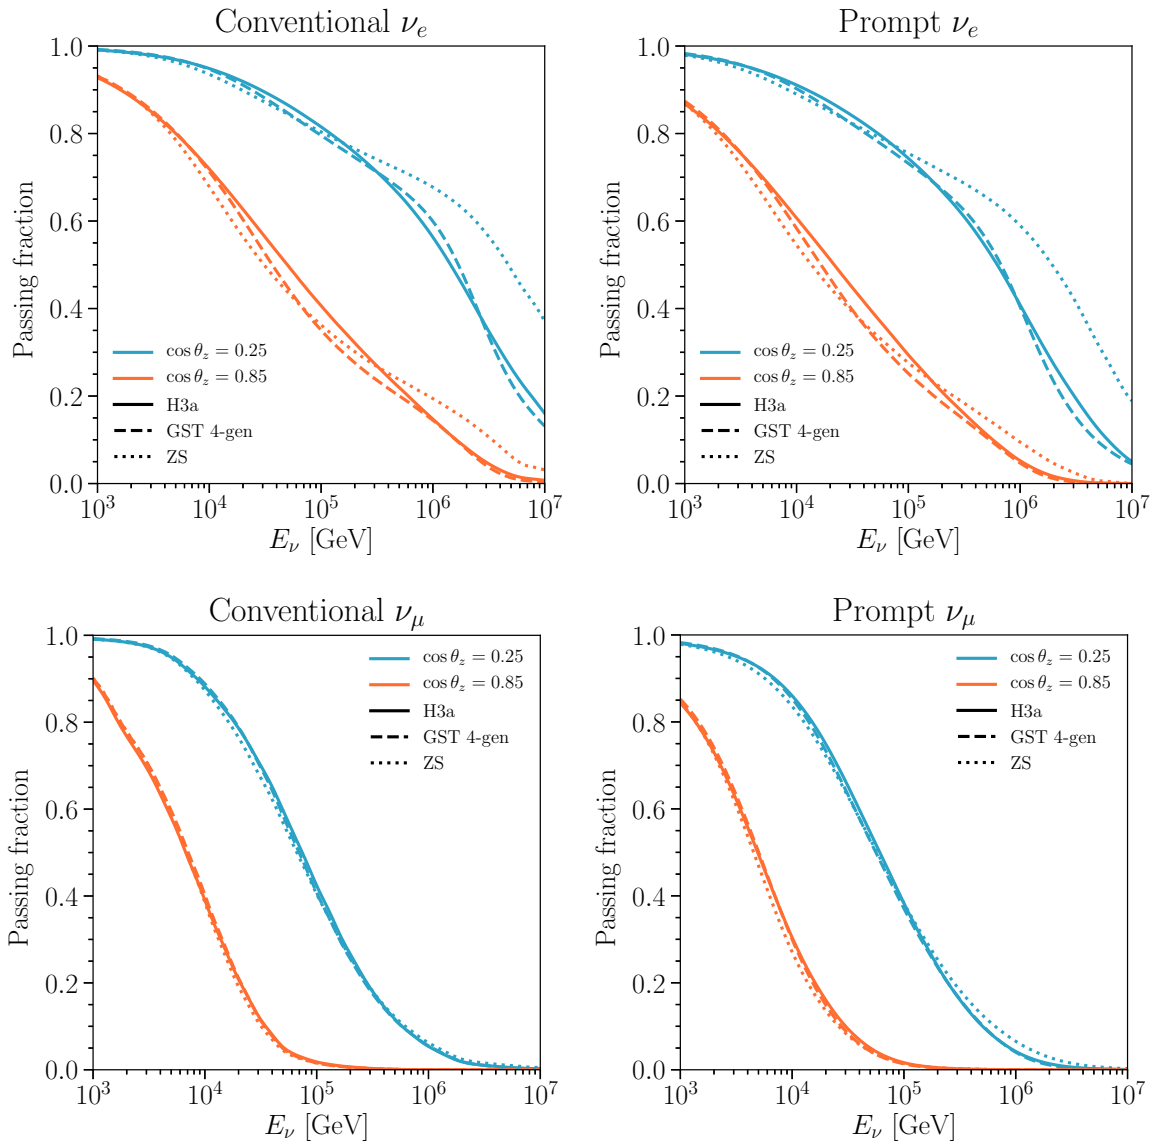

Figure 6.9: ***Passing fractions: effect of primary cosmic-ray spectrum***. Results are shown for two values of $\cos\theta_z$ (from top to bottom): 0.25 (blue) and 0.85 (orange); for four CR models: Hillas-Gaisser H3a [132] (solid), Gaisser, Stanev and Tilav (GST 4-gen) [146] (dashed) and Zatsepin-Solkolskaya (ZS) [149] (dotted). *Top-left panel:* Conventional ν_e passing fraction. *Top-right panel:* Prompt ν_e passing fraction. *Bottom-left panel:* Conventional ν_μ passing fraction. *Bottom-right panel:* Prompt ν_μ passing fraction.

of showers from which the neutrino may have originated is a subdominant effect. The level of variation seen between these models represents a non-negligible uncertainty in the passing fractions for electron neutrinos but is small for muon neutrinos. However, this level of uncertainty alone is likely to be smaller than the statistical uncertainties currently present for the HESE sample.

Here a bracketing approach was used to examine the uncertainties and models chosen to express a range of possibilities currently allowed by experimental measurements. However, this is not a full accounting of the cosmic ray flux uncertainties. The measurement uncertainties can be accounted for in a more satisfactory manner with a non-parametric fit with sufficient degrees of freedom to the available cosmic-ray data. For any point in the parameter space, both the neutrino flux and passing fraction can be computed to provide an expectation for the "apparent" neutrino flux. By directly using estimates from this fit and allowing the cosmic ray model to vary, the uncertainties can be treated appropriately in a combined fit of neutrino data. This method has some computational challenges as it is not feasible to run the necessary MCEq calculations each time a different point in the parameter space is tested. However, some optimizations can be found by carefully examining the basis functions used for the different cosmic ray populations and pre-computing or caching information where applicable. Such a treatment is not used for the HESE 7.5 year analysis but may prove to be useful for future analyses that combine multiple samples with more than 10 years of data.

6.1.2.4.3 Hadronic Interaction Models The muons and neutrinos relevant for this calculation are produced in high energy extensive air showers, the development of which is governed by hadronic interaction physics. However, these showers and hadronic interactions occur at energies that are orders of magnitude beyond what is accessible to collider experiments where these interactions can be measured. Thus, our predictions for how these showers develop depend heavily on hadronic models that may are tuned to collider data but are extrapolations at the energies we are concerned with.

Figure 6.10 shows a comparison of the passing fractions computed, assuming four different hadronic interaction models. A detailed discussion of these models' differences can be found in Section IV.C of [65]. There are non-negligible differences in the passing fractions for electron neutrinos, but the variation is again small for muon neutrinos.

6.1.2.4.4 Depth and Surrounding Medium In these sections, a depth of 1.95 km has been assumed in order to model the veto effects at the center of the IceCube detector. However, IceCube and other neutrino detectors are large enough that the change in depth from top to bottom significantly changes the overburden for a given zenith angle. Increases in the overburden significantly reduce the veto's power, as muons that reach the detector are fewer and less energetic. As a benchmark, we can examine the passing fractions assuming a depth of 3.5 km, corresponding to the depth of KM3NeT-Italy, and compare differences between ice and water.

Figure 6.11 shows these comparisons. Notably, the differences between water and ice are non-negligible but much less dramatic than the effect of depth. As expected, the surrounding medium's effect is also more

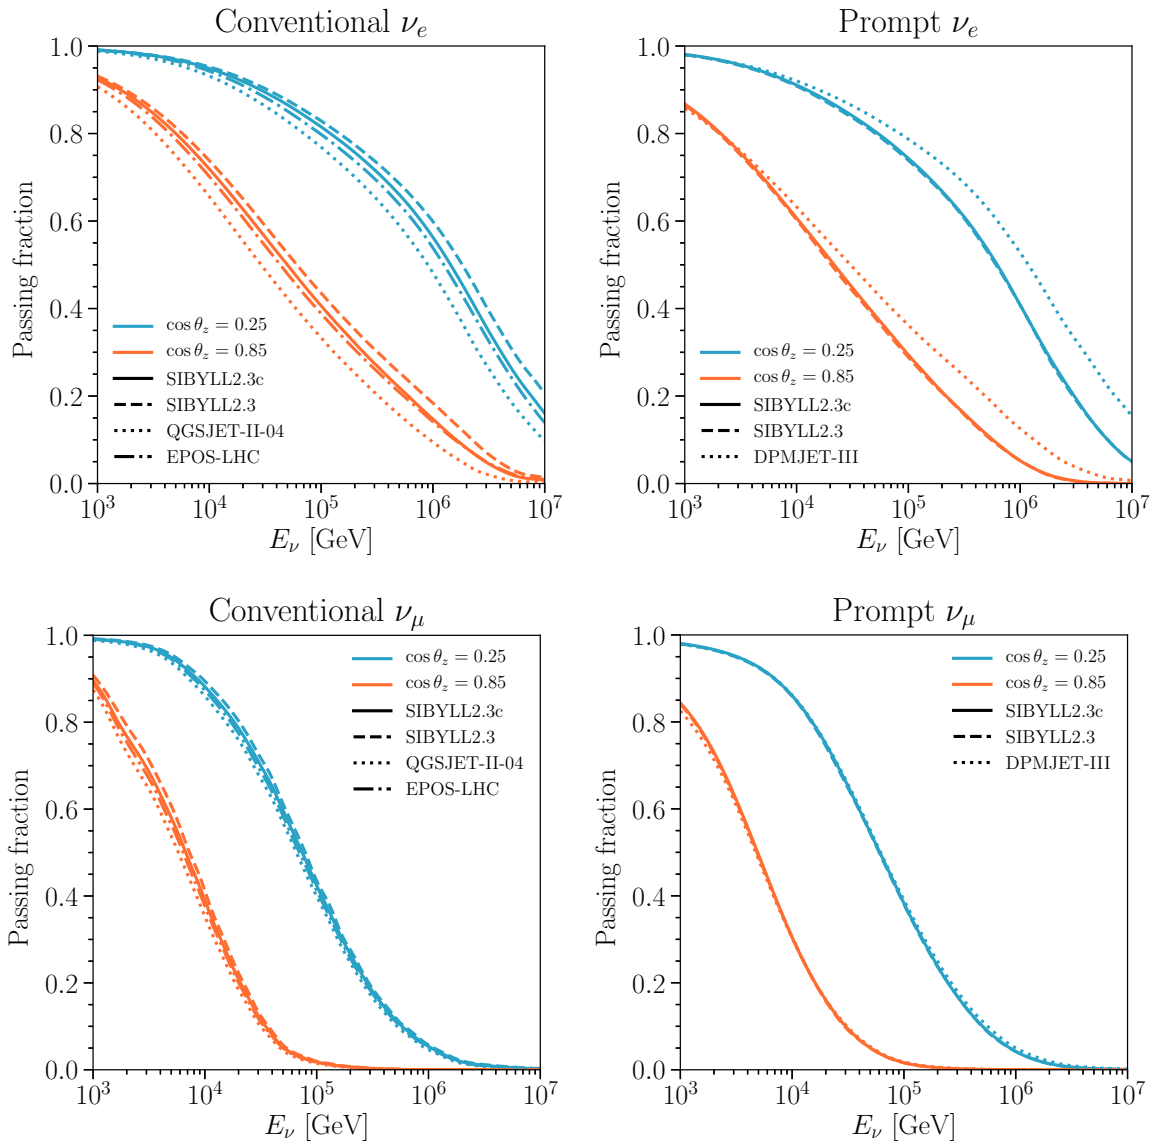

Figure 6.10: ***Passing fractions: effect of hadronic-interaction model***. Results are shown for two values of $\cos\theta_z$ (from top to bottom): 0.25 (blue) and 0.85 (orange); for different hadronic-interaction models: SIBYLL 2.3c [138] (solid), SIBYLL 2.3 [133, 134] (dashed), QGSJET-II-04 [150] (dotted in left panel), EPOS-LHC [151] (dash-dotted in left panel) and DPMJET-III[152] (dotted in right panel). *Top-left panel:* Conventional ν_e passing fraction. *Top-right panel:* Prompt ν_e passing fraction. *Bottom-left panel:* Conventional ν_μ passing fraction. *Bottom-right panel:* Prompt ν_μ passing fraction.

pronounced near the horizon, where the change in effective overburden is larger. The increased depth results in a significantly larger passing fraction, greatly reducing the power of veto techniques. It is, therefore, important to model the depth of events to determine their passing probability accurately. Depth is also an important consideration when evaluating the sensitivity of different neutrino detectors to the astrophysical

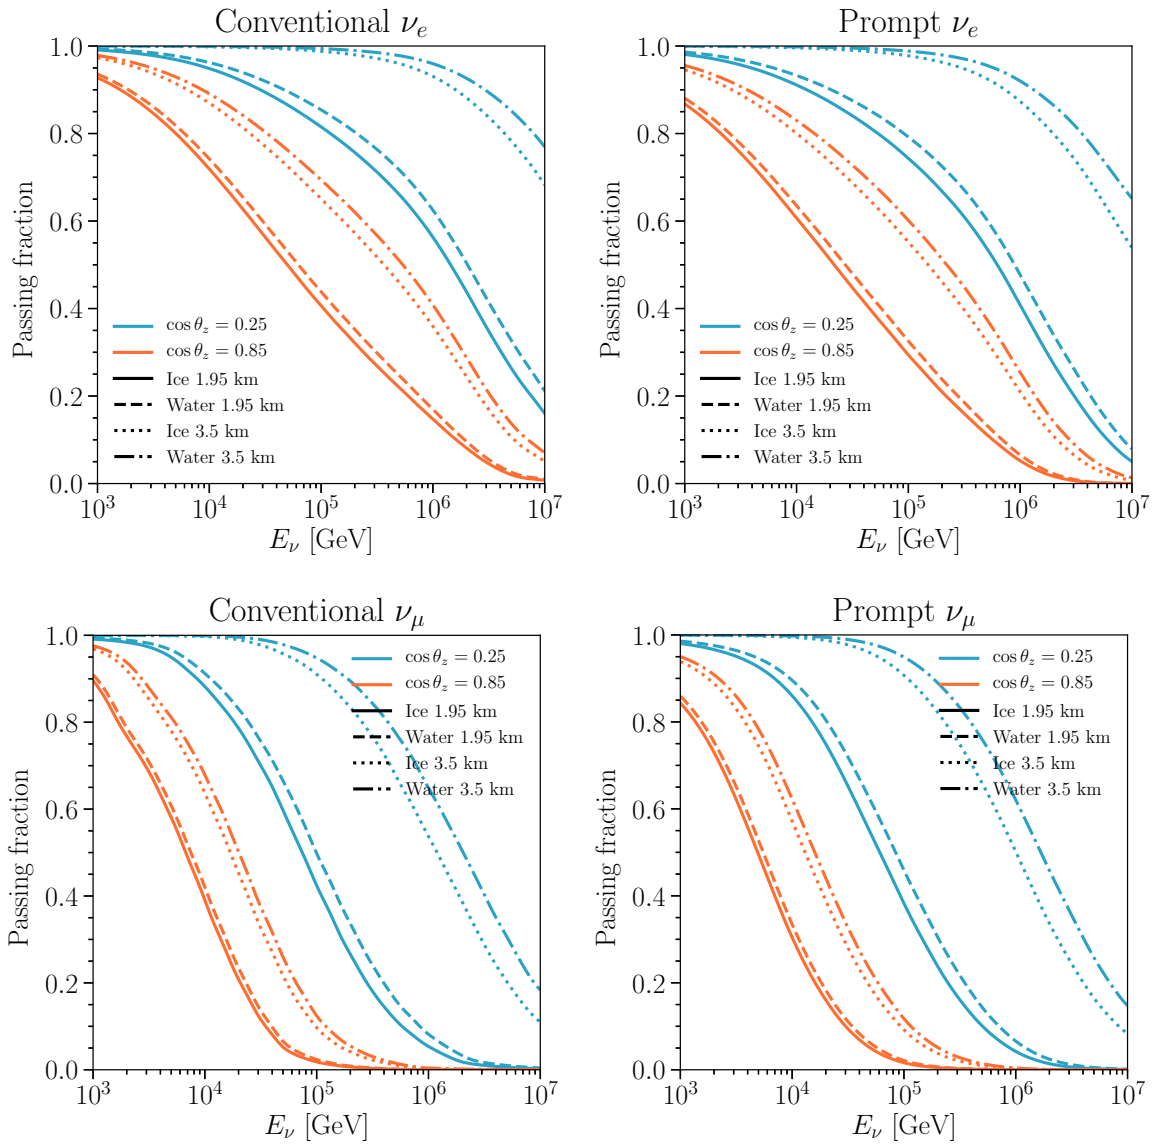

Figure 6.11: ***Passing fractions: effect of depth/medium***. Results are shown for two values of $\cos\theta_z$ (from top to bottom): 0.25 (blue) and 0.85 (orange); for two depths: $d_{\text{det}} = 1.95$ km (solid and dashed) and $d_{\text{det}} = 3.5$ km (dotted and dash-dotted); and for two different media: ice (solid and dotted) and water (dashed and dot-dashed). *Top-left panel:* Conventional ν_e passing fraction. *Top-right panel:* Prompt ν_e passing fraction. *Bottom-left panel:* Conventional ν_μ passing fraction. *Bottom-right panel:* Prompt ν_μ passing fraction.

flux.

6.1.2.4.5 Detector Response The probability of a muon to trigger the veto, $\mathcal{P}_{\text{light}}$, encapsulates the detector response relevant to the passing fraction calculation. So far, in this discussion, all the passing

fractions have been computed with a 1 TeV threshold Heaviside $\mathcal{P}_{\text{light}}$. While the Heaviside parameterization is simple, the real detector response is undoubtedly more complex. The particular form of $\mathcal{P}_{\text{light}}$ will depend on the particulars of the implemented veto. To get a feel for these differences and the uncertainty that detector response modeling could introduce, we look at a few different $\mathcal{P}_{\text{light}}$ functions.

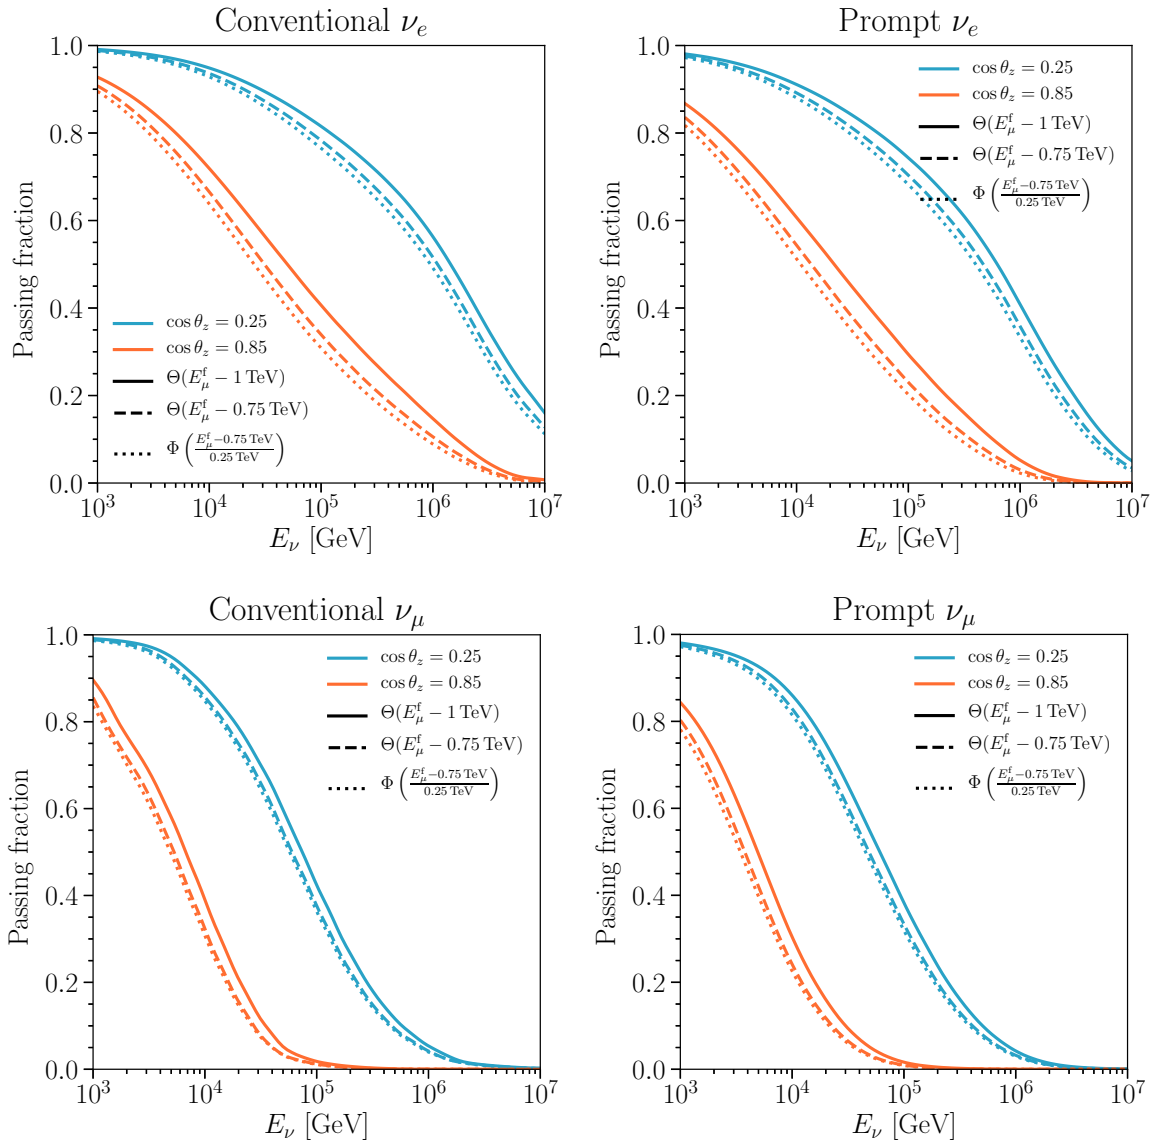

Figure 6.12: ***Passing fractions: effect of $\mathcal{P}_{\text{light}}$***. Results are shown for two values of $\cos\theta_z$ (from top to bottom): 0.25 (blue) and 0.85 (orange); for three different $\mathcal{P}_{\text{light}}(E_\mu^{\text{f}})$: a Heaviside with a muon threshold at 1 TeV, $\Theta\left(E_\mu^{\text{f}} - 1\,\text{TeV}\right)$ (solid), a Heaviside with a muon threshold at 0.75 TeV, $\Theta\left(E_\mu^{\text{f}} - 0.75\,\text{TeV}\right)$ (dashed) and a sigmoid $\Phi\left(\left(E_\mu^{\text{f}} - 0.75\,\text{TeV}\right)/0.25\,\text{TeV}\right)$ (dotted). *Top-left panel:* Conventional ν_e passing fraction. *Top-right panel:* Prompt ν_e passing fraction. *Bottom-left panel:* Conventional ν_μ passing fraction. *Bottom-right panel:* Prompt ν_μ passing fraction.

With straightforward veto implementations, a smooth, monotonic, threshold-like behavior is expected. This can be understood as a consequence of three factors: cuts on observable parameters, fluctuations in observed light, and the positive correlation between muon energy and light emission. The cuts on observable parameters introduce the threshold behavior, choosing to accept and reject events that fall into different observable parameter space regions. However, this hard threshold is smoothed out by fluctuations in the observable parameters. Stochastic energy losses of the muons mean that the emitted light near the veto can vary wildly for identical muons; further variation is caused by photon scattering and the quantum efficiency of the detector PMTs. The Heaviside function models this threshold behavior, but transitions too sharply represent a realistic detector response. An alternative is to use a sigmoid to model the smooth transition; in this case, a logistic function is a reasonable choice. Figure 6.12 shows the passing fractions for three variations of the $\mathcal{P}_{\text{light}}$ function, a Heaviside with 1 TeV threshold, a Heaviside with 0.25 TeV threshold, and finally a logistic function with a 0.25 TeV threshold and 1/0.25 TeV growth parameter. A reduction in the Heaviside function threshold produces a reduction of the passing fraction as expected. The change from Heaviside to sigmoid produces a smaller reduction, although this comparison of the effect's magnitude is arbitrary. Although the change to muon neutrino passing fractions is smaller than that for electron neutrinos, this is the largest change in the muon neutrino passing fraction for a fixed detector geometry. We should stress the importance of correctly modeling the detector response to muons as it will significantly affect both flavors' passing fractions.

6.1.2.5 HESE passing fractions

In this approach, the neutrino properties are known from simulation, and it is sought to average over all the potential properties of the cosmic ray air showers from which the neutrino could have been produced. Ideally, this average over air shower properties should be computed for each neutrino position, direction, and energy because the detector response can vary with all six of these parameters. However, not all of these properties are used directly in the analysis, nor does the detector response depend equally on all of these parameters. For this reason, when performing the efficiency calculation, only the neutrino energy, zenith angle, and depth upon intersection with the detector are considered. Other properties of the neutrino are averaged over. Additionally, in the characterization of the detector response to muons, only dependence on the muon energy and depth are considered. Thus, the computed passing fraction depends on the neutrino energy, the zenith angle, and the incident depth in the detector. However, the detector response and neutrino properties can be factorized to be discussed more generally. In previous analyses, the passing fractions were calculated using an extension of the method described in [130] and bounded at 10 %; details of the method are provided in [59]. Cosmic ray simulations remain a computationally prohibitive way of accounting for the

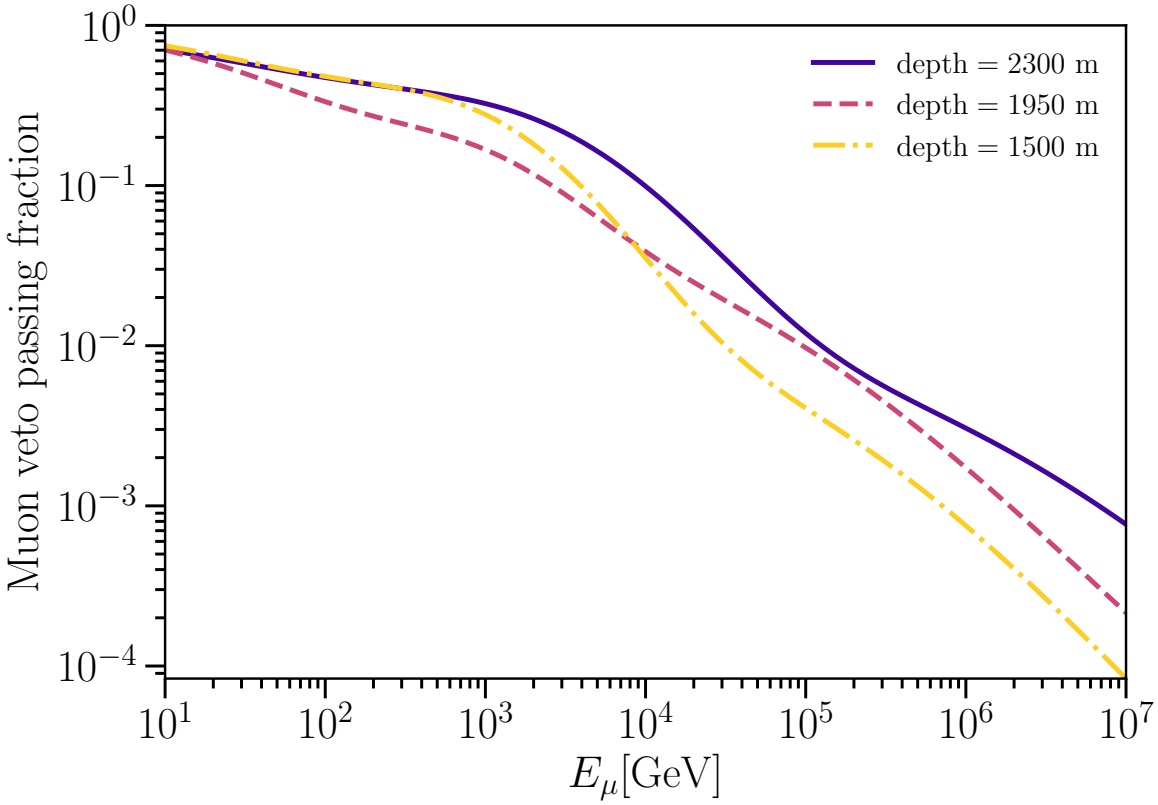

Figure 6.13: ***Muon veto passing fraction.*** Each line shows the fraction of muons of a given energy at the detector edge, E_μ, that pass without triggering the veto when entering the detector at a particular depth. Three depths are shown: 1500, 1950, and 2300 meters from the surface, with lines of darkening color as the depth increases. The veto efficiency increases with the muon energy. Differences at various depths are due to the changing ice properties, and varying acceptance as a function of depth due to the veto region's asymmetric structure. At all depths, a sigmoid function is fit to the results of muon simulation. Above $\sim 100\,\text{TeV}$ the passing fraction is extrapolated.

effects of accompanying muons, so we still rely on calculations of the average passing rate. In this analysis, we use the new calculation described above and in [65] that allows for different cosmic-ray and hadronic models to be used; more importantly for this analysis, any parameterization of the detector veto response to muons can be used in the calculation, as opposed to just an energy threshold. This capability allows us to model the detector response to atmospheric neutrinos more accurately. In Fig. 6.13, we show the probability that a muon will pass the veto as a function of the true muon incident energy for different detector depths.

Using the muon passing fractions in Fig. 6.13 as input and the `nuveto` code provided in [65] the atmospheric passing fraction is calculated for each component and flavor using the Hillas-Gaisser H3a [146, 132, 153] model for the incident cosmic-ray spectra and SIBYLL2.3c [138] for the hadronic interactions in the air shower. Switching to passing fractions derived from alternative cosmic-ray and hadronic interaction models

has sub-leading effects in determining the astrophysical flux [65]. The effects of these systematics were studied by repeating the analysis for different passing fractions that arise from a given combination of cosmic-ray spectrum and hadronic model for various spectra and models available in the literature. We found that the inclusion of these effects in addition to other discrete ice choices mentioned later in Section 3.2 increases the reported uncertainty of the astrophysical parameters by at most 20 % with respect to errors computed without these effects. For this reason, these effects are not included in the analysis and are not reflected in the reported errors of any model parameters. In Fig. 6.14, we show the passing fractions for the conventional and prompt neutrino components. In these figures, the left, center, and right panels correspond to $\cos\theta_z$ values of 0.1, 0.3, and 0.9 respectively; the solid lines correspond to muon neutrinos and the dashed lines to electron neutrinos. From the panels' progression from left to right, one can see the passing fractions become smaller as one approaches vertical directions. Vertical muons have the highest probability of reaching the detector, as the overburden they pass through is the smallest. Though not shown in this figure, the conventional passing fractions differ from neutrinos to anti-neutrinos, see [65] for details; the appropriate passing fractions are used in this analysis. Figs. 6.15 and 6.16 show the distributions of conventional and prompt neutrinos respectively after this correction is applied. This reduction in atmospheric background accounts for much of this analysis's sensitivity to the astrophysical neutrino flux, as the observed down-going atmospheric fluxes in IceCube would otherwise be comparable in magnitude and remain similar in their angular distribution. This is best seen when comparing the atmospheric fluxes before and after the veto to the measured astrophysical flux as shown in Fig. 6.18.

6.2 Muon background estimation

Finally, there is also the possibility of single muons that trigger the event selection without a neutrino interaction in the detector and still pass the veto. The shapes of the atmospheric muon and neutrino fluxes are closely related to each other and bounded by the cosmic-ray flux so that they must be steeply falling. The interaction of muons in the atmosphere and ice further softens the muon spectrum from that of cosmic rays. Although there is uncertainty in the shape of the muon spectrum, the yield of muons from cosmic-ray air showers has more significant modeling uncertainties that stem from uncertainties in the hadronic interaction cross sections [154] and the cosmic-ray composition [155]. As we lack the capability to parameterize both the uncertainty in shape and normalization from first principles, we turn to data-driven techniques to constrain the size of this background. Unfortunately, the data-driven techniques available do not provide us with enough events to determine the muon background's shape. For this reason, we take a pragmatic approach to treat the muon component. We use a simulation estimate of the muon flux shape, which provides a reasonable

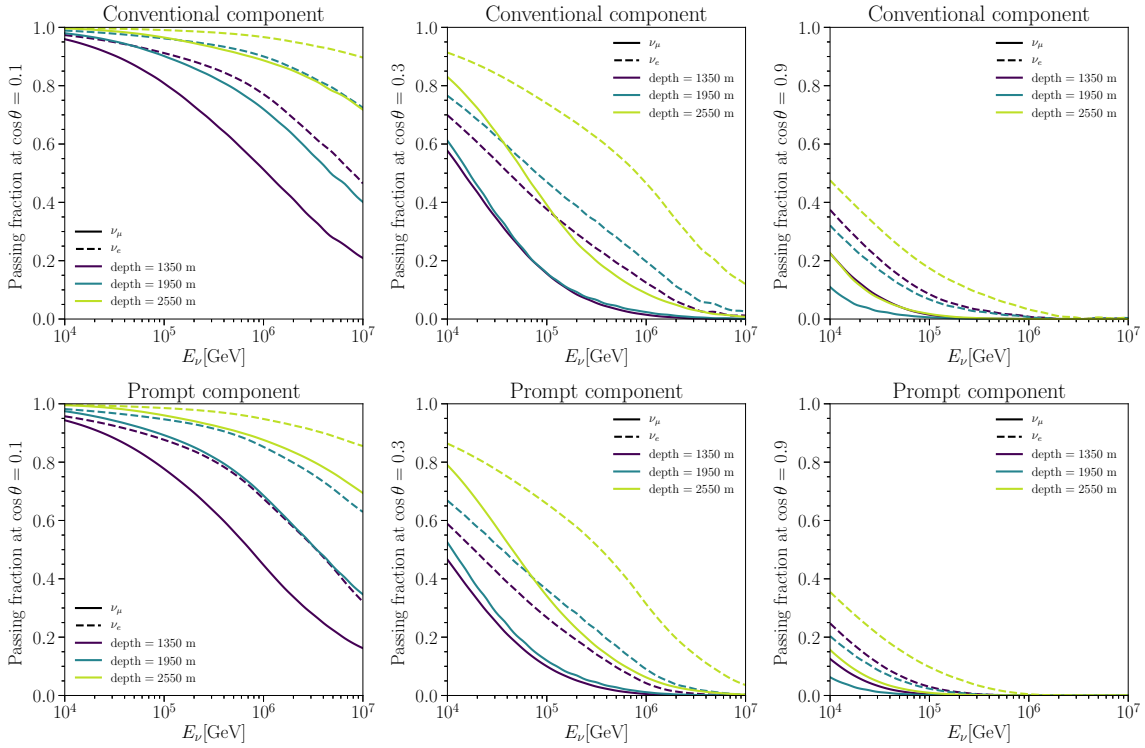

Figure 6.14: ***Conventional and prompt atmospheric component passing fraction.*** The top row of plots shows the atmospheric neutrino passing fraction as a function of the neutrino energy for a flux of neutrinos originating from pions and kaons, assuming the Hillas-Gaisser H3a [146, 132, 153] cosmic-ray model and SIBYLL 2.3c [138] hadronic interaction model. The bottom row of plots shows the atmospheric neutrino passing fraction for a flux of neutrinos originating from charmed hadrons under the same assumptions. Solid lines correspond to muon neutrinos and dashed lines to electron neutrinos. From darkest to lightest, the different colors are for three different detector depths: 1350, 1950, and 2550 meters below the surface. The left, center, and right panel correspond to cosine of the zenith angles 0.1, 0.3, and 0.9 respectively (or zenith angles of 84.3°, 72.5°, and 25.8°).

estimate for a steeply falling muon spectrum but neglects shape uncertainties. The normalization is then constrained using a procedure that tags background muons in data. The spectrum of atmospheric muons from cosmic-ray air showers is modeled by a parameterization of muons from air showers simulated with the `CORSIKA` [156] package assuming the Hillas-Gaisser H4a [146] cosmic-ray flux model and SIBYLL 2.1 [157] hadronic model. A dedicated single muon simulation, called `MUONGUN` [158], is weighted to this flux. A second veto layer inside the original outer veto layer is introduced to construct the data based prior. Events that trigger the outer veto layer, but do not trigger this second inner veto layer, are tagged as muons that pass the inner veto. The muon normalization from simulation is re-scaled from $N_{\texttt{MUONGUN}}$ to $2.1 \cdot N^{\mu}_{\text{tagged}}$ to match the number of tagged muons while accounting for the relative size of the fiducial volumes. Thus, the baseline

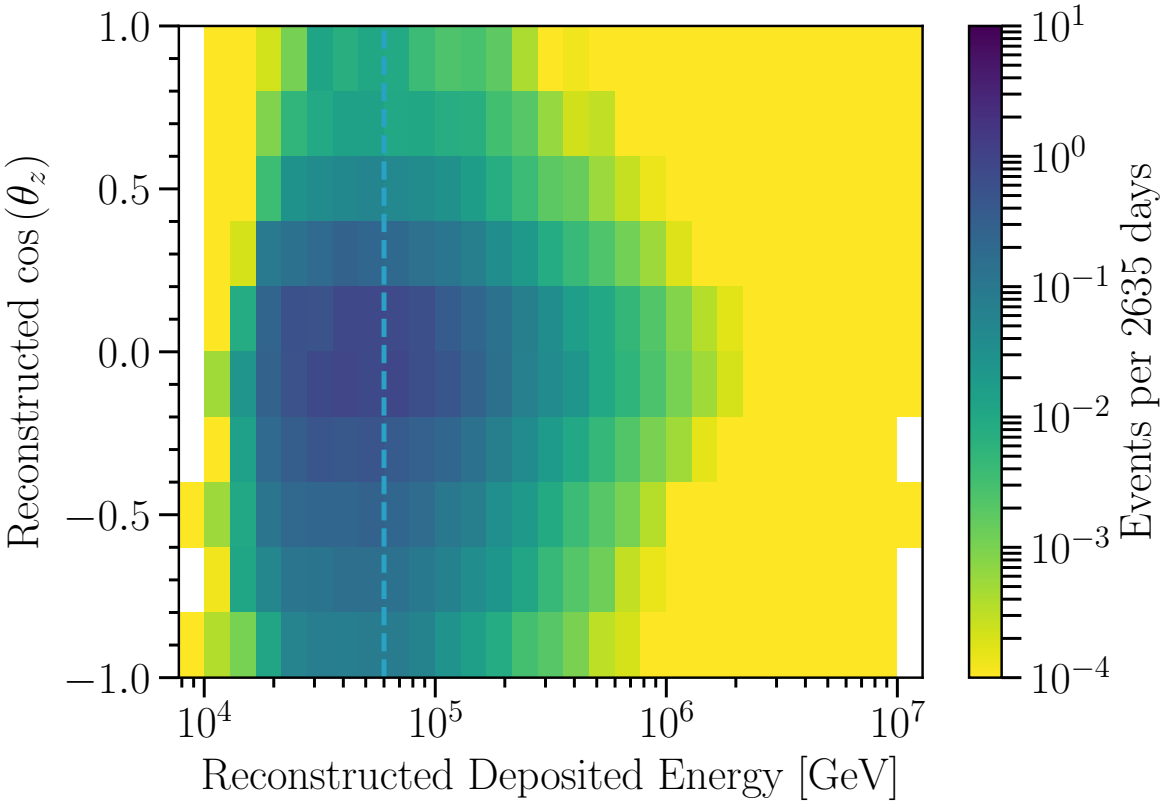

Figure 6.15: ***Expected distribution of atmospheric neutrinos produced by pions and kaons in the sample.*** Distribution of neutrinos that pass the veto as a function of the deposited energy and the cosine of the zenith angle assuming nominal values for the nuisance parameters. The dashed line at 60 TeV marks the low energy cut of the analysis. Suppression in the down-going region is due to the veto, while suppression in the up-going region is due to absorption of neutrinos in the Earth.

expected muon flux is given by

$$\frac{d^3\Phi}{dE_\mu d\theta_{z,\mu} dd_\mu} = \frac{d^3\Phi_{\texttt{GaisserH4a}}}{dE_\mu d\theta_{z,\mu} dd_\mu}(E_\mu, \theta_{z,\mu}, d_\mu) \cdot \frac{2.1 \cdot N^\mu_{\text{tagged}}}{N_{\texttt{MUONGUN}}}, \tag{6.20}$$

where $\Phi_{\texttt{GaisserH4a}}$ is the aforementioned parameterization; and E_μ, $\theta_{z,\mu}$, and d_μ are the muon energy, zenith, and depth at injection respectively. In Table 6.1 we list the number of tagged muons observed per year; in total 17 muons were observed. The expected distribution of passing atmospheric muon events is shown in Fig. 6.19 as a function of the deposited energy and reconstructed cosine of the zenith angle. The prior on the atmospheric muon rate is chosen to be Gaussian with a 50 % standard deviation; this encompasses the statistical uncertainty of our muon background measurement.

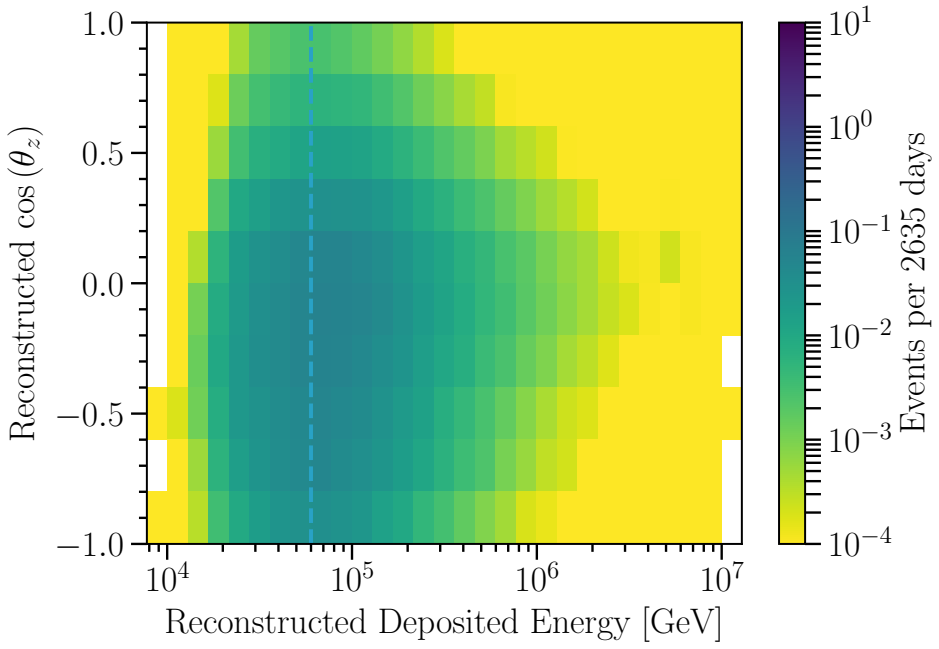

Figure 6.16: ***Expected distribution of atmospheric neutrinos produced by charmed hadrons in the sample.*** Distribution of neutrinos that pass the veto as a function of the deposited energy and the cosine of the zenith angle assuming nominal nuisance parameters and the BERSS flux calculation for neutrinos from charmed hadrons [117]. The dashed line at 60 TeV marks the low energy cut of the analysis. Suppression in the down-going region is due to the veto, while suppression in the up-going region is due to absorption of neutrinos in the Earth.

Season	N^{μ}_{tagged}
2010	2
2011	1
2012	1
2013	1
2014	2
2015	6
2016	2
2017	2
Total	17

Table 6.1: ***Number of tagged muons per season.*** The table shows the number of tagged muons used to construct the muon normalization prior. The first season, 2010, used a partial IceCube configuration with 79 strings, the rest of the seasons took data with the full configuration of 86 strings. The larger number of tagged muons in the 2015 season is believed to be a statistical fluctuation. The last season, 2017, represents only a partial year of data taking in this paper as the 2017 data processing was not yet completed at the time of this analysis.

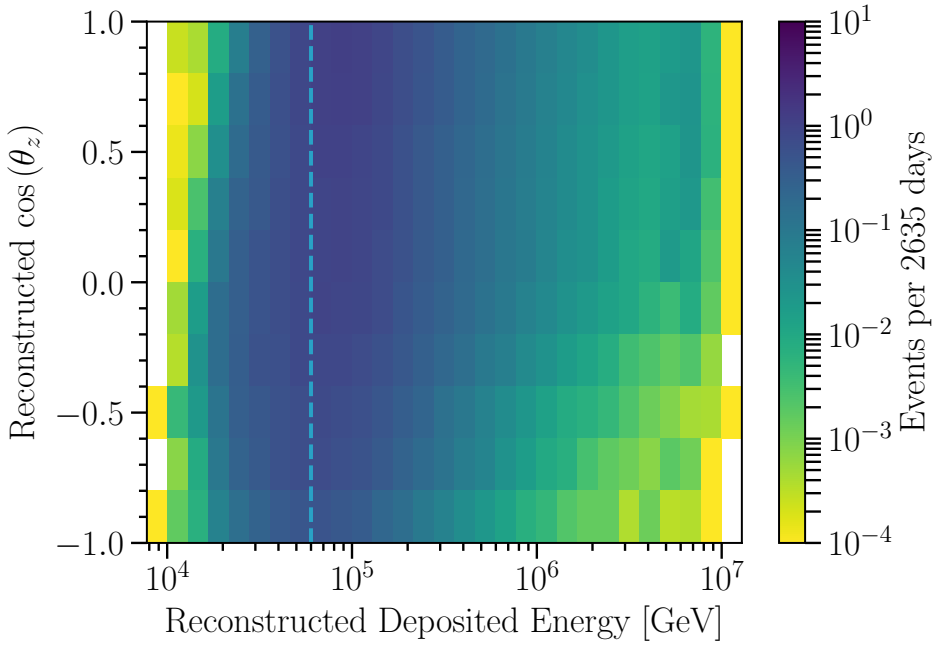

Figure 6.17: ***Expected distribution of astrophysical neutrinos from the best-fit spectrum.*** Distribution of neutrinos that pass the veto as a function of the deposited energy and the cosine of the zenith angle assuming the best-fit parameters of an astrophysical power-law spectrum. The dashed line at 60 TeV marks the low energy cut of the analysis. Suppression in the up-going region is due to absorption of neutrinos in the Earth, while there is no suppression in the down-going region in the absence of accompanying muons.

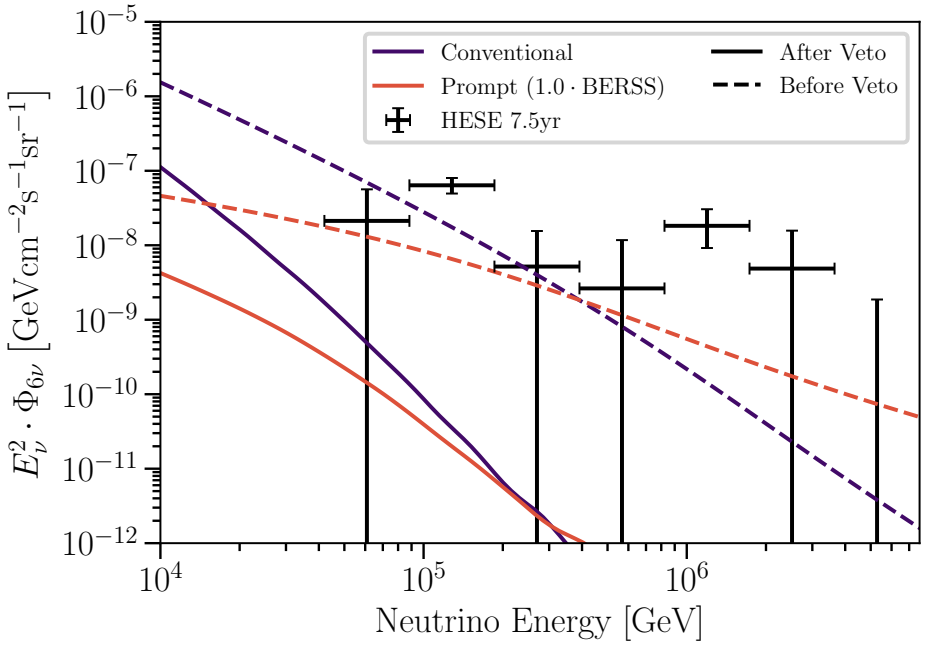

Figure 6.18: ***All-sky astrophysical neutrino flux compared to down-going atmospheric neutrino fluxes before and after the veto.*** The atmospheric neutrino fluxes considered in this analysis are shown as dashed lines. The solid lines show the product of the atmospheric flux with the passing fraction averaged over depth at a zenith angle of 0°. The frequentist segmented power-law fit of the all-sky astrophysical flux assuming isotropy, as described in Section 8.1.1, is shown in black. This comparison demonstrates the effect of the veto in the down-going region, where it is strongest. The suppression of the atmospheric flux becomes weaker towards the horizon and is not present in the up-going region. The dashed lines labeled "before-veto" are equivalent to the up-going atmospheric fluxes, with or without the veto, neglecting Earth absorption effects.

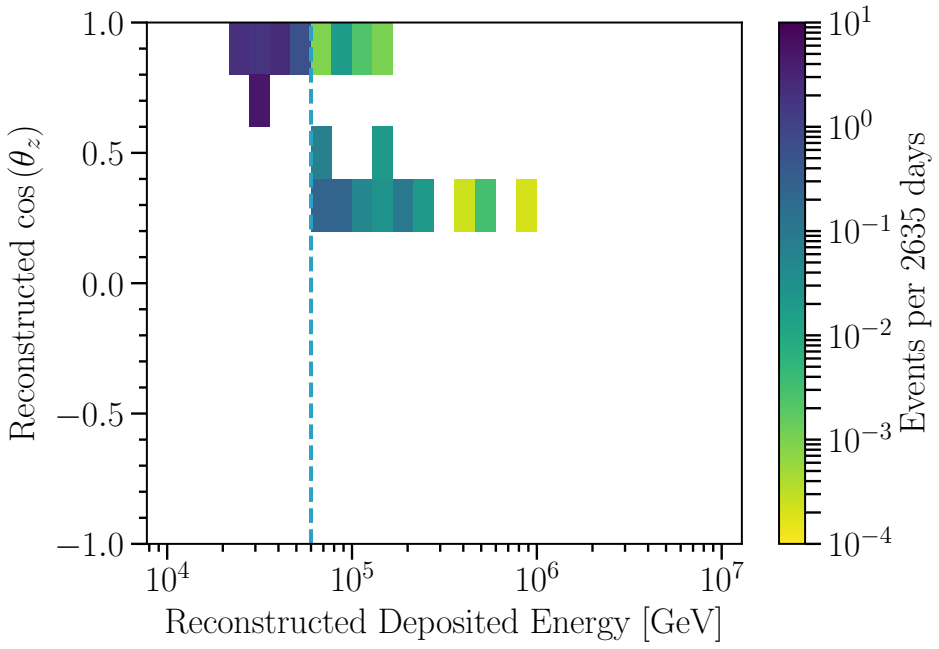

Figure 6.19: ***Expected distribution of atmospheric muons in the sample.*** Distribution of muons that pass the veto as calculated with `MUONGUN` as a function of the deposited energy and the cosine of the zenith angle. The normalization is set to match the data-driven sub-detector study. The dashed line at 60 TeV marks the low energy cut of the analysis.

Chapter 7

Statistics

7.1 Statistical inference

The analysis constructed here sorts events into bins of observable quantities, both to categorize events and so that we have a clear way to perform statistical inferences. Neglecting time variations in astrophysical processes, the binned IceCube observations can be thought of as originating from a multinomial process. There are a finite number of particles that could contribute to IceCube observations in any time interval and the rate of these particles is approximately constant under our assumptions. Each bin is a category in the multinomial process, with one additional category for events not observed. However, the neutrino interaction cross section is extremely small, which means there is a great disparity between the total number of possible events and the expected number of events. For this reason we can approximate the distribution of observed events in any bin as a Poisson process.

Underlying the expected number of events in any bin is a physical model of the particle flux and the detector efficiencies. Detector efficiencies are modeled by simulation of particle interactions and the detector response. This work does not discuss these efficiencies at length, but details of the event selection after detection are given in Chapter 4. The physical particle flux is modeled as several components: atmospheric neutrinos and atmospheric muons which are discussed in Chapter 6, and astrophysical neutrinos which are discussed in Chapter 8.

The model parameters described in this section fall into two categories: parameters of interest ($\vec{\theta}$) and nuisance parameters ($\vec{\eta}$). The former depend on the analysis and the latter include parameters that modify the systematic effects discussed in Section 3.2, as well as physics parameters not being examined. The physics model parameters of interest often refer solely to the astrophysical model parameters, and are discussed in greater detail in Section 8.1. In the case of a single power-law astrophysical flux hypothesis

Parameter	Prior (constraint)	Range	Description
Astrophysical neutrino flux:			
$\Phi_{\texttt{astro}}$	-	$[0, \infty)$	Normalization scale
$\gamma_{\texttt{astro}}$	-	$(-\infty, \infty)$	Spectral index
Atmospheric neutrino flux:			
$\Phi_{\texttt{conv}}$	1.0 ± 0.4	$[0, \infty)$	Conventional normalization scale
$\Phi_{\texttt{prompt}}$	-	$[0, \infty)$	Prompt normalization scale
$R_{K/\pi}$	1.0 ± 0.1	$[0, \infty)$	Kaon-Pion ratio correction
$2\nu/(\nu + \bar{\nu})_{\texttt{atmo}}$	1.0 ± 0.1	$[0, 2]$	Neutrino-anti-neutrino ratio correction
Cosmic ray flux:			
$\Delta\gamma_{\texttt{CR}}$	0.0 ± 0.05	$(-\infty, \infty)$	Cosmic-ray spectral index modification
Φ_{μ}	1.0 ± 0.5	$[0, \infty)$	Muon normalization scale
Detector:			
$\epsilon_{\texttt{DOM}}$	0.99 ± 0.1	$[0.80, 1.25]$	Absolute energy scale
$\epsilon_{\texttt{head-on}}$	0.0 ± 0.5	$[-3.82, 2.18]$	DOM angular response
$a_{\texttt{s}}$	1.0 ± 0.2	$[0.0, 2.0]$	Ice anisotropy scale

Table 7.1: ***Analysis model parameters for the single power-law astrophysical model.*** Prior probabilities (constraints) for analysis parameters used in Bayesian (frequentist) analyses respectively. Priors (constraints) on the parameter are either uniform or Gaussian. Where applicable, the mean, standard deviation, and bounds are given.

these are the astrophysical neutrino flux normalization (Φ_{astro}) and the spectral index of the power-law flux (γ_{astro}). Different parameters of interest are given for other generic astrophysical models in Section 8.1.1 and source-specific models in Section 8.1.3. In the case of searches for new physics, more terms are incorporated into the model parameters; *e.g.* for the dark matter decay search: the dark matter mass and its lifetime. In all cases the systematic treatment of the relevant uncertainties described in Section 3.2 applies.

Results are presented using both frequentist and Bayesian statistical methodologies in this work and associated analyses [159, 160, 161]. These two methodologies provide distinct information [72]. In the frequentist framework, we report the parameters that most likely explain the data. We also report intervals in parameter space constructed such that they contain the true value of the parameter some fraction of the time for repeated experiments. These constructions free us from dependence on priors, but do not make probabilistic statements about the model parameters; see [162, 163, 164, 165, 166] for further discussion on confidence intervals. On the other hand, the Bayesian framework makes statements about the model by invoking Bayes theorem, at the cost of a dependence on prior choice. In this framework, we report the most probable model parameters given the observed data, and the preferred regions of parameter space. Thus, frequentist and Bayesian methods can both be applied and provide complementary information about the model and the data; see [167] for a review and comparison of these methods in the context of neutrino experiments.

Here both of these approaches make use of the likelihood function, which reflects the plausibility of model parameters given observed data and is defined as $\mathcal{L}(\vec{\theta},\vec{\eta}|\texttt{data}) = p(\texttt{data}|\vec{\theta},\vec{\eta})$. Where $p(\texttt{data}|\vec{\theta},\vec{\eta})$ is the probability of the data given the model parameters. External knowledge of the model parameters is also included with the term $\Pi(\vec{\theta},\vec{\eta})$, which is the constraint (prior) on the parameters in the frequentist (Bayesian) interpretation. Details of $\Pi(\vec{\theta},\vec{\eta})$ are given in Table 7.1.

Frequentist results are presented with the best-fit parameters and their errors using the profile likelihood technique. Dropping the explicit notational dependence on data, the profile likelihood function is defined as

$$\tilde{\mathcal{L}}^{\texttt{profile}}(\vec{\theta}) = \max_{\vec{\eta}} \mathcal{L}(\vec{\theta},\vec{\eta}) \cdot \Pi(\vec{\theta},\vec{\eta}), \tag{7.1}$$

where the negative log of the function is minimized in place of maximizing the function. This minimization is performed over continuous nuisance parameters using the L-BFGS-B algorithm [168]. In the frequentist statistical treatment, the constraint $\Pi(\vec{\theta},\vec{\eta})$ is the likelihood of the model parameters given external data. For some parameters no external data is available, and so the constraint $\Pi(\vec{\theta},\vec{\eta})$ is constant with respect to those parameters. Maximizing $\tilde{\mathcal{L}}^{\texttt{profile}}$ over all parameters defines the best-fit point $\hat{\vec{\theta}}$. Frequentist results are then presented assuming Wilks' theorem [169], with the appropriate degrees of freedom; see [170] for a recent summary of the conditions under which this theorem holds. Although analyses presented with the asymptotic approximation described in [169] violate some conditions of Wilks' theorem, the introduction of many nuisance parameters helps to alleviate differences between the real and approximated test-statistic distributions as demonstrated in [170]. As a result, this asymptotic approximation is sufficient for presentation here. The required model parameter test-statistic (TS) is defined as

$$\text{TS}(\vec{\theta}) = -2\log\left(\frac{\tilde{\mathcal{L}}^{\texttt{profile}}(\vec{\theta})}{\tilde{\mathcal{L}}^{\texttt{profile}}(\hat{\vec{\theta}})}\right). \tag{7.2}$$

The product of $\mathcal{L}(\vec{\theta},\vec{\eta})$ and $\Pi(\vec{\theta},\vec{\eta})$ can be normalized to form a probability distribution of the model parameters known as the posterior distribution. The posterior distribution encodes information about the model parameters after being updated by the observed data [171], and is used to present many of the Bayesian results of this work. This probabilistic interpretation allows one to determine the regions of parameter space that have the largest probability of containing the parameter [172]. Integrating the posterior over the nuisance parameters, we obtain the marginal posterior

$$\mathcal{P}(\vec{\theta}) = \frac{\int d\vec{\eta}\ \mathcal{L}(\vec{\theta},\vec{\eta})\Pi(\vec{\theta},\vec{\eta})}{\int d\vec{\theta}d\vec{\eta}\ \mathcal{L}(\vec{\theta},\vec{\eta})\Pi(\vec{\theta},\vec{\eta})}. \tag{7.3}$$

Practically, this is achieved using a Markov Chain Monte Carlo (MCMC) called `EMCEE` [173] to sample the posterior distribution, and examining the distribution of samples in $\vec{\theta}$ [171].

In the case that the model parameter posterior is confined to a compact region we report the highest-posterior-density (HPD) credible region of that parameter, and its maximum *a posteriori* (MAP) estimation.

However, there are some cases where credible regions are ill-defined, or a natural choice of prior is not immediately clear. In these scenarios we report our results using the Bayes factor as a function of the model parameter [174, 175], *i.e.* the ratio of the evidence between the alternative physics model and the null hypothesis,

$$\mathcal{B}_{10} = \frac{\int d\vec{\theta}' d\vec{\eta}' \ \mathcal{L}_1(\vec{\theta}', \vec{\eta}') \cdot \Pi_1(\vec{\theta}', \vec{\eta}')}{\int d\vec{\theta} d\vec{\eta} \ \mathcal{L}_0(\vec{\theta}, \vec{\eta}) \cdot \Pi_0(\vec{\theta}, \vec{\eta})}, \tag{7.4}$$

where the evidence is the average value of the likelihood function with respect to the prior distribution over all model parameters. To compute the model evidence we use the `MultiNest` package [176]. Given a Bayes factor it is customary to assign a qualitative description. For this we use Jeffreys' scale [177].

In order to evaluate the likelihood function, it is necessary to compute the expected number of events in each observable bin given the model parameters. This expectation is obtained through Monte Carlo simulation of the detector. The IceCube Monte Carlo is computationally expensive at high energies, so much so that it is prohibitive to produce background MC such that the statistical fluctuations of the MC are much smaller than the data fluctuations for the atmospheric muon background. In order to avoid making incorrect statements due to the large MC statistical uncertainty in some bins, the analyses described in Section 8.1 use a modified Poisson likelihood function, $\mathcal{L}_{\text{Eff}}$, that incorporates additional uncertainty from the limited MC sample size [66]. This treatment produces similar results to others available in the literature [178, 179, 180], but provides improved coverage properties, is numerically more stable, and is computationally more efficient [66]. The likelihood for this analysis is given by

$$\mathcal{L}(\vec{\theta}, \vec{\eta}) = \prod_j^n \mathcal{L}_{\text{Eff}}(\mu_j(\vec{\theta}, \vec{\eta}), \sigma_j(\vec{\theta}, \vec{\eta}); d_j), \tag{7.5}$$

and the priors or constraints, depending on the context, are given by

$$\Pi(\vec{\theta}, \vec{\eta}) = \left(\prod_r \Pi_r(\theta_r) \right) \cdot \left(\prod_s \Pi_s(\eta_s) \right), \tag{7.6}$$

where j refers to the bin number, r indexes the parameters of interest, and s indexes the nuisance parameters. The variables θ_r and η_r denote the parameters of interest and nuisance parameters, respectively. The arguments of the likelihood μ_j and σ_j are the expected number of events and MC statistical uncertainty of that quantity, respectively, while d_j is the number of observed data events in that bin. The parameters $\vec{\theta}$ and $\vec{\eta}$ have priors or constraints which are represented in Eq. (7.6) by $\Pi_r(\theta_r)$ and $\Pi_s(\eta_s)$, respectively, and are enumerated in Table 7.1. For parameters with improper uniform priors in the Bayesian treatment we apply no external constraint in the frequentist treatment, otherwise the prior and constraint are the same. This results in equivalent functional forms of the product $\mathcal{L}(\vec{\theta}, \vec{\eta})\Pi(\vec{\theta}, \vec{\eta})$ for the frequentist and Bayesian analyses.

The bin widths of the analysis histogram are chosen to be comparable to the detector resolution. There are 840 bins in observable quantities used for the analysis. Events are first separated by their inferred

Morphology	Observable	Bin Edge Condition	Bin Edge	Binning Minimum	Binning Maximum
Cascades	Energy	$\log_{10}(E^{i+1}_{\texttt{bin edge}}/E^{i}_{\texttt{bin edge}}) = 0.111$	$E^{0}_{\texttt{bin edge}} = 60\,\text{TeV}$	$E^{\texttt{min}} = 60\,\text{TeV}$	$E^{\texttt{max}} = 10\,\text{PeV}$
	Zenith	$\cos\theta^{i}_{z,\texttt{bin edge}} - \cos\theta^{i+1}_{z,\texttt{bin edge}} = 0.2$	$\cos\theta^{0}_{z,\texttt{bin edge}} = 0$	$\cos\theta^{\texttt{min}}_{z} = -1$	$\cos\theta^{\texttt{max}}_{z} = 1$
Tracks	Energy	$\log_{10}(E^{i+1}_{\texttt{bin edge}}/E^{i}_{\texttt{bin edge}}) = 0.111$	$E^{0}_{\texttt{bin edge}} = 60\,\text{TeV}$	$E^{\texttt{min}} = 60\,\text{TeV}$	$E^{\texttt{max}} = 10\,\text{PeV}$
	Zenith	$\cos\theta^{i}_{z,\texttt{bin edge}} - \cos\theta^{i+1}_{z,\texttt{bin edge}} = 0.2$	$\cos\theta^{0}_{z,\texttt{bin edge}} = 0$	$\cos\theta^{\texttt{min}}_{z} = -1$	$\cos\theta^{\texttt{max}}_{z} = 1$
Double Cascades	Energy	$\log_{10}(E^{i+1}_{\texttt{bin edge}}/E^{i}_{\texttt{bin edge}}) = 0.111$	$E^{0}_{\texttt{bin edge}} = 60\,\text{TeV}$	$E^{\texttt{min}} = 60\,\text{TeV}$	$E^{\texttt{max}} = 10\,\text{PeV}$
	Length	$\log_{10}(l^{i+1}_{\texttt{bin edge}}/l^{i}_{\texttt{bin edge}}) = 0.1$	$l^{0} = 10\,\text{m}$	$l^{\texttt{min}} = 10\,\text{m}$	$l^{\texttt{max}} = 1000\,\text{m}$

Table 7.2: ***Binning of observable quantities.*** The conditions used to construct the bin edges in each observable for each morphology are presented in this table. An initial bin edge is given, and other bin edges are defined using a relationship between bin edges. The lowest and highest bins are truncated if their bin edges extend beyond the defined boundaries. For each inferred morphology the binning is defined by the Cartesian product of the binning in two observables. This gives 210, 210, and 420 bins for cascades, tracks, and double cascades, respectively, for a total of 840 bins.

morphology, and then binned in two observables. Tracks and cascades are binned in reconstructed energy and reconstructed zenith angle, whereas double cascades are binned in reconstructed energy and the reconstructed separation between cascades. Details of how the bin edges are defined are given in Table 7.2.

7.2 Detailed likelihood and weighting description

The analysis likelihood is given by

$$\mathcal{L}(\vec{\theta},\vec{\eta}) = \left(\prod_{j}^{n} \mathcal{L}_{\text{Eff}}(\mu_j(\vec{\theta},\vec{\eta}),\sigma_j(\vec{\theta},\vec{\eta});d_j)\right)\cdot\left(\prod_{r}\Pi_r(\theta_r)\right)\cdot\left(\prod_{s}\Pi_s(\theta_s)\right), \tag{7.7}$$

where $1 \leq j \leq n$ refers to the bin number, r indexes the parameters of interest, s indexes the nuisance parameters, θ_r are the different parameters of interest, η_s are the different nuisance parameters, Π_r are the priors on the parameters of interest, Π_s are the priors on the nuisance parameters, μ_j refers to the expectation in bin j, σ_j refers to the standard deviation of the expectation in bin j, d_j refers to the number of data events in that bin, $\vec{\theta}$ refers to all the parameters of interest, $\vec{\eta}$ refers to all the nuisance parameters, and $\mathcal{L}_{\text{Eff}}$ is the effective likelihood described in [66]. The arguments of the likelihood are then

$$\mu_j(\vec{\theta},\vec{\eta}) = \sum_i w_i^j(\vec{\theta},\vec{\eta}) \text{ and } \sigma_j(\vec{\theta},\vec{\eta}) = \sqrt{\sum_i \left(w_i^j(\vec{\theta},\vec{\eta})\right)^2}, \tag{7.8}$$

where w_i^j are the weights in bin j. The event weights have contributions from each flux component, which we enumerate as `conv`, `prompt`, `muon`, and `astro` for the conventional atmospheric neutrino, prompt atmospheric neutrino, atmospheric muon, and astrophysical neutrino fluxes, respectively. We also split the weights into their flux dependence, $\alpha_i^{\texttt{component}}$, and systematic corrections, $\beta_i^{\texttt{component}}$. As neutrino events and atmospheric muons are simulated separately, $w_i^{\texttt{muon}}$ is zero if the other weight components are non-zero and vice-versa. The conventional and prompt systematic corrections are applied in the same manner regardless of neutrino

type, whereas the astrophysical systematic corrections are applied on a per-flavor basis. In symbolic notation, this can be written as:

$$w_i = w_i^{\texttt{conv}} + w_i^{\texttt{prompt}} + w_i^{\texttt{muon}} + w_i^{\texttt{astro}}, \tag{7.9}$$

$$w_i^{\texttt{conv}} = \alpha_i^{\texttt{conv}} \beta_i^{\texttt{conv}}, \tag{7.10}$$

$$w_i^{\texttt{prompt}} = \alpha_i^{\texttt{prompt}} \beta_i^{\texttt{prompt}}, \tag{7.11}$$

$$w_i^{\texttt{muon}} = \alpha_i^{\texttt{muon}}, \tag{7.12}$$

$$w_i^{\texttt{astro}} = \alpha_i^{\texttt{astro}} \beta_i^{\texttt{astro},\|p_i\|}, \tag{7.13}$$

where p_i denotes the particle type, and $\|p_i\|$ denotes the particle type irrespective of whether it is an anti-particle or not.

To correct for differences between the simulated event generation probability distributions and those of another hypothesis given by $\vec{\theta}$ and $\vec{\eta}$, we must concern ourselves with the true simulated properties of each simulation-event's primary particle. We denote p_i as the particle type of the event's primary particle, E_i as the primary particle's initial energy, θ_i^z as the primary's zenith angle, and d_i as the primary's depth at the first intersection with a cylinder centered around and containing the detector. The main correction is between the generated distribution of neutrinos and a baseline flux of neutrinos for a particular livetime. This correction necessitates a weighting factor $(\Phi \cdot T)/(N \cdot P)$, where Φ is the differential flux, T the livetime of the sample, N is the number of generated events, and P is the probability density of event generation. Additional corrections can account for deviations from the baseline model, and effects not simulated.

The conventional component considers the neutrino flux from pions (π) and kaons (K) separately. For a neutrino of type p_i, of energy E_i^ν, and with zenith angle $\theta_i^{\nu,z}$, the differential flux of such neutrinos for pions and kaons is $\Phi_{\texttt{HONDA}}^{\pi,p_i}(E_i^\nu, \theta_i^{\nu,z})$ and $\Phi_{\texttt{HONDA}}^{K,p_i}(E_i^\nu, \theta_i^{\nu,z})$, respectively [114]. The overall normalization of the conventional atmospheric neutrino flux is modified via the parameter $\Phi_{\texttt{conv}}$ and the relative normalizations of the components from pions and kaons are controlled by the parameter $R_{K/\pi}$. To modify the slope of the conventional atmospheric neutrino spectrum the parameter $\Delta\gamma_{\texttt{CR}}$ acts as a spectral index correction about a 2020 GeV pivot point. The proportion of atmospheric neutrinos and anti-neutrinos is allowed to vary via the parameter $2\nu/(\nu+\bar{\nu})_{\texttt{atmo}}$ which can range from zero (all anti-neutrinos) to two (all neutrinos), where the nominal value of one gives the relative neutrino to anti-neutrino content specified in the chosen atmospheric models. Finally, as we only simulate single neutrinos and not all the products of the cosmic-ray air showers from which they originate, a correction factor is needed to account for the probability that a neutrino event may be rejected by the presence of an accompanying muon. To account for this the weight is multiplied by the probability that the neutrino is not rejected by an accompanying muon $P_{\texttt{passing}}^{\texttt{conv},p_i}(E_i^\nu, \theta_i^{\nu,z}, D_i^\nu)$. In

symbolic notation, this is given by:

$$\alpha_i^{\texttt{conv}} = \Phi_{\texttt{conv}} \left(\frac{\Phi^{\pi,p_i}_{\texttt{HONDA}}(E_i^\nu, \theta_i^{\nu,z}) T_{\texttt{sample}}}{N_{\texttt{gen}} P^{p_i}_{\texttt{gen}}(E_i^\nu, \theta_i^{\nu,z})} + R_{K/\pi} \frac{\Phi^{K,p_i}_{\texttt{HONDA}}(E_i^\nu, \theta_i^{\nu,z}) T_{\texttt{sample}}}{N_{\texttt{gen}} P^{p_i}_{\texttt{gen}}(E_i^\nu, \theta_i^{\nu,z})} \right) \cdot \left(\frac{E_i^\nu}{2020\,\text{GeV}} \right)^{-\Delta\gamma_{\texttt{CR}}}$$
$$\cdot P^{\texttt{conv},p_i}_{\texttt{passing}}(E_i^\nu, \theta_i^{\nu,z}, D_i^\nu) \cdot \begin{cases} 2\nu/\left(\nu+\bar{\nu}\right)_{\texttt{atmo}} & p_i \texttt{ is } \nu \\ 2 - 2\nu/\left(\nu+\bar{\nu}\right)_{\texttt{atmo}} & p_i \texttt{ is } \bar{\nu} \end{cases} . \tag{7.14}$$

The prompt neutrino component considers the flux of atmospheric neutrinos from charmed hadrons $\Phi^{p_i}_{\texttt{BERSS}}(E_i^\nu, \theta_i^{\nu,z})$ [117]. Similar corrections as those defined for the conventional component are used for the prompt component. In this case the normalization of the prompt neutrino flux is controlled by $\Phi_{\texttt{prompt}}$. The parameters $\Delta\gamma_{\texttt{CR}}$ and $2\nu/\left(\nu+\bar{\nu}\right)_{\texttt{atmo}}$ serve the same purpose in the prompt neutrino component, except that the pivot point is chosen to be 7887 GeV. A term must also be included to account for the probability of the neutrino being rejected due to accompanying muons: $P^{\texttt{prompt},p_i}_{\texttt{passing}}(E_i^\nu, \theta_i^{\nu,z}, D_i^\nu)$. In symbolic notation, this is given by:

$$\alpha_i^{\texttt{prompt}} = \Phi_{\texttt{prompt}} \left(\frac{\Phi^{p_i}_{\texttt{BERSS}}(E_i^\nu, \theta_i^{\nu,z}) T_{\texttt{sample}}}{N_{\texttt{gen}} P^{p_i}_{\texttt{gen}}(E_i^\nu, \theta_i^{\nu,z})} \right) \cdot \left(\frac{E_i^\nu}{7887\,\text{GeV}} \right)^{-\Delta\gamma_{\texttt{CR}}}$$
$$\cdot P^{\texttt{prompt},p_i}_{\texttt{passing}}(E_i^\nu, \theta_i^{\nu,z}, D_i^\nu) \cdot \begin{cases} 2\nu/\left(\nu+\bar{\nu}\right)_{\texttt{atmo}} & p_i \texttt{ is } \nu \\ 2 - 2\nu/\left(\nu+\bar{\nu}\right)_{\texttt{atmo}} & p_i \texttt{ is } \bar{\nu} \end{cases} . \tag{7.15}$$

The flux of atmospheric muons from cosmic-ray air showers is modelled by a parameterization of muons from air showers simulated with the CORSIKA package assuming the Hillas-Gaisser H4a [146] cosmic-ray flux model and SIBYLL 2.1 [157] hadronic model. This parameterized flux is denoted by $\Phi_{\texttt{GaisserH4a}}(E_i^\nu, \theta_i^{\nu,z}, d_i^\nu)$. As for other fluxes a normalization factor Φ_μ is included. The only other correction is to shift the baseline flux to the center of the data derived prior. This is accomplished with the $2.1 \cdot N^\mu_{\texttt{tagged}}/N_{\texttt{MUONGUN}}$ factor. Where 2.1 is the ratio of detection volumes for the full and reduced volume event selections (not accounting for differences in efficiency), $N^\mu_{\texttt{tagged}}$ is the number of tagged muons, and $N_{\texttt{MUONGUN}}$ is the number of expected events in the baseline atmospheric muon model before re-scaling. Namely,

$$\alpha_i^{\texttt{muon}} = \Phi_\mu \frac{\Phi_{\texttt{GaisserH4a}}(E_i^\mu, \theta_i^{\mu,z}, d_i^\mu) T_{\texttt{sample}}}{N_{\texttt{gen}} P^{p_i}_{\texttt{gen}}(E_i^\mu, \theta_i^{\mu,z}, d_i^\mu)} \cdot \frac{2.1 \cdot N^\mu_{\texttt{tagged}}}{N_{\texttt{MUONGUN}}}. \tag{7.16}$$

The astrophysical component is modeled with a single power law as a baseline, although this flux can be replaced with other models as has been done in Section 8.1. The factors accounting for the generation are the same as for the other neutrino fluxes, and a normalization factor $\Phi_{\texttt{astro}}$ is included. The baseline normalization is chosen to be $10^{-18}[\text{GeV}^{-1}\text{cm}^{-2}\text{s}^{-1}\text{sr}^{-1}]$ at 100 TeV. The factor of 2π stems from the uniform azimuthal distribution of the astrophysical flux, and $\gamma_{\texttt{astro}}$ governs the index of the spectrum. In symbolic notation, this is given by:

$$\alpha_i^{\texttt{astro}} = \Phi_{\texttt{astro}} \left(\frac{10^{-18}[\text{GeV}^{-1}\text{cm}^{-2}\text{s}^{-1}\text{sr}^{-1}] 2\pi T_{\texttt{sample}}}{N_{\texttt{gen}} P^{p_i}_{\texttt{gen}}(E_i^\nu, \theta_i^{\nu,z})} \right) \cdot \left(\frac{E_i^\nu}{100\,\text{TeV}} \right)^{-\gamma_{\texttt{astro}}}. \tag{7.17}$$

Similarly to true simulated properties, the reconstructed event properties are also needed for the purpose of binning and systematic corrections. We denote m_i^R as the inferred morphology of event i, E_i^R as the reconstructed deposited energy of the event, $\theta_i^{R,z}$ as the reconstructed event zenith angle, and l_i^R as the reconstructed distance between energy depositions in the double cascade reconstruction. The effects of detector systematic parameters $\epsilon_{\texttt{DOM}}$, $\epsilon_{\texttt{head-on}}$, and a_s are assumed to be independent, and are each accounted for by corrections to the expectation stored in b-splines. For each systematic the corrections for tracks and cascades are applied using the same combination of observables, while double cascades differ. These are organized as:

$$\beta_i^x = \begin{cases} \beta_i^{\texttt{dc},x} & m_i^R = \texttt{double cascade} \\ \beta_i^{\texttt{t/c},x} & m_i^R = \texttt{track or } m_i^R = \texttt{cascade} \end{cases} \tag{7.18}$$

Parameterized expectations for each systematic s are denoted by f_s, and corrections are ratios of f_s evaluated at a specific systematic parameter value to f_s evaluated at the nominal systematic parameter value denoted by s_0. These correction factors are explicitly given by:

$$\beta_i^{\texttt{dc},x} = \left(\frac{f^{x,\texttt{dc}}_{\epsilon_{\texttt{DOM}}}(E_i^R,\theta_i^{R,z},\epsilon_{\texttt{DOM}})}{f^{x,\texttt{dc}}_{\epsilon_{\texttt{DOM}}}(E_i^R,\theta_i^{R,z},\epsilon_{\texttt{DOM},0})}\right)\cdot\left(\frac{f^{x,dc}_{\epsilon_{\texttt{head-on}}}(E_i^R,l_i^R,\epsilon_{\texttt{head-on}})}{f^{x,dc}_{\epsilon_{\texttt{head-on}}}(E_i^R,l_i^R,\epsilon_{\texttt{head-on},0})}\right)\cdot\left(\frac{f^{x,dc}_{a_s}(l_i^R,a_s)}{f^{x,dc}_{a_s}(l_i^R,a_{s,0})}\right), \tag{7.19}$$

$$\beta_i^{\texttt{t/c},x} = \left(\frac{f^{x,m_i^R}_{\epsilon_{\texttt{DOM}}}(E_i^R,\theta_i^{R,z},\epsilon_{\texttt{DOM}})}{f^{x,m_i^R}_{\epsilon_{\texttt{DOM}}}(E_i^R,\theta_i^{R,z},\epsilon_{\texttt{DOM},0})}\right)\cdot\left(\frac{f^{x,m_i^R}_{\epsilon_{\texttt{head-on}}}(E_i^R,\theta_i^{R,z},\epsilon_{\texttt{head-on}})}{f^{x,m_i^R}_{\epsilon_{\texttt{head-on}}}(E_i^R,\theta_i^{R,z},\epsilon_{\texttt{head-on},0})}\right)\cdot\left(\frac{f^{x,dc}_{a_s}(E_i^R,a_s)}{f^{x,dc}_{a_s}(E_i^R,a_{s,0})}\right). \tag{7.20}$$

7.3 Dealing with limited simulation samples

The contents of this section is reproduced here with minor modifications from a collaborative work with Carlos A. Argüelles, and Tianlu Yuan [66]. This work was originally motivated by the desire to account for simulation statistical uncertainties that originate from the muon background component, and our failure to find a treatment in the literature with the desired properties. The simulation statistical uncertainties of the muon component are comparable to the statistical uncertainties of data in the same bins, but little to no information is available outside of the bins that contain simulated muons. The treatment described in this section is able to account for the inherent uncertainty in these bins, but neglects potential contributions from muon in bins not covered by the simulation template. We would hope to improve upon the treatment of these empty bins in future work. Ultimately these differences do not change the inferences made in this analysis because of other constraints present in the data, but this treatment is applicable to a wide range of analyses and we plan to use it in the future.

The use of Monte Carlo (MC) techniques to calculate nontrivial theoretical quantities and expectations in complex experimental settings is common practice in particle physics. An MC event is a single representation of what can be detected in data and is typically generated from a single realization of the underlying physics

parameters, $\vec{\theta}_g$. These events are often binned in some observable space and compared with the data. Since the generation process is stochastic, a particular $\vec{\theta}_g$ used for generating the MC can lead to different outputs. This stochasticity introduces an uncertainty in the MC distributions. Furthermore, as production of large MC is often time-consuming, reweighting is used to move from one hypothesis to another. In reweighting, each MC event is assigned a new weight, $w(\vec{\theta})$, that accounts for the difference between the generation parameters $\vec{\theta}_g$ and the hypothesis parameters $\vec{\theta}$ [181]. It follows that MC uncertainties will be hypothesis dependent; thus, to do hypothesis testing, it is important to account for them. This is especially important for small-signal searches, performed in the small sample limit, where a modified-χ^2 may not be suitable [182]. A Poisson likelihood is a more appropriate statistical description of event counts [183], but in that case a proper treatment of MC statistical uncertainties is less straightforward. Solutions to this problem have been discussed in the literature in the context of frequentist statistics by adding nuisance parameters [184, 185, 178], as well as detailed probabilistic treatment of MC weights [179]. However, [184, 178, 179] add additional time complexity, and [185] does not provide a full exposition on how to incorporate weighted MC. This section presents a new treatment that is valid in the large and small limit of the data sample size, suited for frequentist and Bayesian analyses, based on the Poisson likelihood. Our likelihood accounts for statistical uncertainties due to MC, allows for arbitrary event-by-event reweighting, and is computationally efficient. A test statistic based on the proposed likelihood is found to follow a distribution closer to the asymptotic form expected from Wilks' theorem. An implementation of the likelihood described in this work can be found in [186].

Section 7.3.1 briefly reviews two common treatments available in the literature to account for MC statistical uncertainty. Section 7.3.2 defines and discusses a new likelihood. Section 7.3.3 studies the performance of the likelihood through an example and compares it to other likelihoods in the literature. Section 7.3.4 provide our conclusions. A summary of the likelihoods discussed, including the main result of this work, is given in Appendix 7.3.5.

7.3.1 The Poisson likelihood and previous work

In order to compare MC with data, events are often binned into distributions across a set of observables. For simplicity, we focus on a single bin. In the absence of cross-bin-correlated systematic uncertainties the generalization to multiple bins is simply a product over the likelihood in all bins. This is assumed for the remainder of the discussion. It is well known that the count of independent, rare natural processes can be described by the Poisson likelihood, given by

$$\mathcal{L}(\vec{\theta}|k) = \text{Poisson}(k;\lambda(\vec{\theta})) = \frac{\lambda(\vec{\theta})^k e^{-\lambda(\vec{\theta})}}{k!}, \tag{7.21}$$

where $\lambda(\vec{\theta})$ is the expected bin count for a hypothesis and k is the number of observed data events. Equation (7.21) requires exact knowledge of the expected bin count, $\lambda(\vec{\theta})$. In the case of complex experiments it is often not possible to obtain $\lambda(\vec{\theta})$ exactly and MC techniques are used to estimate the expected distributions. For weighted MC, often a direct substitution of $\lambda(\vec{\theta})$ by $\sum_i w_i(\vec{\theta})$ is used, where w_i are the weights of each of the MC events in the bin. Then Eq. (7.21) can be approximated as

$$\mathcal{L}_{\text{AdHoc}}(\vec{\theta}|k) = \frac{\left(\sum_i w_i(\vec{\theta})\right)^k e^{-(\sum_i w_i(\vec{\theta}))}}{k!}. \tag{7.22}$$

This ad hoc treatment assumes that the MC estimate of the expected bin counts exactly matches the true expectation rate of the model, neglecting the stochastic nature of MC. In the case of large MC, Eq. (7.22) converges to Eq. (7.21) for the hypothesis given by $\vec{\theta}$.

7.3.1.1 The Barlow-Beeston likelihood

To treat MC statistical uncertainties in the small sample limit, a modification of the Poisson likelihood was introduced in [184], which is briefly covered below. First, note that the expectation in a single bin is given by contributions from different physical processes, which we index by j. Then, the number of expected events can be written as

$$\lambda(\vec{\theta}) = \sum_{j=1}^{s} \bar{n}_j(\vec{\theta}), \tag{7.23}$$

where $\bar{n}_j$ is the expected number of MC events from process j that fall in the bin and s is the total number of relevant processes. Substituting Eq. (7.23) into Eq. (7.21) gives the Poisson likelihood for observing k data events. For stochastic models, $\bar{n}_j$ is unknown. Instead, the MC outcome can be modeled as having drawn n_j events from a random process that simulates the physical process. When MC generation is expensive, we can approximate n_j as being drawn from a Poisson process with mean $\bar{n}_j$ [1]. Profiling on the true number of MC events per process in the bin results in the Barlow-Beeston (BB) likelihood, given by [184]

$$\mathcal{L}_{\text{BB}}(\vec{\theta}|k) = \max_{\{\bar{n}_j\}} \frac{\lambda(\vec{\theta})^k e^{-\lambda(\vec{\theta})}}{k!} \prod_{j=1}^{s} \frac{\bar{n}_j^{n_j} e^{-\bar{n}_j}}{n_j!}, \tag{7.24}$$

where $\lambda(\vec{\theta})$ is given by Eq. (7.23), n_j and $\bar{n}_j$ are the estimated and true MC counts in the bin respectively, and $\{\bar{n}_j\}_{j=1}^{s}$ denotes the s nuisance parameters we have profiled over.

[1] The MC generation is a binomial process where we generate a fixed number of events for each process, N_j, and accept them into the bin of interest with probability $\beta_j(\vec{\theta})$, such that $\bar{n}_j(\vec{\theta}) = \beta_j(\vec{\theta})N_j$. In the limit of both rare processes ($\beta_j \ll 1$) and many generated events ($N_j \gg 1$), the total number of observed events can be approximated as Poisson distributed with mean $\lambda(\vec{\theta}) = \sum_j \beta_j(\vec{\theta})N_j = \sum_j \bar{n}_j(\vec{\theta})$.

In the above formalism MC has been produced at the natural rate, but this is not the case for weighted MC. The prescription is given by replacing Eq. (7.23) with

$$\lambda(\vec{\theta}) = \sum_{j=1}^{s} \eta_j(\vec{\theta})\bar{n}_j, \tag{7.25}$$

where $\eta_j(\vec{\theta})$ is a scale factor for process j that accounts for the differences in the MC generation and the target hypothesis of interest. In this case, the likelihood definition is still given by Eq. (7.24); an explicit formula for $s = 1$ is given in the appendix. However, for arbitrary weight distributions per physical process $\mathcal{L}_{\text{BB}}$ may not be appropriate as it neglects the variance from a sum of weights [184]. It remains valid only in the case where the distribution of weights for each process is narrow.

A computational issue with this approach arises for $s \geq 1$. Although the likelihood may be expressed simply as the maximum with respect to the per-bin nuisance parameters, this maximization becomes computationally cumbersome for large s. The primary result of [184] is a more efficient approach that reduces the global maximization into s root finding problems. However, when this numerical root-finding approach introduces discrete jumps in the numerical approximation of the likelihood. These discrete jumps in turn cause problems for the global maximization of the likelihood with respect to physical parameters, and can cause this procedure to fail.

7.3.1.2 Uncertainties in the large-sample limit

In the large-sample regime, the Gaussian distribution is an appropriate description of the observed data. In this limit, the use of Pearson's χ^2 as a test-statistic [187] is common practice. For a single analysis bin, Pearson's χ^2 is defined as

$$\chi^2(\vec{\theta}) = \frac{(k - \lambda(\vec{\theta}))^2}{\lambda(\vec{\theta})}, \tag{7.26}$$

where we continue to use the approximation $\lambda(\vec{\theta}) = \sum_i w_i(\vec{\theta})$ and w_i are the weights of each of the MC events. The form of Pearson's χ^2 arises from the fact that the Gaussian distribution of k is the large-sample limit of a Poisson distribution for which the expected statistical variance of the observation is given by $\lambda(\vec{\theta})$. Systematic uncertainties, under the assumption that they follow a Gaussian distribution and are independent between bins, can be included as

$$\chi^2(\vec{\theta}) = \frac{(k - \lambda(\vec{\theta}))^2}{\lambda(\vec{\theta}) + \sigma^2_{\text{syst.}}}. \tag{7.27}$$

However, this method of incorporating systematic uncertainties tends to overestimate them in shape-only analyses; see [188] for a recent discussion in the context of reactor neutrino anomalies. Similarly, one can include uncertainties to account for statistical fluctuations of the MC in the test-statistic. In doing so, the

Gaussian behavior is implicit and the modified χ^2 reads

$$\chi^2_{\text{mod}}(\vec{\theta}) = \frac{(k - \lambda(\vec{\theta}))^2}{\lambda(\vec{\theta}) + \sigma^2_{\text{syst.}} + \sigma^2_{\text{mc}}}, \tag{7.28}$$

where σ^2_{mc} is the MC statistical uncertainty in the bin given by

$$\sigma^2_{\text{mc}}(\vec{\theta}) \equiv \sum_{i=1}^{m} w_i(\vec{\theta})^2. \tag{7.29}$$

Note that this test-statistic definition is not appropriate in the small-sample regime, as the data is no longer well described by a Gaussian distribution. If one uses a χ^2 test-statistic in the small-sample regime, one ought to calculate the test-statistic distribution properly to achieve appropriate coverage [189]. Calculation of this test statistic distribution presents some challenges which are discussed further in Section 7.4.

7.3.2 Generalization of the Poisson likelihood

Ideally we would like to obtain the expected event count for any hypothesis, $\lambda(\vec{\theta})$, however we are considering problems where this relationship is not known and λ is instead estimated by MC. The key difference here is that instead of using exact knowledge of λ we want to perform Bayesian inference to obtain $\mathcal{P}(\lambda|\vec{\theta})$ using the MC available. Assuming the weights are functions of $\vec{\theta}$, we have

$$\mathcal{L}_{\text{General}}(\vec{\theta}|k) = \int_0^\infty \frac{\lambda^k e^{-\lambda}}{k!} \mathcal{P}\left(\lambda|\vec{w}(\vec{\theta})\right)\, d\lambda, \tag{7.30}$$

where the distribution of λ, $\mathcal{P}\left(\lambda|\vec{w}(\vec{\theta})\right)$, is inferred from the MC. The likelihood, $\mathcal{L}_{\text{AdHoc}}$, in Eq. (7.22) is recovered when $\mathcal{P}(\lambda|\vec{w}(\vec{\theta})) = \delta\left(\lambda - \sum_i w_i(\vec{\theta})\right)$, but clearly this is an unrealistic assumption as it presumes perfect knowledge of the parameter $\lambda(\vec{\theta})$ from a finite number of realizations. Instead, it is more appropriate to construct $\mathcal{P}(\lambda|\vec{w}(\vec{\theta}))$ based on the MC realization. This is given by

$$\mathcal{P}\left(\lambda|\vec{w}(\vec{\theta})\right) = \frac{\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))\mathcal{P}(\lambda)}{\int_0^\infty \mathcal{L}(\lambda'|\vec{w}(\vec{\theta}))\mathcal{P}(\lambda')\, d\lambda'}, \tag{7.31}$$

where $\mathcal{P}(\lambda)$ is a prior on λ that must be chosen appropriately and $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ is the likelihood of λ given $\vec{w}(\vec{\theta})$. This is similar to [184, 185], but instead of fitting λ as a nuisance parameter as in $\mathcal{L}_{\text{BB}}$ in Eq. (7.24), we marginalize over it in Eq. (7.30) as informed by the MC weights. When $\mathcal{L}_{\text{General}}$ is used under a frequentist approach, the marginalization over λ implies a hybrid Bayesian-frequentist construction, similar to the treatment of nuisance parameters described in [190] and employed in [191, 192].

This section is organized as follows. First $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ is derived assuming identical weights in Sec. 7.3.2.1, and is then extended to arbitrary weights in Sec. 7.3.2.2. With this in hand, we can calculate an analytic expression for Eq. (7.30) using Eq. (7.31) under a uniform $\mathcal{P}(\lambda)$ prior in Sec. 7.3.2.3. Section 7.3.2.4 briefly discusses a family of distributions as possible alternative priors. Section 7.3.2.5 shows that the effective

likelihood converges to Eq. (7.22) in the limit of large MC size. Finally, Sec. 7.3.2.6 provides some intuition on the behavior of the generalized likelihood. Equation (7.45), along with the definitions of μ and σ^2 given in Eq. (7.32), constitutes the primary result of this work.

7.3.2.1 Derivation of $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ for identical weights

In this section $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ is derived for identical weights. We will show that $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ can be written in terms of two quantities

$$\mu \equiv \sum_{i=1}^{m} w_i \text{ and } \sigma^2 \equiv \sum_{i=1}^{m} w_i^2 \tag{7.32}$$

for a bin with m MC events.

For identical weights, $w \equiv w_i \ \forall i$, the following equalities hold:

$$\mu = wm,\ \sigma^2 = w^2 m,\ w = \sigma^2/\mu, \text{ and } \ m = \mu^2/\sigma^2. \tag{7.33}$$

Assume that m is the outcome of sampling a Poisson-distributed random variable M with probability mass function

$$\text{Poisson}(M = m; \bar{m}) = \frac{e^{-\bar{m}}\bar{m}^m}{m!}, \tag{7.34}$$

where $\bar{m}$ is the mean of the distribution. Further, assume that the expected number of data events $\lambda = w\bar{m}$ so that $\bar{m} = \lambda/w$. Substituting back into Eq. (7.34), we can interpret $\text{Poisson}(M = m; \bar{m})$ as a likelihood function of λ

$$\mathcal{L}(\lambda|\vec{w}(\vec{\theta})) = \mathcal{L}(\lambda|\mu,\sigma) = \frac{e^{-\lambda\mu/\sigma^2}\left(\lambda\mu/\sigma^2\right)^{\mu^2/\sigma^2}}{(\mu^2/\sigma^2)!}, \tag{7.35}$$

as μ and σ fully specify $\vec{w}(\vec{\theta})$ for identical weights.

7.3.2.2 Extension to arbitrary weights

The derivation above assumed identical weights. For arbitrary weights, μ is an outcome sampled from a compound Poisson distribution (CPD), which can be approximated by a scaled Poisson distribution (SPD) by matching the first and second moments of the two distributions [193]. In order to make the connection, first rewrite μ and σ^2 as

$$\mu = w_{\text{Eff}} m_{\text{Eff}} \text{ and } \sigma^2 = w_{\text{Eff}}^2 m_{\text{Eff}}, \tag{7.36}$$

where m_{Eff} is the effective number of MC events and w_{Eff} the effective weight. From Eq. (7.33) these are given by: $m_{\text{Eff}} = \mu^2/\sigma^2$ and $w_{\text{Eff}} = \sigma^2/\mu$. Next, assume $\bar{m} = \lambda/w_{\text{Eff}}$ and

$$\mathcal{L}(\bar{m}|m_{\text{Eff}}) = \frac{e^{-\bar{m}}\bar{m}^{m_{\text{Eff}}}}{\Gamma(m_{\text{Eff}}+1)}, \tag{7.37}$$

where λ again is the expected number of events in data. Equation (7.37) can be written as a likelihood function of λ,

$$\mathcal{L}(\lambda|\vec{w}(\vec{\theta})) = \mathcal{L}(\lambda|\mu,\sigma) = \frac{e^{-\lambda\mu/\sigma^2}\left(\lambda\mu/\sigma^2\right)^{\mu^2/\sigma^2}}{\Gamma(\mu^2/\sigma^2+1)}, \tag{7.38}$$

which is identical to Eq. (7.35) except the denominator is now a gamma function instead of a factorial. However, since the denominator does not depend on λ it cancels out in Eq. (7.31).

To understand this approximation, note that the maximum likelihood in Eq. (7.37) occurs when $\bar{m} = m_{\text{Eff}}$. The first and second moments of the SPD random variable $w_{\text{Eff}}M$, where $M \sim \text{Poisson}(m_{\text{Eff}})$, are given by

$$\begin{aligned} \text{E}[w_{\text{Eff}}M] &= w_{\text{Eff}}m_{\text{Eff}} \\ &= \mu, \end{aligned} \tag{7.39}$$

and

$$\begin{aligned} \text{Var}[w_{\text{Eff}}M] &= w_{\text{Eff}}^2 m_{\text{Eff}} \\ &= \sigma^2. \end{aligned} \tag{7.40}$$

This shows that the SPD, under the maximum likelihood solution for the given MC realization, has first and second moments that match the sample mean, μ, and variance, σ^2, respectively. These are equal to the first and second moments of the CPD as described in [193]. By assuming that μ is drawn from a SPD, we can treat μ and σ as outcomes that fix the likelihood function of the underlying scaled expectation λ, analogous to the case of identical weights. Because both the first and second moments are matched, this approximation accounts for the variance of the CPD unlike $\mathcal{L}_{\text{BB}}$, which only accounts for the mean. Thus, while $\mathcal{L}_{\text{BB}}$ is valid only for the case of narrow weight distributions, this approximation remains valid for broader distributions.

7.3.2.3 The effective likelihood

Now that we have an expression for $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ from the MC, we can proceed to compute Eq. (7.30) under a uniform $\mathcal{P}(\lambda)$. To simplify the notation, let

$$\alpha = \frac{\mu^2}{\sigma^2} + 1 \text{ and } \beta = \frac{\mu}{\sigma^2}. \tag{7.41}$$

Then, assuming a uniform $\mathcal{P}(\lambda)$ and substituting Eq. (7.38) for $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ in Eq. (7.31) we obtain

$$\begin{aligned} \mathcal{P}(\lambda|\vec{w}(\vec{\theta})) &= \beta\frac{e^{-\lambda\beta}(\lambda\beta)^{\alpha-1}}{\Gamma(\alpha)} \\ &= \frac{e^{-\lambda\beta}\lambda^{\alpha-1}\beta^{\alpha}}{\Gamma(\alpha)} \\ &= \mathcal{G}(\lambda;\alpha,\beta), \end{aligned} \tag{7.42}$$

where Γ is the gamma function and $\mathcal{G}$ the gamma distribution with shape parameter α and inverse-scale parameter β. Note that in going from Eq. (7.38) to Eq. (7.42) μ and σ^2 go from random variates for a particular λ to parameters that govern the probability density of λ. With this choice of $\mathcal{L}(\lambda|\vec{w}(\vec{\theta}))$ and $\mathcal{P}(\lambda)$, we can rewrite $\mathcal{L}_{\text{General}}$ from Eq. (7.30) as

$$\mathcal{L}_{\text{Eff}}(\vec{\theta}|k) = \int_0^\infty \frac{\lambda^k e^{-\lambda}}{k!}\mathcal{G}(\lambda;\alpha,\beta)\, d\lambda \tag{7.43}$$

$$= \frac{\beta^\alpha \Gamma(k+\alpha)}{k!\,(1+\beta)^{k+\alpha}\,\Gamma(\alpha)} \tag{7.44}$$

$$= \left(\frac{\mu}{\sigma^2}\right)^{\frac{\mu^2}{\sigma^2}+1} \Gamma\left(k+\frac{\mu^2}{\sigma^2}+1\right)\left[k!\left(1+\frac{\mu}{\sigma^2}\right)^{k+\frac{\mu^2}{\sigma^2}+1}\Gamma\left(\frac{\mu^2}{\sigma^2}+1\right)\right]^{-1}, \tag{7.45}$$

where μ and σ^2 depend on $\vec{\theta}$ through $\vec{w}$.

7.3.2.4 A family of likelihoods

It is possible to generalize the choice of α and β in Eq. (7.41) by choosing a particular form of $\mathcal{P}(\lambda)$. Since the distribution of interest is a Poisson distribution, a well-motivated choice of $\mathcal{P}(\lambda)$ is a gamma distribution (the conjugate prior of the Poisson distribution) [194]; also see [179] for a recent discussion. Thus, we set $\mathcal{P}(\lambda) = \mathcal{G}(\lambda;a,b)$, where a and b are the shape and inverse-scale parameters of the gamma distribution, respectively. These hyper-parameters dictate the distribution of the Poisson parameter λ [195]. In line with our previous discussion, the gamma distribution prior implies that Eq. (7.41) becomes

$$\alpha = \frac{\mu^2}{\sigma^2} + a \text{ and } \beta = \frac{\mu}{\sigma^2} + b. \tag{7.46}$$

The rest of the likelihood derivation remains the same. This allows the choice of specific values for a and b to satisfy certain properties. Equation (7.41) is obtained with $a = 1$ and $b = 0$, corresponding to the uniform prior discussed above. Another interesting choice is to require that the mean and variance of $\mathcal{P}(\lambda|\mu,\sigma)$ match μ and σ^2, respectively. This can be achieved by setting $a = b = 0$, and we refer to this parameter assignment as $\mathcal{L}_{\text{Mean}}$. In the case of identical weights, $\mathcal{L}_{\text{Mean}}$ is equivalent to Eq. (20) in [179]. Both choices are improper priors, as technically they are limiting cases of the gamma distribution. However, we can use them to obtain proper $\mathcal{P}(\lambda|\mu,\sigma)$ distributions.

In [179], a convolutional approach is suggested for handling arbitrary weights. We refer to this likelihood as $\mathcal{L}_{\text{G}}$. Each weighted MC event has $\mathcal{P}(\lambda_i|w_i) = \mathcal{G}(\lambda_i;1,1/w_i)$, corresponding to the prior $\mathcal{P}(\lambda_i) = \mathcal{G}(\lambda_i;0,0)$, such that $\lambda = \sum_i^m \lambda_i$. The likelihood $\mathcal{L}_{\text{Mean}}$ is a good analytic approximation of the more computationally expensive calculation given in [179] for $\mathcal{L}_{\text{G}}$. The latter has time complexity $\mathcal{O}(k^2 m)$ where k and m are the number of data and MC events in the bin respectively. When assuming uniform priors, the convolutional approach does not recover Eq. (7.42) for identical weights, so it cannot be used as a generalization of $\mathcal{L}_{\text{Eff}}$.

7.3.2.5 Convergence of the effective likelihood

In this section we will show that, if the relative uncertainty of the bin content vanishes as MC size increases, $\mathcal{L}_{\text{Eff}}$ and $\mathcal{L}_{\text{Mean}}$ both converge to $\mathcal{L}_{\text{AdHoc}}$.

For positive weights w_i, the relative uncertainty σ/μ is bounded between zero and one. Uncertainty as large as the estimated quantity, $\sigma/\mu = 1$, occurs if and only if $m = 1$. In the limit that σ/μ goes to zero, Eq. (7.42) converges to $\delta(\lambda - \mu)$, and $\mathcal{L}_{\text{Eff}}$ and $\mathcal{L}_{\text{Mean}}$ both go to $\mathcal{L}_{\text{AdHoc}}$. We can see this by noting that the shape parameter, α, goes to infinity as the MC relative uncertainty goes to zero, turning the gamma distribution into a Gaussian distribution of mean α/β and variance α/β^2. This Gaussian converges to $\delta(\lambda - \mu)$ in the limit of vanishing σ/μ. Substituting into Eq. (7.30), we recover Eq. (7.22), which converges to Eq. (7.21) in the large MC limit.

It remains to be shown that the relative uncertainty of the bin content vanishes as MC size increases. For identical weights,

$$\lim_{m\to\infty} \frac{\sigma_{\text{identical}}}{\mu_{\text{identical}}} = \lim_{m\to\infty} \frac{1}{\sqrt{m}} = 0. \tag{7.47}$$

For arbitrary weights, the limit can be written in terms of the running average of w_i and w_i^2 as

$$\lim_{m\to\infty} \frac{\sigma}{\mu} = \lim_{m\to\infty} \frac{\sqrt{\langle w^2\rangle_m}}{\langle w\rangle_m \sqrt{m}}, \tag{7.48}$$

where $\langle w\rangle_m$ is the average over w_i and $\langle w^2\rangle_m$ the average over w_i^2 for $i \leq m$. This shows that as long as $\langle w^2\rangle_m$ does not grow much faster than $\langle w\rangle_m^2$, the limit will converge to zero. For weight distributions with positive support and finite, non-zero mean, this should be the case.

7.3.2.6 Behavior of the effective likelihood

It is instructive to examine the behavior of $\mathcal{L}_{\text{Eff}}$ for a single bin. It is standard to work with the log-likelihood $l(\mu, \sigma|k) \equiv -2 \ln \mathcal{L}(\mu, \sigma|k)$ and we do so here. Figure 7.1 shows the contour lines for $l_{\text{Eff}}(\mu, \sigma|k = 100)$. Since μ and σ are both dependent on the same underlying parameters, $\vec{\theta}$, a minimization over $\vec{\theta}$ can be thought of as a constrained minimization over μ and σ. This is visualized as the gray region in Fig. 7.1, which indicates where μ and σ are allowed to vary for some physics model[2]. Similarly, we can also visualize the standard Poisson log-likelihood, $l_{\text{Poisson}}(\mu|k = 100)$, which is simply l_{Eff} constrained along the line $\sigma = 0$.

To further illustrate the effect of the accessible region, we minimize l_{Eff} over μ for two possible constraints: fixed σ and fixed σ/μ. In terms of Eq. (7.32), a sufficient but not necessary condition for constant σ/μ with varying μ is equal weights, and a necessary but not sufficient condition for constant σ with varying μ is $m \geq 2$. For a standard Poisson likelihood, $\hat{\mu}_{\text{Poisson}} \equiv \min_\mu l_{\text{Poisson}}(\mu|k) = k$. Figure 7.2 shows

[2] A general bound for positive weights is $\sigma \leq \mu \leq \sigma\sqrt{m}$ which can be seen from their definitions.

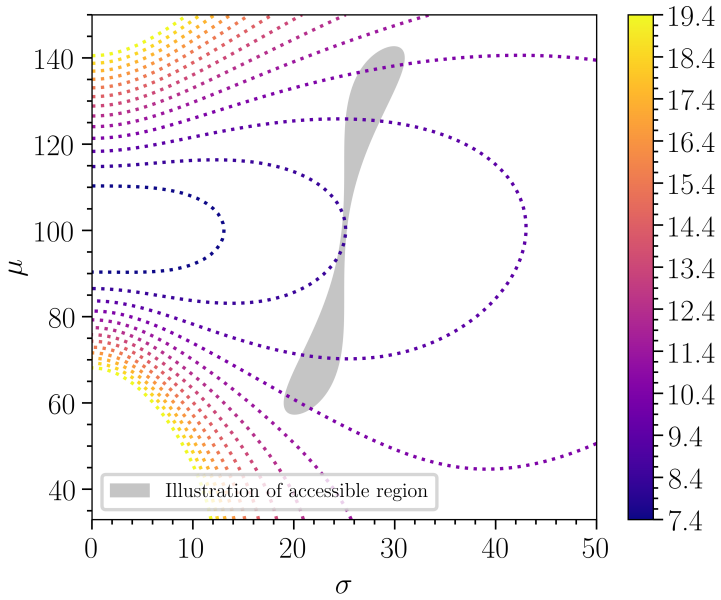

Figure 7.1: ***Likelihood contours and accessible region.*** Contours of constant $l_{\text{Eff}}(\mu, \sigma|k = 100)$. The accessible region (gray) illustrates where values of μ and σ may lie for a hypothetical physics model. A minimization over $\vec{\theta}$ can be thought of as a constrained minimization over the accessible region in μ and σ. Note that as σ increases the contours broaden.

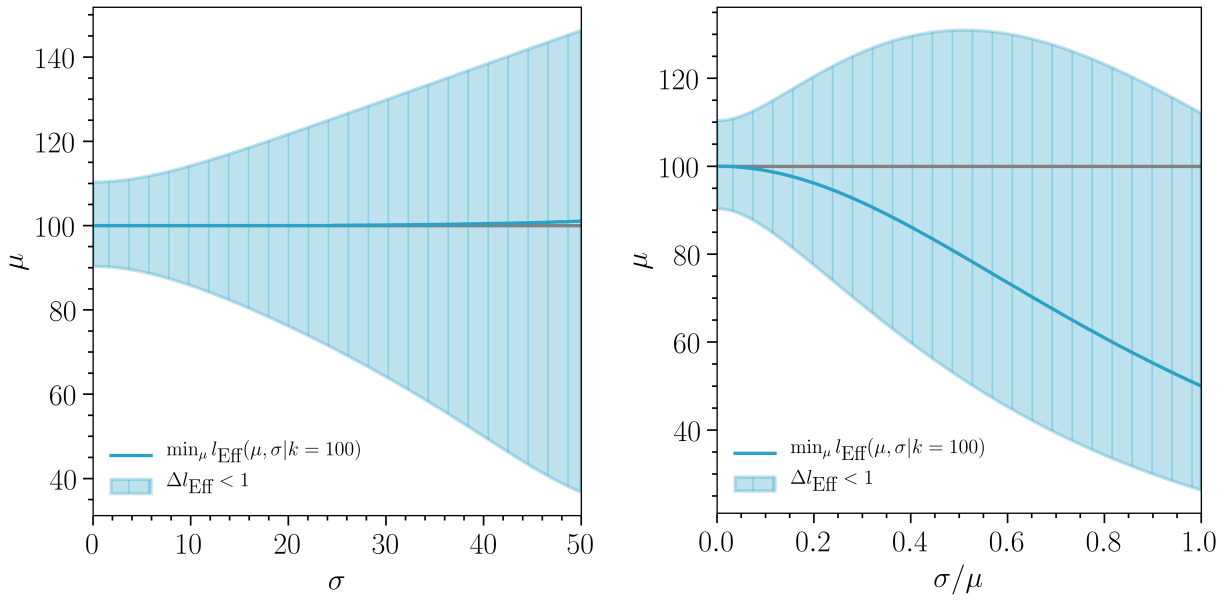

Figure 7.2: ***Slices of l_{Eff} for two accessible regions.*** This figure shows $l_{\text{Eff}}(\mu, \sigma|k = 100)$ minimized over μ while σ (left) and σ/μ (right) are held fixed. The minimum, $\hat{\mu}_{\text{Eff}}$, is shown as the solid blue line running through the center of the shaded regions. The shaded regions indicate where $l_{\text{Eff}}(\mu, \sigma|k) - l_{\text{Eff}}(\hat{\mu}_{\text{Eff}}, \sigma|k) < 1$. As σ goes to zero, the Poisson best-fit $\hat{\mu}_{\text{Poisson}} = 100$ is obtained. For fixed σ/μ, $\hat{\mu}_{\text{Eff}}$ deviates from $\hat{\mu}_{\text{Poisson}}$ as σ/μ increases.

$\hat{\mu}_{\text{Eff}} \equiv \min_\mu l_{\text{Eff}}(\mu, \sigma|k = 100)$ as well as the region where $l_{\text{Eff}}(\mu, \sigma|k) - l_{\text{Eff}}(\hat{\mu}, \sigma|k) < 1$ for fixed σ (left) and fixed σ/μ (right). Note that the shaded regions for fixed σ are calculated without requiring that $\mu \geq \sigma$, which would be the case for Eq. (7.32). As σ goes to zero, the Poisson best-fit and Wilks' 1σ interval are recovered. As σ or σ/μ increases, the shaded region becomes wider, as expected. For fixed σ, $\hat{\mu}_{\text{Eff}}$ does not deviate much from $\hat{\mu}_{\text{Poisson}}$, while for fixed σ/μ, $\hat{\mu}_{\text{Eff}}$ deviates from $\hat{\mu}_{\text{Poisson}}$ as σ/μ increases. The shaded regions correspond to the 1σ interval assuming the asymptotic approximation from Wilks' theorem and give a sense of the shape of $\mathcal{L}_{\text{Eff}}$ projected onto one-dimensional slices.

7.3.3 Example and performance

In practice, likelihoods such as those discussed above are used to estimate physical parameters from data. As discussed at the beginning of Sec. 7.3, weighted MC is often used to compute the likelihood of a particular physical scenario given the observed data. Statements are then made about the physical scenarios either by maximizing the likelihood or by examining the posterior distribution assuming some priors. We examine a toy experiment where we measure the mode, Ω, and normalization, Φ, of a Gaussian-distributed signal against a steeply falling inverse power-law background. The performance of $\mathcal{L}_{\text{Eff}}$ is evaluated and compared against other likelihoods.

For this toy experiment, we generate the true energies, E_t, of synthetic data events from a background falling as $(E_t/100\text{GeV})^{-\gamma_t^b}$, where $\gamma_t^b = 3.07$, and a Gaussian signal centered at $\Omega_t = 125\,\text{GeV}$ with width of $\sigma_t = 2\,\text{GeV}$ and normalization $\Phi_t = 5013$ for a fixed number of expected events. Our imaginary detector is sensitive in the $100 - 160\,\text{GeV}$ range. To simulate the effect of a real detector, the true energy, E_t, is smeared by 5 % for background and 3 % for signal to obtain event-by-event reconstructed energies, E_r. We generate a total number of MC events, N_{MC}, split evenly between the components. Generation is performed assuming inverse power-law distributions of $(E_t/100\text{GeV})^{-\gamma_g}$ for signal and $(E_t/100\text{GeV})^{-\gamma_g^b}$ for background. We choose $\gamma_g = 1$ and $\gamma_g^b = 2$. Reweighting of the MC can then be performed as a function of E_t and forward-folded onto distributions in E_r over which the events are histogrammed and likelihoods evaluated. A diagram of the steps described above is shown in Fig. 7.3. For all toy experiments, the background component, (Φ^b, γ^b), and the signal width, σ, are kept fixed to their true values. Only the signal mean, Ω, and normalization, Φ, are treated as free parameters.

7.3.3.1 Point estimation

Figure 7.4 shows the expectation in E_t as well as the data and $\mathcal{L}_{\text{Eff}}$ best-fit distributions in E_r. The leftmost panel shows the expectation for both signal and background assuming no smearing in E_t. The three other

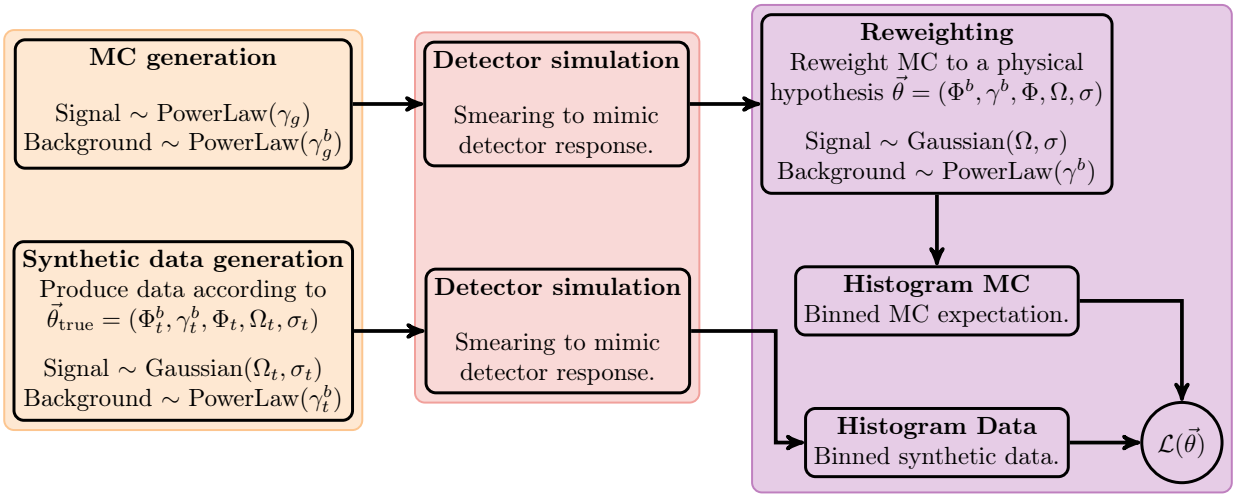

Figure 7.3: ***Diagram of toy experiment steps.*** The three colored boxes indicate the three steps of our toy experiment. The left box (almond) summarizes the MC and data generation. The center box (salmon) indicates the step in which we apply the detector response. The right box (lilac) summarizes the MC reweighting, data and MC histogramming, and final likelihood evaluation from the histograms. This final lilac box is repeated for each likelihood evaluation.

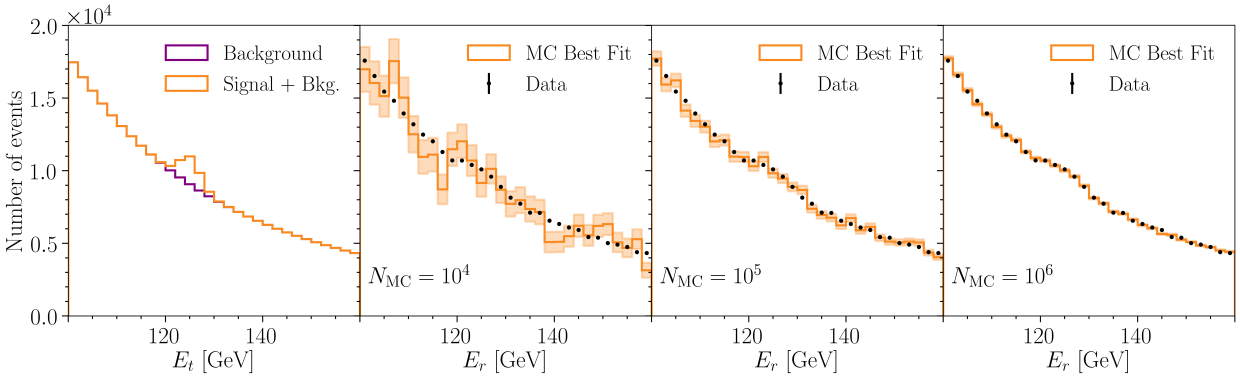

Figure 7.4: ***Benchmark scenario.*** Leftmost panel: the underlying distribution that data is drawn from with background (purple) and total rate (orange). Right three panels: observed data (black), $\mathcal{L}_{\text{Eff}}$ best-fit MC distributions (orange), and MC uncertainties (orange band), with increasing MC size from left to right. The background component and signal width are fixed, while the mean and normalization of the signal peak are fit by maximizing the likelihood with respect to those parameters.

panels show the smeared, E_r, distribution for data (black) and the best-fit result from $\mathcal{L}_{\text{Eff}}$ for three different MC datasets (orange) of varying MC size. The smeared shape of the signal peak is clearly visible in data, but not in the smallest size MC. As the MC increases in size, the best-fit MC distribution can be seen to converge to the data distribution.

Likelihood	$N_{\text{MC}} = 10^4$	10^5	10^6
$\mathcal{L}_{\text{AdHoc}}$	$(127.0, 6368.0)$	$(124.7, 5655.7)$	$(125.1, 4888.5)$
$\mathcal{L}_{\text{Eff}}$	$(127.1, 6077.1)$	$(124.7, 5576.0)$	$(125.1, 4889.4)$

Table 7.3: ***Best-fit parameters*** For the toy experiment shown in Fig. 7.4, best-fit parameters using $\mathcal{L}_{\text{AdHoc}}$ and $\mathcal{L}_{\text{Eff}}$ are shown. The columns in the table are for the different MC sizes. The two numbers in parentheses for each entry correspond to Ω and Φ, respectively.

The best-fit values for the example shown in Fig. 7.4 are given in Table 7.3 for $\mathcal{L}_{\text{Eff}}$ and $\mathcal{L}_{\text{AdHoc}}$. As point estimators, both likelihoods return similar values. This is driven by the fact that the same underlying MC distribution is used to fit to the data. The effect of convoluting $\mathcal{P}(\lambda|\vec{w}(\vec{\theta}))$ mostly serves to broaden the likelihood space, while preserving the maximum within the constraints described in Sec. 7.3.2.6. In the large MC limit, both likelihoods can be used for unbiased point estimation, provided that the likelihood space is smooth enough for standard minimization techniques to probe the global minimum.

7.3.3.2 Coverage

Due to the higher computational cost of computing frequentist confidence intervals by generating pseudodata to estimate the test-statistic ($\mathcal{TS}$) distribution, it is common to use the approximation given by Wilks' theorem for the cases where the underlying assumptions hold. In the case of small MC, a likelihood description that neglects MC uncertainties may lead to under-coverage even for a large data sample. In this section, we will use $\mathcal{TS} = \Delta l = l(\vec{\theta}_{\text{true}}) - l(\hat{\vec{\theta}})$, where $\vec{\theta}_{\text{true}}$ and $\hat{\vec{\theta}}$ correspond to the true and best-fit (Ω, Φ), respectively. We evaluate the coverage properties, computed using the asymptotic approximation given by Wilks' theorem, of the two-dimensional fit over (Ω, Φ) for several likelihood constructions. These include the modified-χ^2, $\mathcal{L}_{\text{AdHoc}}$, $\mathcal{L}_{\text{BB}}$, $\mathcal{L}_{\text{Mean}}$, and $\mathcal{L}_{\text{Eff}}$. These five test-statistics were chosen on the basis of their computation speed and as tests of different approaches towards the treatment of weighted MC. Note that using Wilks' theorem is an approximation and in general coverage tests should be performed to determine if it holds for a given experimental setup.

Several configurations were tested, all under the assumptions of the toy experiment described in Sec. 7.3.3.1. The MC was generated for two different settings of the total number of events: 10^3 and 10^6. For each setting, 500 toy experiments were generated, their best-fits found, and their Δl evaluated. Each toy experiment was

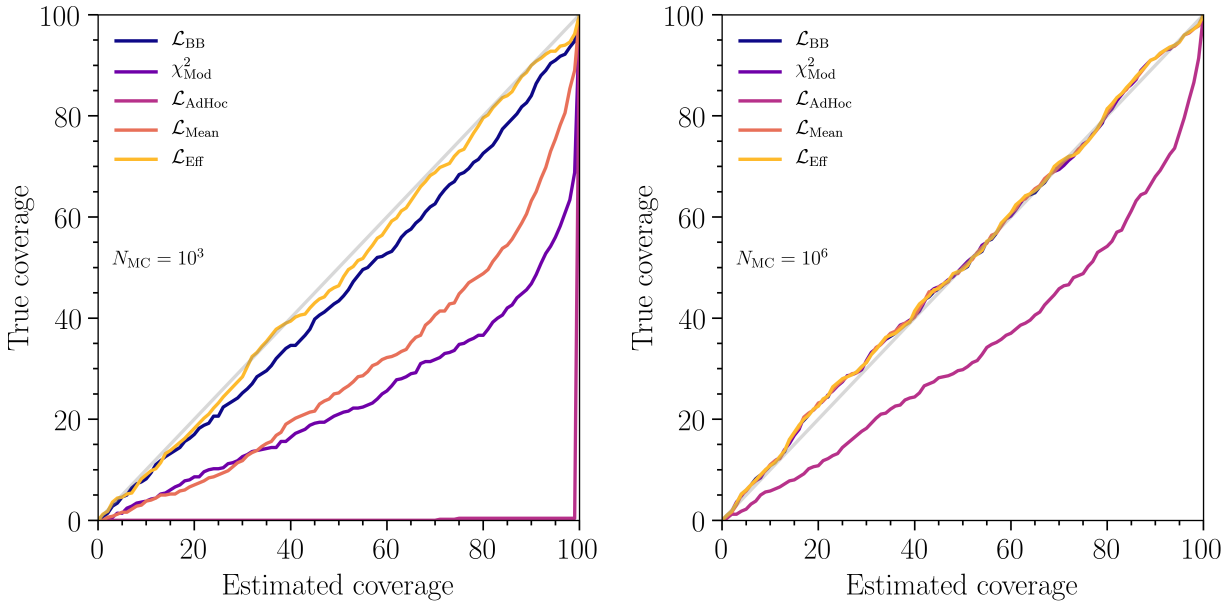

Figure 7.5: ***Coverage properties.*** True coverage of the Wilks' confidence interval for two sets of toy experiments with different MC sizes: 10^3 events (left) and 10^6 events (right). The ad hoc Poisson likelihood severely under-covers. The modified-χ^2, $\mathcal{L}_{\mathrm{BB}}$, and $\mathcal{L}_{\mathrm{Mean}}$ also undercover for small MC size. The effective likelihood, $\mathcal{L}_{\mathrm{Eff}}$, derived in Sec. 7.3.2.1 performs best.

classified as covering $\vec{\theta}_{\mathrm{true}}$ at a specified level p if $\Delta l < I(p; 2)$, where I is the inverse of the χ^2 cumulative density function and 2 indicates the number of degrees of freedom.

Figure 7.5 shows the percentage of times the true parameters were within the confidence intervals at level p as a function of the estimated coverage percentile for that level. First note that, as expected, the true coverage is highly dependent on MC size, with higher MC size leading towards improved agreement. In the case of $N_{\mathrm{MC}} = 10^3$, $\mathcal{L}_{\mathrm{BB}}$, $\mathcal{L}_{\mathrm{Mean}}$, modified-$\chi^2$, and $\mathcal{L}_{\mathrm{AdHoc}}$ all undercover to varying degrees of severeness. For $N_{\mathrm{MC}} = 10^6$, $\mathcal{L}_{\mathrm{AdHoc}}$ still under-covers, which is not surprising as it presumes zero MC uncertainty, but the other likelihoods exhibit good agreement. In this benchmark test, $\mathcal{L}_{\mathrm{Eff}}$ exhibits the best coverage properties. Note that these two examples represent extremely low simulation sample sizes in comparison to the data, and more realistic scenarios may not show such large deviations. Further tests have shown that these coverage properties hold true for cases where the background simulation size is small and the signal simulation size more appropriate; this is case for most IceCube analyses. We do not expect these coverage properties to hold for cases in which the asymptotic approximation breaks down for infinite simulation sample size. Finally, in the unlikely scenario where signal simulation samples are small and background simulation samples are large, the coverage of this method in the asymptotic approximation breaks down. In this last case other methods exist that may provide better coverage at the cost of additional computational complexity [180].

7.3.3.3 Posterior distributions

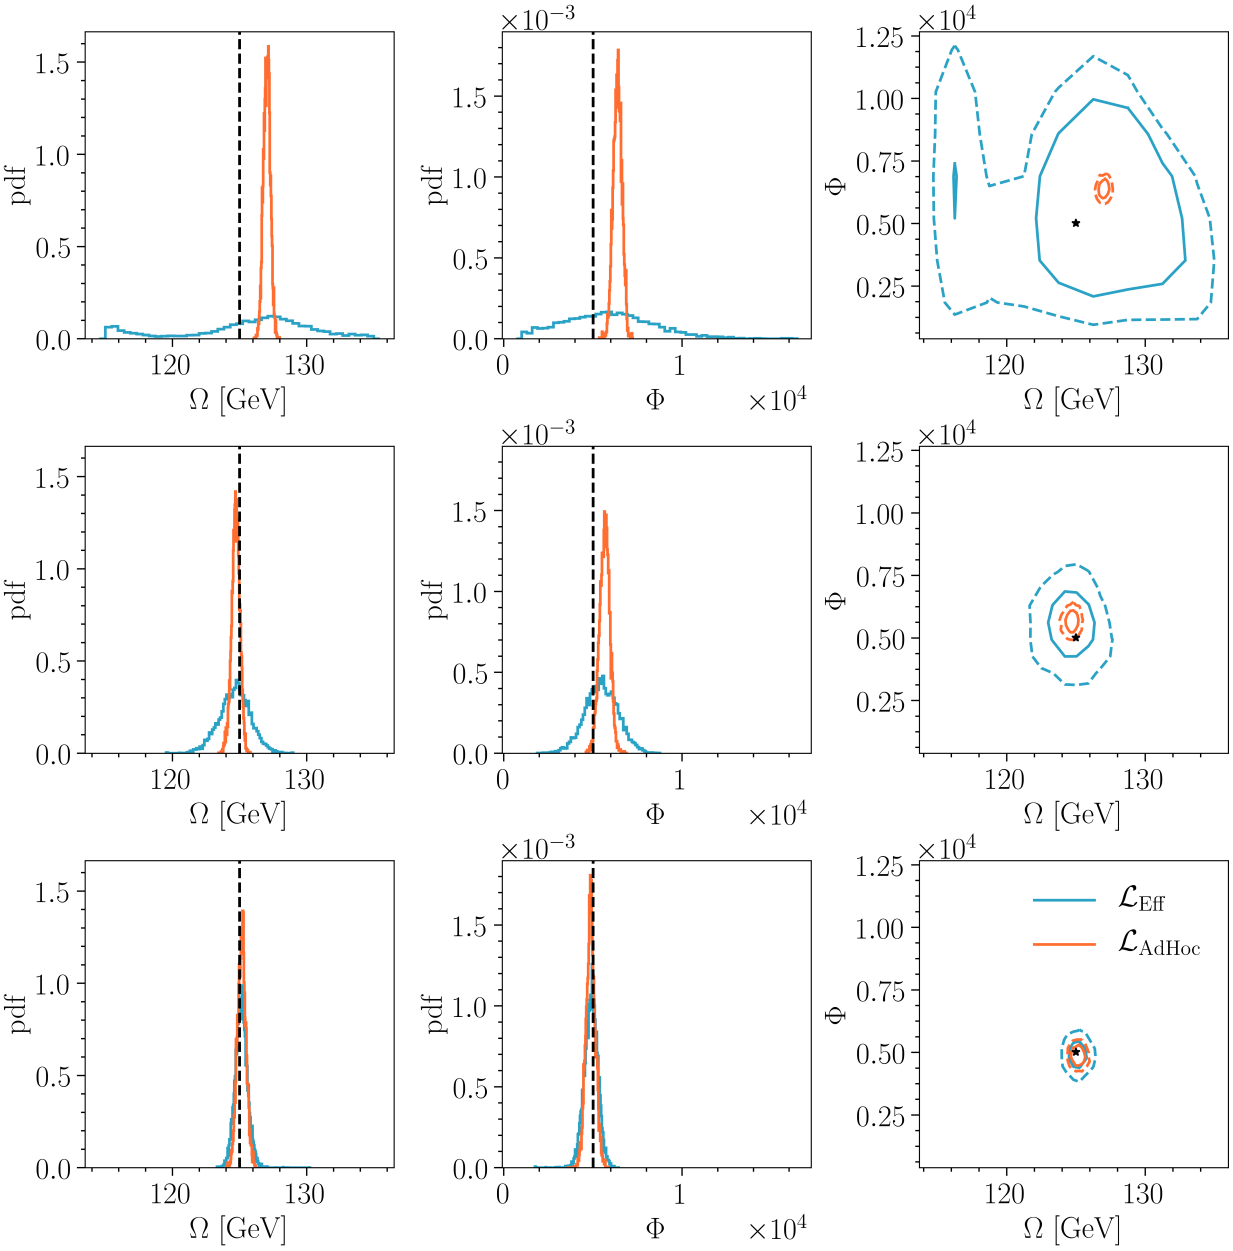

Figure 7.6: ***Posterior distributions in parameter space.*** Comparison of $\mathcal{P}(\vec{\theta}|k)$ for $\mathcal{L}_{\mathrm{Eff}}$ (blue) and $\mathcal{L}_{\mathrm{AdHoc}}$ (orange). Each horizontal row above uses a different MC set size, with $N_{\mathrm{MC}} = 10^4$, 10^5, and 10^6 from top to bottom. The left and center column show the marginal posterior distribution for the mass, Ω, and normalization, Φ, respectively. The true value is indicated by the dashed, vertical line. The rightmost column shows the joint posterior distribution with 68% (solid) and 95% (dashed) contours. The true values are indicated by the star.

It is also possible to use $\mathcal{L}_{\mathrm{Eff}}$ in a Bayesian approach. Using Bayes' theorem, the posterior is

$$\mathcal{P}(\vec{\theta}|k) \propto \mathcal{L}(\vec{\theta}|k)\pi(\vec{\theta}), \tag{7.49}$$

where $\pi(\vec{\theta})$ is a prior on the parameters. As evaluation of the normalization factor can by challenging, $\mathcal{P}(\vec{\theta}|k)$ can be approximated using a Markov Chain Monte Carlo (MCMC). For our toy example, we used `emcee` [173] to sample $\mathcal{P}(\vec{\theta}|k)$ under a uniform box prior for two different likelihood functions: $\mathcal{L}_{\text{Eff}}$ and $\mathcal{L}_{\text{AdHoc}}$. The sampling was performed using the data and MC sets described in Sec. 7.3.3.1.

Figure 7.6 shows the posterior distributions of Ω and Φ. For each comparison, $\mathcal{L}_{\text{Eff}}$ (blue) and $\mathcal{L}_{\text{AdHoc}}$ (orange) were sampled using the same underlying data and MC. We used 20 walkers with 300 burn-in steps followed by 1000 steps as settings for `emcee`. The left and center column show the marginal posterior distribution for the mass, Ω, and normalization, Φ, respectively. The true value is indicated by the dashed, vertical line. The rightmost column shows the joint posterior distribution with 68% (solid) and 95% (dashed) contours. The true values are indicated by the star. With $\mathcal{L}_{\text{AdHoc}}$, the true value of the parameter is highly improbable for the lower MC-size cases of the top and middle rows. In contrast, the posterior evaluated using $\mathcal{L}_{\text{Eff}}$ has increased width due to the reduced MC size. Even for $N_{\text{MC}} = 10^6$ (bottom row), the shape of the posterior evaluated using $\mathcal{L}_{\text{AdHoc}}$ is narrower than that using $\mathcal{L}_{\text{Eff}}$. Credible regions estimated using $\mathcal{L}_{\text{AdHoc}}$ would bias the result.

7.3.3.4 Performance

In this section we compare our performance with other treatments available in the literature in terms of the runtime cost per likelihood evaluation for a single bin. We perform our tests using a single `Intel® Core™ i5-8350U CPU @ 1.70GHz` running code compiled with `clang version 6.0.0-1ubuntu2`. We compute the likelihood CPU-evaluation time for the following likelihoods: $\mathcal{L}_{\text{AdHoc}}$, modified-$\chi^2$, $\mathcal{L}_{\text{G}}$ [179], $\mathcal{L}_{\text{BB}}$ [184], and $\mathcal{L}_{\text{Eff}}$. For each of them we consider increasing number of MC events from 10^2 to 10^6, increasing number of background components from 1 to 10^3, and increasing counts of data events from 10^1 to 10^4. Figure 7.7 shows the behavior of the runtime with respect to these quantities. All likelihoods have runtime that increases with the number of MC events, as seen in the leftmost panel of Fig. 7.7, as each likelihood must compute the sum of event weights which incurs an $\mathcal{O}(m)$ cost, where m is the number of MC events in the bin. Additionally, at low MC sample sizes the modified-χ^2 is faster than $\mathcal{L}_{\text{Eff}}$ since $\mathcal{L}_{\text{Eff}}$ requires the evaluation of more expensive special functions, however at larger MC sample sizes this additional cost is negligible compared to that of summing the MC weights. In the middle panel of Fig. 7.7 it can be seen that all likelihoods except $\mathcal{L}_{\text{BB}}$ are constant with respect to the number of background components as they only depend on summary statistics of the weight distribution. The Barlow-Beeston likelihood, $\mathcal{L}_{\text{BB}}$, incurs an $\mathcal{O}(bd \log d)$ cost for solving a single root finding problem per physical component, where b is the number of background components and d is the number of digits of precision, and therefore is not constant in runtime with respect to the number of components. However, one key difference between $\mathcal{L}_{\text{BB}}$ and $\mathcal{L}_{\text{Eff}}$ is that $\mathcal{L}_{\text{Eff}}$ must compute two summations

(the sum of the weights and sum of the square weights), while $\mathcal{L}_{\mathrm{BB}}$ needs only to compute a single summation of the MC weights. The rightmost panel of Fig. 7.7 shows the runtime as a function of the number of data events; for most likelihoods the number of data events, k, enters only in the evaluation of some special functions which for all practical applications are approximately constant in runtime. $\mathcal{L}_{\mathrm{G}}$ evaluates a special function which for these purposes can only be computed in $\mathcal{O}(k^2 m)$ time, resulting in the dependence on the number of data events. The $\mathcal{L}_{\mathrm{AdHoc}}$ treatment is always the fastest, but it does not incorporate MC statistical uncertainties in any way.

These tests make the behavior clear for single likelihood evaluations, and we expect similar trends for the runtime when applied to inference problems that are solved using MCMC techniques. However, additional complications can arise for inference techniques that make use of the profile likelihood. Global minimization can be affected by the numerical precision of the likelihood evaluation and any discontinuities that result. We have found such performance issues for $\mathcal{L}_{\mathrm{BB}}$ and $\mathcal{L}_{\mathrm{G}}$, while other techniques seem to have better numerical stability. While it may be possible to overcome these issues for $\mathcal{L}_{\mathrm{BB}}$ and $\mathcal{L}_{\mathrm{G}}$, currently known techniques have fallen short.

7.3.4 Conclusion

The use of MC to estimate expected outcomes of physical processes is nowadays standard practice. By construction, MC distributions are sample observations and subject to statistical fluctuations. MC events are also typically weighted to a particular physics model, and these weights may not be uniform across all events in an observable bin. A direct comparison of MC distributions to data is typically performed using $\mathcal{L}_{\mathrm{AdHoc}}$ or χ^2, where the expectation from MC is computed as a sum over weights in a particular observable bin. Such likelihoods neglect the intrinsic MC fluctuations and may lead to vastly underestimated parameter uncertainties in the case of low MC size. A better approach is to use a likelihood that accounts for MC statistical uncertainties.

Along with the definitions of μ and σ^2 in Eq. (7.32), the main result of this work is given in Eq. (7.45). This new $\mathcal{L}_{\mathrm{Eff}}$ is motivated by treating the MC realization as an observation of a Poisson random variate, computing the likelihood of the expectation using the MC and marginalizing the Poisson probability of observed data over all possible expectations. It is an analytic extension of the Poisson likelihood that accounts for MC statistical uncertainty under a uniform prior, $\mathcal{P}(\lambda)$. By assuming that the number of MC events per bin is the outcome of sampling a Poisson-distributed random variable, and that the SPD is a good approximation of the CPD for arbitrary weights, $\mathcal{L}\left(\lambda|\vec{w}(\vec{\theta})\right)$ can be written in terms of μ and σ^2 as shown in Eq. (7.38). This allows us to calculate $\mathcal{L}_{\mathrm{Eff}}$, given in Eq. (7.45), which can be directly substituted in

favor of $\mathcal{L}_{\text{AdHoc}}$. This construction is computationally efficient, exhibits proper limiting behavior, and has excellent coverage properties. In the tests performed here, it outperforms other treatments of MC statistical uncertainty.

7.3.5 Summary of likelihood formulas

Parameters	$\mu \equiv \sum_{i=1}^{m} w_i,\ \sigma^2 \equiv \sum_{i=1}^{m} w_i^2$
$\mathcal{L}_{\text{AdHoc}}$	$\frac{\mu^k e^{-\mu}}{k!}$
χ^2_{mod}	$\frac{(k-\mu)^2}{\mu+\sigma^2}$
$\mathcal{L}^{s=1}_{\text{BB}}$	$\max_{\bar{m}} \left\{ \frac{1}{k!m!} \left(\frac{\mu\bar{m}}{m}\right)^k \bar{m}^m e^{-\frac{\mu\bar{m}}{m}-\bar{m}} \right\}$
$\mathcal{L}_{\text{Mean}}$	$\left(\frac{\mu}{\sigma^2}\right)^{\frac{\mu^2}{\sigma^2}} \Gamma\left(k+\frac{\mu^2}{\sigma^2}\right) \left[k!\left(1+\frac{\mu}{\sigma^2}\right)^{k+\frac{\mu^2}{\sigma^2}} \Gamma\left(\frac{\mu^2}{\sigma^2}\right)\right]^{-1}$
$\mathcal{L}_{\text{Eff}}$	$\left(\frac{\mu}{\sigma^2}\right)^{\frac{\mu^2}{\sigma^2}+1} \Gamma\left(k+\frac{\mu^2}{\sigma^2}+1\right) \left[k!\left(1+\frac{\mu}{\sigma^2}\right)^{k+\frac{\mu^2}{\sigma^2}+1} \Gamma\left(\frac{\mu^2}{\sigma^2}+1\right)\right]^{-1}$

Table 7.4: ***Table of likelihood formulas.*** The likelihood functions discussed in this paper are given in each row. They are written in terms of μ and σ, whose explicit formulas are given in the top row, and the number of observed events, k, in the bin. In the case of $\mathcal{L}_{\text{BB}}$ we write the likelihood for the single-process case. Our main result and recommended likelihood, $\mathcal{L}_{\text{Eff}}$, is given in the last row.

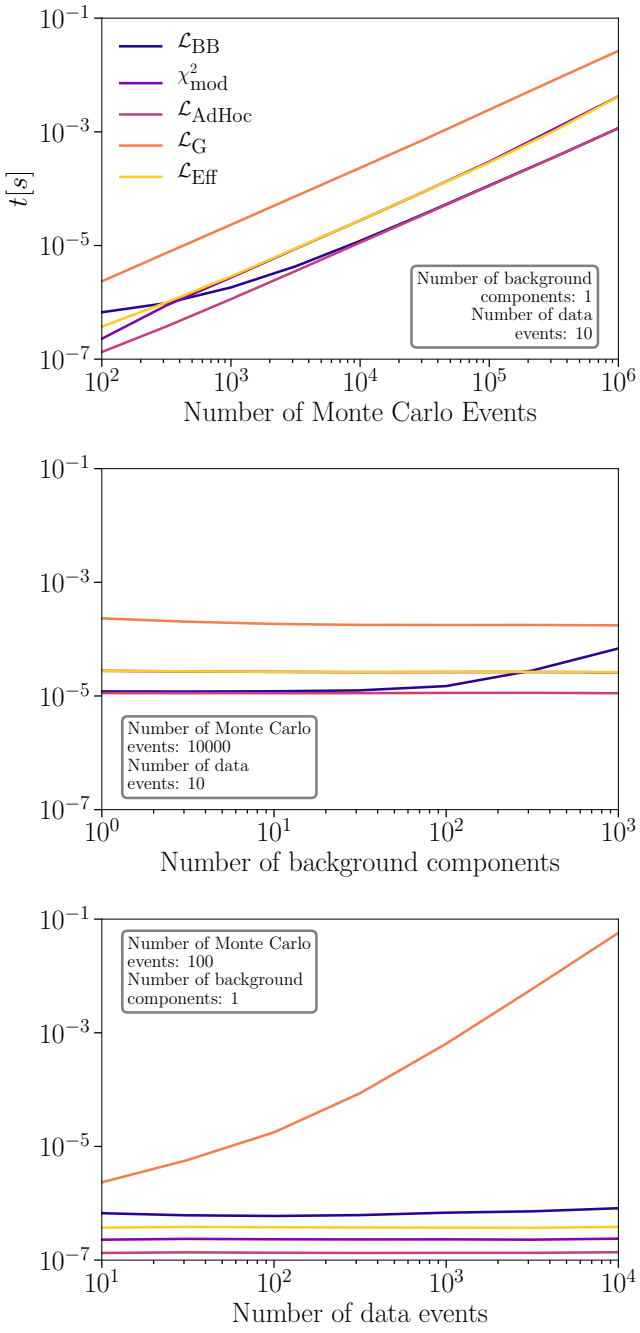

Figure 7.7: ***Likelihood function performance.*** Average single likelihood evaluation time is shown in the vertical axis in seconds. Different line colors show different likelihoods. Top-most panel: the number of MC events used is shown on the horizontal axis. Center panel: the number of background components is shown on the horizontal axis. Bottom-most panel: the number of data events is show on the horizontal axis.

7.4 Frequentist confidence intervals with nuisance parameters and limited simulation

Frequentist and Bayesian techniques deal with different two different kinds of probability. In frequentist statistics, the relevant probability is the frequency of the outcome of a repeatable experiment. Under this framework the important concepts are parameter estimation, confidence intervals, and statistical tests. In Bayesian statistics, the relevant probabilities come from the application of Bayes theorem which means we can define the probability density of parameters. This definition of the parameter p.d.f. is applicable to the same problems of parameter estimation, interval construction, and statistical tests but comes at the cost of defining "prior belief" about parameters.

In this section we will ignore the problem of statistical tests, instead focusing on the common features that underpin parameter estimation and interval construction. Generally in parameter estimation and interval construction there are two sets of parameters, parameters of interest $\vec{\theta}$ and nuisance parameters $\vec{\eta}$. Fundamentally there is no distinction between these two kinds of parameters. The difference is only in which parameters we want to infer information about.

For both parameter estimation and interval construction the likelihood function is central. The likelihood function reflects the plausibility of model parameters given observed data and is defined as $\mathcal{L}(\vec{\theta}, \vec{\eta}|\text{data}) = p(\text{data}|\vec{\theta}, \vec{\eta})$. Where $p(\text{data}|\vec{\theta}, \vec{\eta})$ is the probability of the data given the model parameters. A useful technique to eliminate nuisance parameters is the profile likelihood technique. Dropping the explicit notational dependence on data, the profile likelihood function is defined as

$$\tilde{\mathcal{L}}^{\texttt{profile}}(\vec{\theta}) = \max_{\vec{\eta}} \mathcal{L}(\vec{\theta}, \vec{\eta}), \tag{7.50}$$

where often the negative log of the function is maximized in place of the function for computational reasons. The profile likelihood is then only a function of the parameters of interest. Parameter estimation can be performed by maximizing the profile likelihood to obtain the "best-fit" parameters

$$\hat{\vec{\theta}} = \arg\max_{\vec{\theta}} \tilde{\mathcal{L}}^{\texttt{profile}}(\vec{\theta}). \tag{7.51}$$

This best-fit point in the parameter space is a derived property of the likelihood function. However, the same procedure can be performed with other functions to the same effect. In general a minimization procedure is used, and we refer to these functions as "test-statistics" (TS). A particularly useful TS is derived directly from the profile likelihood technique,

$$\text{TS}(\vec{\theta}) = -2\log\left(\frac{\tilde{\mathcal{L}}^{\texttt{profile}}(\vec{\theta})}{\tilde{\mathcal{L}}^{\texttt{profile}}(\hat{\vec{\theta}})}\right). \tag{7.52}$$

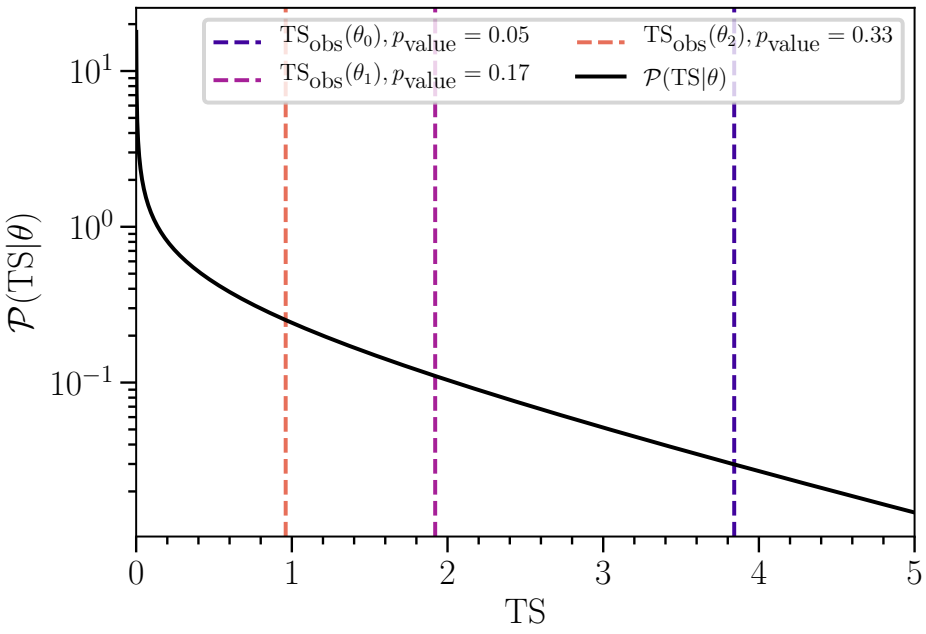

Figure 7.8: ***Test statistic distribution.*** An example of a test statistic distribution. Such distributions tend to have the bulk of their mass close to the lower boundary with a long tail. Lower values indicate better statistical compatibility with the data.

Using this TS to perform parameter estimation through minimization is mathematically equivalent to maximizing the likelihood, however, this form will prove to be uniquely useful for interval construction.

Since frequentist statistics deals with the frequency of outcomes from repeated experiments, we can use the TS that results from repeated experiments to construct probabilities. Consider for a moment a single point in the parameter space $\vec{\theta}_0$. At this point in the parameter space there is a distribution of data that can be observed, and therefore a distribution of TS functions. Instead of considering the distribution of TS functions originating from this point, we can simplify the picture by looking at the TS function only evaluated at this point $\mathrm{TS}(\vec{\theta}_0)$. This gives us a distribution of TS values for this point in the parameter space that may look like Fig. 7.8. It is important to note that for the profile likelihood TS and similar statistics a smaller TS value indicates better compatibility with the data. For this reason many statistical tests are constructed using a single tail significance, by comparing the TS from a single experiment to a background TS distribution and reporting a p-value that is the fraction of the TS distribution greater than the observed TS.

This procedure can be extended to construct intervals by considering the TS distributions of every point in parameter space and comparing to the observed TS function. Consider the one-dimensional case where there is a TS distribution for each value of the parameter, illustrated in Fig. 7.9 We can construct an interval

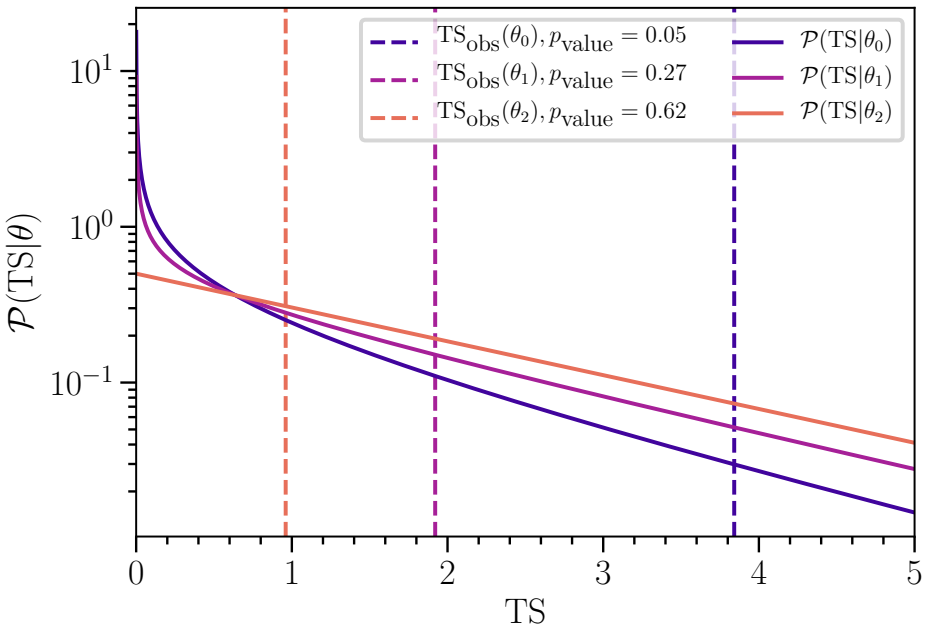

Figure 7.9: ***One-dimensional test-statistic distribution comparison.*** An example of the test statistic distributions as a function of a single parameter.

that will contain the true value of the parameter a fraction of the time α for repeated experiments. This interval is the collection of points in the one-dimensional parameter space where the TS at that point is greater than the α quantile of the corresponding TS distribution. If the TS distribution is the same for all points in parameter space, the interval construction can procedurally be thought of as drawing a horizontal line at the appropriate threshold and only including points that lie below the line. Varying TS distributions modify this procedure to the comparison of two curves. This procedure is not limited to one-dimension but can be extended to an arbitrary number of parameters of interest to construct n-dimensional regions with the same properties.

There is however an important caveat to this construction that appears when we consider nuisance parameters. In order for the intervals to have the desired properties, the observed TS must be greater than the threshold for all possible values of the nuisance parameters. This ultimatum presents several challenges. Nuisance parameters can often have a broad or even unbounded range of allowed values, meaning if the effect of nuisance parameters does not taper off at the extrema then almost all intervals are guaranteed to be empty. From a practical standpoint, computing the TS distributions for many points in parameter space is often done via Monte-Carlo and is computationally expensive. Adding additional dimensions to the parameter space for which we must compute TS distributions exponentially increases the computation time.

To combat these issues we can limit our interval construction to be valid for values of the nuisance parameters that are "reasonable". There are several methods for doing this, but we can split them into two categories: pure frequentist, and frequentist-Bayesian hybrid. In the pure frequentist approaches we can either choose a single value of the nuisance parameters, or work with a limited range of the nuisance parameter values. For the single value approach either nominal values are chosen before looking at the data, or estimators of the nuisance parameters are used to choose their values. This approach benefits from simplicity, but fails if the test-statistic distributions vary rapidly with changes to the nuisance parameters for values that we might consider "reasonable". A more expensive but robust approach is to explore the behavior of TS distributions for a limited range of the nuisance parameter values, which can be chosen *a priori* or from data-based bounds on the nuisance parameters. If we are willing to consider a hybrid approach, then some more pragmatic options are available.

Although Bayesian methods could be used to choose a single point in parameter space from which to generate the TS distributions, the more interesting application is one that uses a distribution in parameter space. In Bayesian statistics we can directly assign a probability density to the points in parameter space, either based on our prior information, or directly informed by the posterior distribution, from this extended perspective the probability of certain nuisance parameter values is of interest when considering the frequency of TS values for different parameters of interest. The prior case is simple in that we sample from the nuisance parameter priors when generating the TS distribution which allows us to account for variability introduced by the nuisance parameters without relying on hard cutoffs or biasing our inferences with parameter values that are unrealistic. This prior based technique is well motivated if the priors are derived from external observations, however in the case where nuisance parameters have broad or "uninformative" priors this motivation and benefit may break down. In some cases we expect nuisance parameters to be heavily constrained by the same data sample used to investigate the parameters of interest, so a different approach is merited. The alternative is to use the posterior distribution to construct our nuisance parameter p.d.f. Ideally a posterior distribution would be computed for each point in the parameters of interest space by fixing those parameters of interest. In this way the nuisance parameter posterior used for sampling depends on the point in parameter space we are examining.

With the possible solutions available, we can now look at the problem of limited simulation size when generating test statistic distributions. As explored in Section 7.3 for binned Poisson likelihood problems, the real expectation in data or simulation for the number of events in a bin is not a known quantity. Because the real expectations are not known, it is impossible to exactly model the distribution of TS that are expected for a particular point in the parameter space. However, as Section 7.3 also explored, limited simulation can be modeled with nuisance parameters so the techniques discussed above can be applied directly to the problem.

The "single point in parameter space" approach fails to address the additional uncertainty present in this case. Allowing for unbounded variation of the bin expectations fails as it is guaranteed to produce empty intervals. Bounding of the bin expectations within a reasonable range provides manageable intervals, but the dimensionality of the problem makes this computationally unfeasible beyond a handful of bins. Unfortunately this excludes all the "classic" frequentist solutions to this problem. The hybrid Bayesian-frequentist methods in this case provide a tractable solution that accounts for the additional uncertainty. We can make use of the treatment described in Section 7.3.2.3, where the bin expectation is derived to be gamma distributed, and the expected number of data events modeled to be Poisson distributed once this expectation is known. Practically this can be achieved by sampling data events from $\mathcal{L}_{\text{Eff}}$, or through a two-step process where the expectation is sampled from a gamma distribution $\mathcal{G}(\lambda; \alpha, \beta)$ where $\alpha = \frac{\mu^2}{\sigma^2} + 1$ and $\beta = \frac{\mu}{\sigma^2}$, and the data events are sampled from a Poisson distribution $\frac{\lambda^k e^{-\lambda}}{k!}$. It is important to note that this procedure only applies to variations in the data and should not be used to vary simulation expectations. This is because the TS distribution is intended to model variations in the data, whereas the simulation used for analysis is fixed. Combined with a similar hybrid treatment for other nuisance parameters, this provides a more complete accounting of the uncertainties given the available modeling.

Chapter 8

Analysis

The goal of the analysis in this chapter is to characterize the astrophysical neutrino spectrum using a high-purity sample of events. On its own, this may not provide the great insight that we desire into larger problems of astrophysical neutrino origin or cosmic ray acceleration mechanisms. However, the detailed measurements performed here are a vital piece of the puzzle that may help us to answer these larger questions in the long run. In pursuit of these measurements, a handful of new techniques were also developed, and these will certainly improve the other measurements we make going forward. This chapter combines the analysis components outlined in previous chapters with different definitions of the astrophysical neutrino flux and applies the statistical inference techniques of Section 7.1 to the analysis structure of Fig. 1.2. With this, a wide variety of tests and measurements of the astrophysical flux are explored, and their results examined.

8.1 Characterization of the astrophysical neutrino flux

IceCube has reported evidence of neutrino emission associated with a blazar [196, 197], as well as a 2.9σ excess with respect to background from a starburst galaxy [198]. However, these specific associations represent approximately 1 % of the astrophysical neutrino flux above 200 TeV [197], leaving the origin of the vast majority of the flux still unassociated with sources [199, 200, 201]. To cover different ideas about the origin of the astrophysical neutrino flux, we take a two-pronged approach to the spectrum's characterization. Section 8.1.1 considers generic forms for the spectrum, which could arise from numerous physical scenarios, while Section 8.1.3 tests a small sample of specific spectra from the literature. Separate from the astrophysical analyses, Section 8.1.2 explores the atmospheric flux of neutrinos from charmed hadrons. We provide the data, Monte Carlo, and tools necessary for these tests in the data release outlined in Appendix Section B.1 and encourage readers to perform their own tests of model compatibility [202].

In addition to the spectral models chosen, three key assumptions are made about the astrophysical

neutrino flux in this section: the flux incident on Earth is isotropic, it is the same between neutrinos and anti-neutrinos, and the same for each neutrino flavor. An isotropic flux is expected in models where the dominant contribution is from distant sources. We focus on the isotropic flux hypothesis as it is compatible with the available neutrino data. Finally, IceCube is insensitive to the differences between neutrino and anti-neutrino interactions on an event-by-event basis for most energies. So differences between the neutrino and anti-neutrino flux content do not modify the expectation of detected events in this analysis for most energies. The one region where this does not hold is for electron anti-neutrinos near 6.3 PeV, where their resonant interaction with atomic electrons called the Glashow Resonance (GR) [75, 76] can occur. This resonance enhances the expectation of detected down-going events near the resonance energy because of the larger interaction probability in the detection volume but reduces the expectation for up-going events because of increased absorption in the Earth [203, 204]. However, because the spectrum is steeply falling, as observed in previous analyses [60, 58, 205, 107, 64], the expected number of GR events is small in comparison to the rest of the sample; approximately 3.6 % above 60 TeV, assuming the best-fit spectrum from Section 8.1.1.1.

As described in Chapter 6, backgrounds from cosmic-ray showers produced in the Earth's atmosphere are a small but non-negligible contribution to the events observed in the sample; quantitatively above 60 TeV, 1.5 % atmospheric muons and 16 % atmospheric neutrinos. The background model used throughout this section includes atmospheric neutrinos from pions, kaons, and charmed hadrons, as well as atmospheric muons. Table 7.1 describes the parameters of the fit for the single power-law model and the priors (or constraints) associated with them. In addition to the astrophysical model and backgrounds with their respective nuisance parameters, detector systematics are also included. Only the astrophysical neutrino flux parameters differ between the models described in this section, the background models and detector systematics remain the same.

8.1.1 Generic models

Many well-motivated models of the astrophysical neutrino flux come in the form of power laws. This commonality stems from the possibility that astrophysical neutrinos and cosmic rays may share a common origin, and that we observe cosmic rays at Earth with a power-law spectrum. To examine the possibility of a power-law-like flux, we study a few generic scenarios for the astrophysical flux in the following sections.

The first scenario, called the "single power law" (SPL), is an unbroken power law across all energies with a freely varied normalization and spectral index. This is the simplest model as it only has two free parameters for the astrophysical spectrum: spectral index and normalization. It is also motivated by Fermi-acceleration, which predicts a power-law energy spectrum [206, 207].

The second scenario called the "double power law" (DPL), is the sum of two unbroken power-law spectra, both with freely varying normalizations and spectral indices. This potentially describes scenarios in which there are two populations of sources, two production mechanisms for high-energy astrophysical neutrinos that produce different power-law fluxes, or where neutrinos and anti-neutrinos have different spectra [208, 209].

A variety of source production models predict a high-energy cutoff in the neutrino spectrum, whether from limitations of the source energetics, a drop in pion production efficiency, energy losses of secondary pions and muons, or other mechanisms. To accommodate this possibility in a functionally simple way, we define a model with a single power-law astrophysical flux that has an exponential suppression at high energies. This third scenario, called the "exponential cutoff," has an additional parameter describing the cutoff's energy scale.

The fourth scenario, called the "log parabola" (LP), is a simple extension to the power law that adds a changing spectral index. This model is often used to describe gamma-ray spectra across many orders of magnitude of gamma-ray energy, which could otherwise be described as power laws in smaller energy ranges [210].

8.1.1.1 Single power-law flux

For the single power-law-flux scenario, an isotropic flux of astrophysical neutrinos is assumed incident on the Earth with a total differential all-flavor neutrino-plus-anti-neutrino spectrum given by

$$\frac{d\Phi_{6\nu}}{dE} = \Phi_{\texttt{astro}}\left(\frac{E_\nu}{100\,\text{TeV}}\right)^{-\gamma_{\texttt{astro}}} \cdot 10^{-18}\ \text{GeV}^{-1}\text{cm}^{-2}\text{s}^{-1}\text{sr}^{-1}, \tag{8.1}$$

where $\Phi_{6\nu}$ is the flux of the six neutrino species combined, $\Phi_{\texttt{astro}}$ is the normalization, and $\gamma_{\texttt{astro}}$ is the common spectral index. These two parameters are incorporated as arguments of the likelihood, according to Section 7.1. To better understand the relationship between the data and the neutrino flux contributions, we first look at projections in the two observables most different between neutrino fluxes (zenith, and energy) and compare data to the expectation from Monte Carlo assuming the nuisance parameters from the best-fit $(\hat{\vec{\theta}}, \hat{\vec{\eta}}) = \arg\max_{\vec{\theta},\vec{\eta}} \mathcal{L}(\vec{\theta}, \vec{\eta}) \cdot \Pi(\vec{\theta}, \vec{\eta})$. The bottom panel of Fig. 8.1 shows the data and expected number of events in bins of the cosine of the reconstructed zenith angle. In the down-going region, the data are well described with the addition of an isotropic astrophysical neutrino flux. Atmospheric components alone cannot describe the data well because the atmospheric neutrino components are suppressed in the down-going region by accompanying muons; see Section 8.1.2 for details. In fact, the null hypothesis (atmospheric only) is rejected with respect to the alternative hypothesis that includes an astrophysical component at approximately 8σ with this sample. The top panel of Fig. 8.1 shows the data and expected number of events in bins of reconstructed deposited energy. The region below 60 TeV is not included in the analysis because of larger

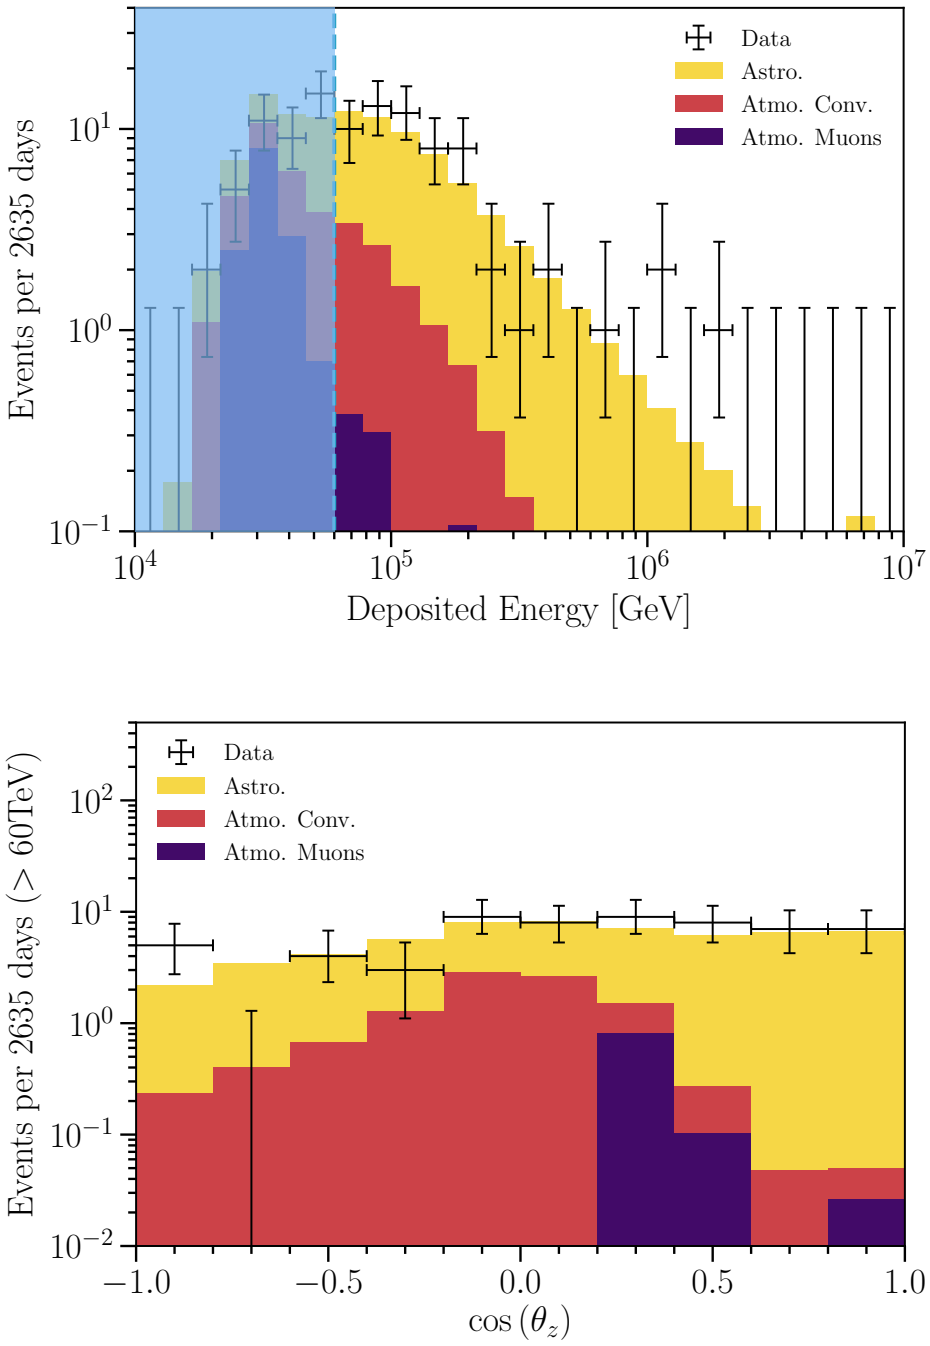

Figure 8.1: ***Deposited energy and reconstructed*** $\cos\theta_z$ ***distributions.*** In these panels the data is shown as crosses and the best-fit expectation as a stacked histogram with each color specifying a given flux component: astrophysical neutrinos (golden), conventional atmospheric neutrinos (red), and penetrating atmospheric muons (purple). Top: distributions of events and expected event count assuming best-fit parameters as a function of the deposited energy; events below 60 TeV (light blue vertical line) are ignored in the fit. Bottom: distribution of events with energy greater than 60 TeV in the cosine of their reconstructed zenith angle. Up-going events are on the left side of this panel and down-going events on the right. The expected number of events are split by components and displayed as a stacked histogram. The normalization of the prompt atmospheric neutrino component fits to zero and so is not shown in the stacked histogram. The distribution of data events appears to be largely flat as a function of cosine zenith with a small decline towards the up-going region. The lower event rate in the up-going region is expected as a result of the Earth's absorption of the neutrino flux, and appears to be compatible with the Monte Carlo expectation.

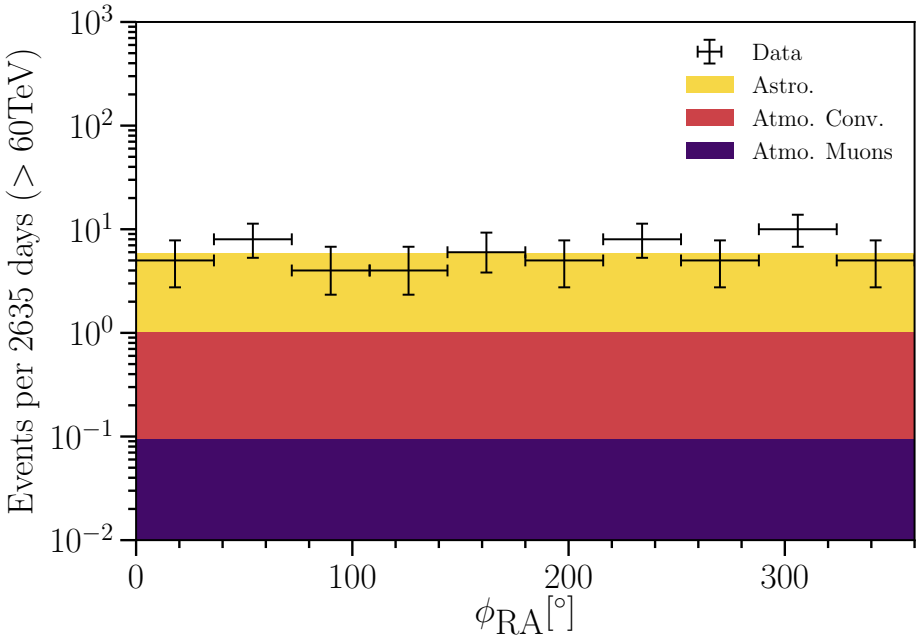

Figure 8.2: ***Right ascension distribution.*** Events above 60 TeV are shown together with the stacked expectation of different components using simulation weighted with the best-fit parameters for the single power-law astrophysical model. This figure has the same format as Fig. 8.1.

background uncertainties. However, we present the data to MC comparison in this region to demonstrate the level of agreement below the cut. From the stacked histogram, it is clear that above 60 TeV, the sample is dominated by the astrophysical component. Although not used directly in the analysis, we can look at another observable, the reconstructed right ascension. Under the assumption of isotropy for the astrophysical neutrino flux, this distribution should be uniform. Figure 8.2 is structured in the same manner as Fig. 8.1 but for the right ascension. From this distribution, the data appears to be compatible with the assumption of isotropy. Table 8.1 provides a comparison of the expected number of events in observable categories from the different components and compares this to the observed data.

8.1.1.1.1 Frequentist analysis The frequentist analysis of the single power law gives a best-fit point across all parameters, one-dimensional confidence intervals of each parameter, and two-dimensional confidence regions for the astrophysical normalization and spectral index. The one-dimensional results are summarized in Table 8.2, and are obtained by assuming that the TS is χ^2 distributed with one degree of freedom. We obtain a best-fit spectral index of $\gamma_{\texttt{astro}} = 2.87^{+0.20}_{-0.19}$. Fig. 8.3 shows the one-dimensional TS for $\gamma_{\texttt{astro}}$, $\Phi_{\texttt{astro}}$, and $\Phi_{\texttt{prompt}}$ on the diagonal panels as well as the bounds of the one-dimensional 68.3 % confidence regions plotted as vertical lines. Fig. 8.4 and the non-diagonal panels of Fig. 8.3 show the 68.3 % and 95.4 % confidence

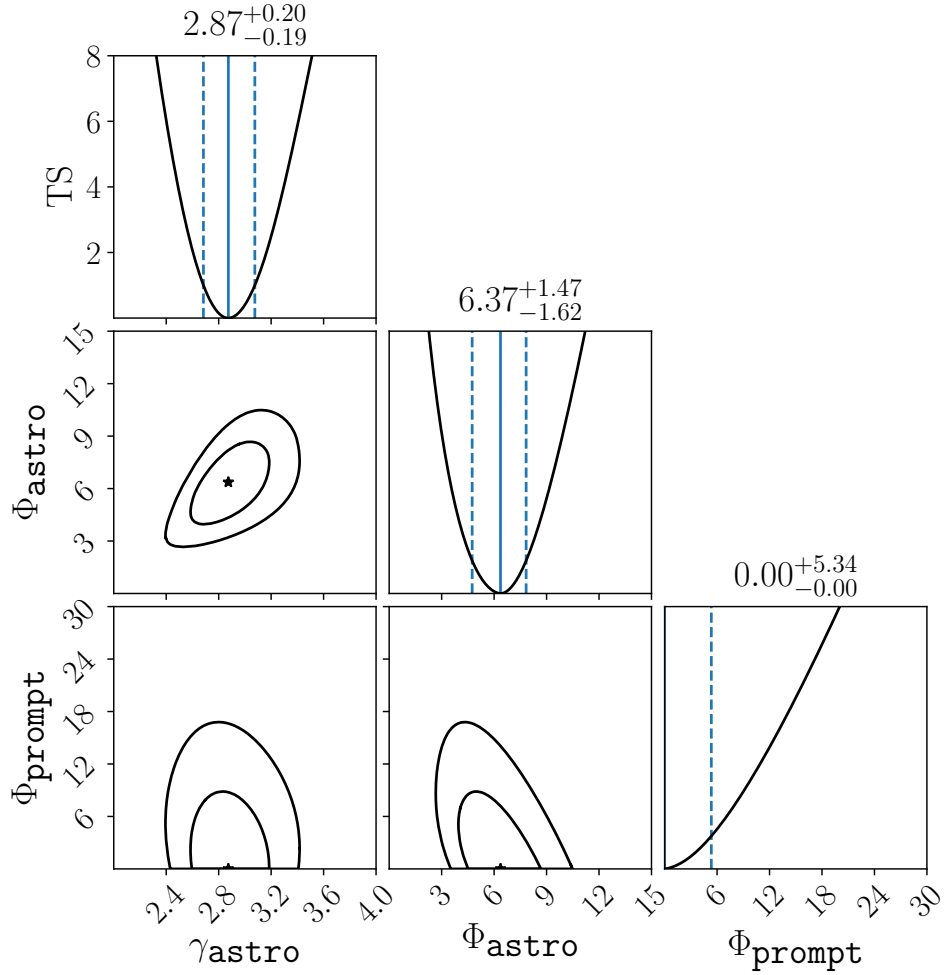

Figure 8.3: ***Single power-law profile likelihood.*** Diagonal panels show the TS, as a function of different model parameters, and the one sigma intervals assuming Wilks' theorem. Other panels show the best-fit point and two-dimensional contours. Solid (dashed) contours represent the 68.3 % (95.4 %) confidence regions assuming Wilks' theorem. The parameter $\gamma_{\mathtt{astro}}$ is the single power-law spectral index, $\Phi_{\mathtt{astro}}$ is a scaling factor of the astrophysical flux at 100 TeV, and $\Phi_{\mathtt{prompt}}$ is a scaling factor of the BERSS prompt neutrino flux calculation [117]; further descriptions of these parameters are provided in Section 7.1, Eq. (8.1), and Eq. (6.1)

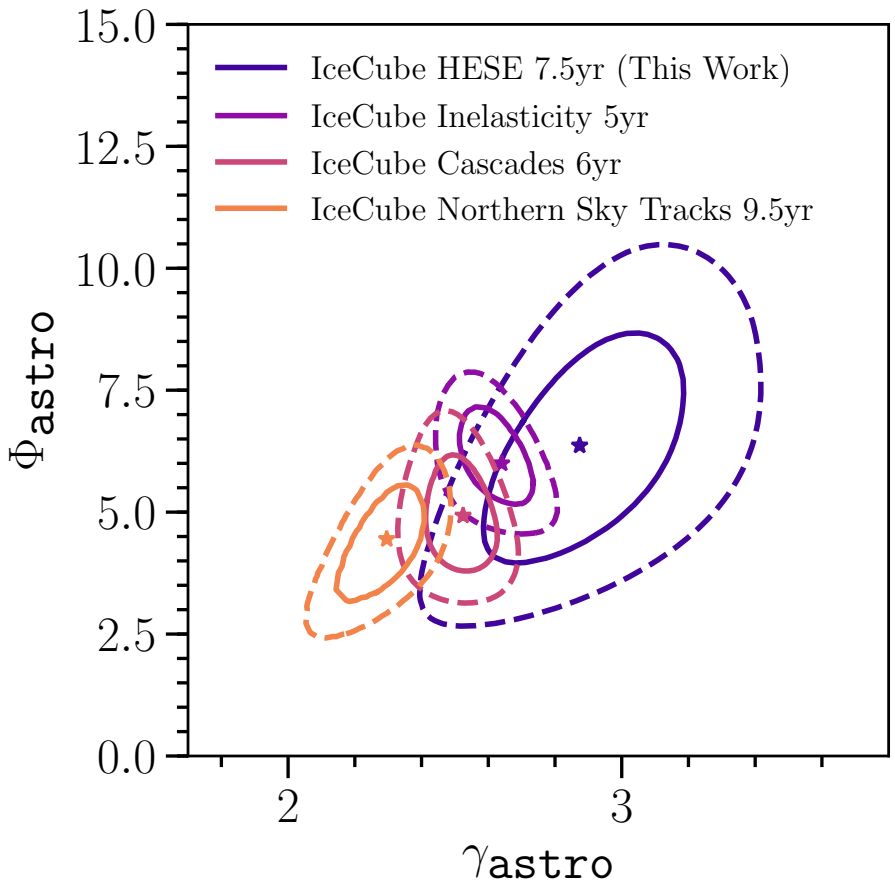

Figure 8.4: ***Comparison of single power-law parameters from different analyses.*** Assuming an unbroken single power-law model for the astrophysical neutrino flux, results from different IceCube samples are shown. The horizontal axis is the spectral index of the model and the vertical axis is six-neutrino flux normalization at 100 TeV given as a dimensionless multiplicative factor relative to 10^{-18} $\mathrm{GeV}^{-1}\mathrm{sr}^{-1}\mathrm{s}^{-1}\mathrm{cm}^{-2}$. The stars denote the different best-fit points, solid contours show the 68.3 % confidence region using the asymptotic approximation given by Wilks' theorem, and dashed contours show the 95.4 % confidence regions. Blue represents results from this work, while the purple shows results from IceCube's 5yr inelasticity measurement [205], salmon shows results from IceCube's 6yr cascade sample [64], and orange shows IceCube's 9.5yr Northern track sample preliminary result [107]. The differing preferred regions of parameter space for the astrophysical flux between the samples suggest a level of discrepancy, however a small region of parameter space is compatible with all samples at the 95.4 % level. Many checks have been performed for possible explanations of the discrepancy without definitive conclusions.

$E < 60\,\text{TeV}$	N_μ	N_{conv}	N_{prompt}	N_{astro}	Total	Data
Total Events	14.1	12.3	0.0	20.7	47.2	42.0
Up	0.0	7.6	0.0	10.4	18.0	19.0
Down	14.1	4.7	0.0	10.4	29.2	23.0
Cascade	1.2	6.9	0.0	17.5	25.6	30.0
Track	12.9	5.2	0.0	2.8	20.9	10.0
Double Cascade	0.0	0.2	0.0	0.4	0.7	2.0

$E > 60\,\text{TeV}$	N_μ	N_{conv}	N_{prompt}	N_{astro}	Total	Data
Total Events	0.9	9.1	0.0	48.4	58.4	60.0
Up	0.0	5.5	0.0	18.1	23.5	21.0
Down	0.9	3.6	0.0	30.3	34.9	39.0
Cascade	0.0	4.4	0.0	38.0	42.4	41.0
Track	0.9	4.5	0.0	8.2	13.7	17.0
Double Cascade	0.0	0.2	0.0	2.1	2.3	2.0

Table 8.1: ***Single power-law best-fit event expectations*** The left-most column indicates the event category, which may correspond to a particular choice of morphology, or a direction. The right-most column shows the number of data events observed for a given category. Each intermediate column corresponds to the expected number of events in the sample for a given source category: atmospheric muons (N_μ), conventional atmospheric neutrinos (N_{conv}), prompt atmospheric neutrinos (N_{prompt}), and astrophysical neutrinos (N_{astro}).

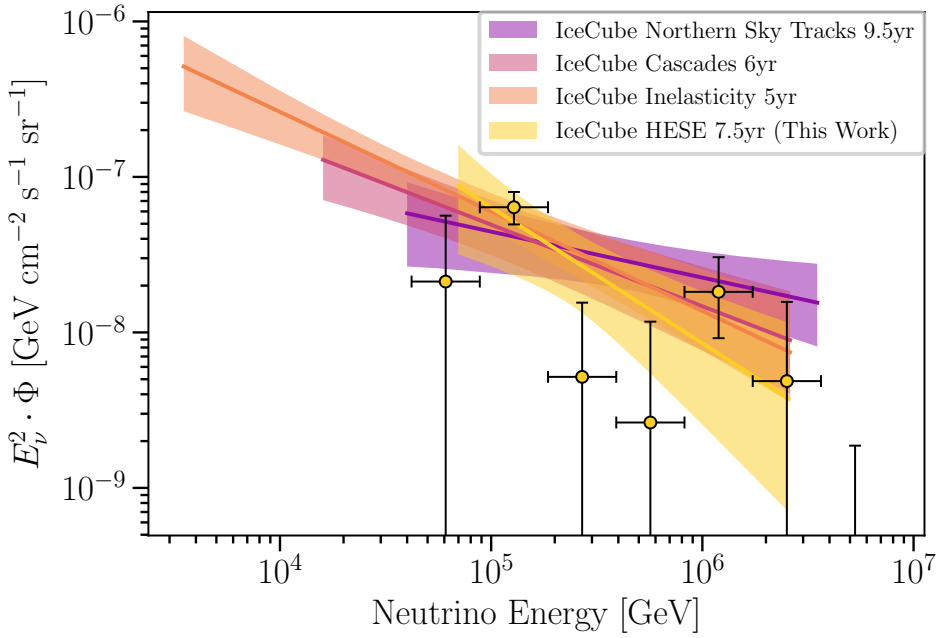

Figure 8.5: ***Comparison of single power-law flux from different analyses.*** As a complement to Fig. 8.4, this figure shows the best-fit single power-law flux for a variety of neutrino samples in their relevant energy ranges. A band is also plotted which encompasses the range of fluxes possible within the two-dimensional 68.3 % confidence regions of the astrophysical parameters. The yellow data points denote the best-fit parameters from the segmented power-law model described in Section 8.1.1.5, and the error bars are approximate 68.3 % confidence intervals for those normalization parameters.

regions for the two variables on the horizontal and vertical axes assuming two degrees of freedom. The impact of the systematics on the parameters of this model are shown in Fig. 8.6. The most relevant systematic affecting the astrophysical normalization is the DOM efficiency and the relative contribution of neutrinos from charmed hadrons. The astrophysical spectral index is more weakly affected by these systematics, but the normalization of the neutrino flux from charmed hadrons has the largest effect.

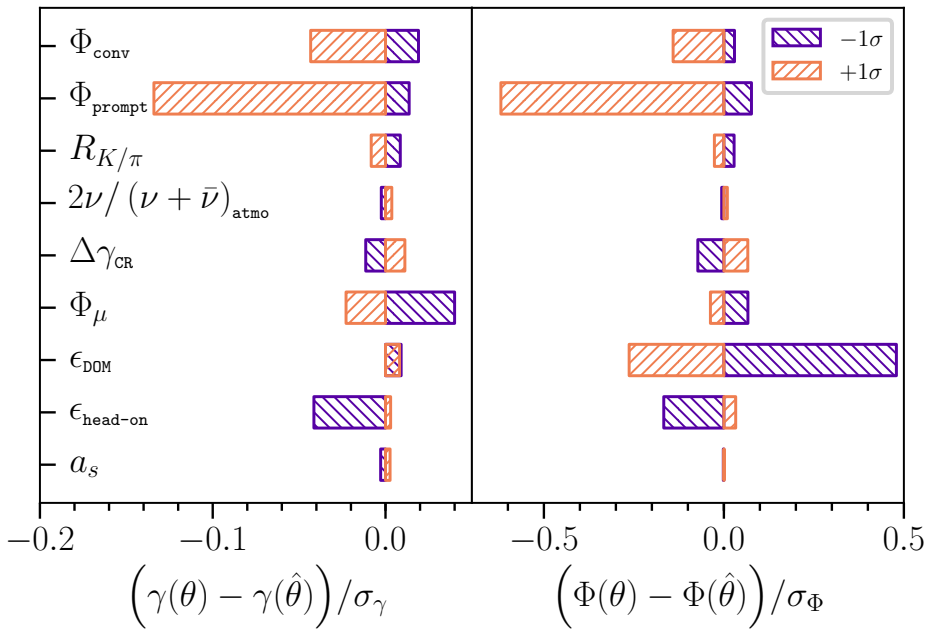

Figure 8.6: ***Impact of systematic uncertainties on the single power-law parameters.*** Each panel shows the impact of the systematic on the astrophysical spectral index (left panel) and normalization (right panel). The impact, horizontal axis, is defined as the change in the parameter of interest relative to its uncertainty when modifying one systematic nuisance parameter. Orange bars indicate the effect of increasing the value of the nuisance parameter from its maximum *a posteriori* (MAP) value by 1σ as defined by the nuisance parameter's 68.3 % highest posterior density region, while purple bars indicate the corresponding reduction of the parameter. Systematic parameter values are given in Table 8.2. The prompt normalization (Φ_{prompt}) and DOM efficiency (ϵ_{DOM}) have the largest effect on the astrophysical parameters. All other systematics pull the astrophysical parameters by significantly less than 0.5σ.

8.1.1.1.2 Crosschecks Our results agree with a previous iteration of this analysis [60] within the 2σ confidence regions of the astrophysical power-law parameters. The previous analysis obtained a best-fit spectral index of $\gamma_{\text{astro}}^{3\,\text{years}} = 2.3^{+0.3}_{-0.3}$, compared to $\gamma_{\text{astro}}^{7.5\,\text{years}} = 2.87^{+0.20}_{-0.19}$ in this analysis. This difference is primarily driven by a higher number of low-energy events observed in the latter 4.5 years compared to the first 3 years. The analysis energy range was also extended from 3 PeV to 10 PeV. Without events observed in this region, the spectral index shifted to a softer flux by ~ 0.1. Further extension of the analysis energy range produces negligible changes.

To investigate the shift in spectral index between analysis iterations, an *a posteriori* analysis of the data time-dependence was performed. Specifically, we compared a null hypothesis of a constant flux to a time-dependent spectrum with different astrophysical spectra for each of the two data partitions (first 3 years and latter 4.5 years), where each spectrum is modeled as a single power law. We performed a likelihood ratio based model comparison test, which disfavors the null hypothesis with a p-value of ~ 0.13. We conclude that

there is no evidence for time dependence in this data sample.

Additionally, we tested the effect of different systematics on the fit. We found that the inclusion or exclusion of any individual systematic or tested combination of systematics did not appreciably affect the fit result or uncertainties.

Other crosschecks were performed with the sample: comparing the spectrum of tracks and cascades, looking for differences between the up-going and down-going spectra, examining the summer and winter spectra, comparing the spectra from events in different regions of the detector, checking the charge distributions of events across many categorizations, looking for differences between charge calibrations, and checking for pulls resulting from reconstruction and simulation changes. None of these checks showed any statistically significant differences.

Although the uncertainty on $\gamma_{\texttt{astro}}$ is numerically similar between this analysis and the 3 years analysis, this is not the result of any additional systematic uncertainty or analysis change. This is a direct result of the change in the best-fit spectral index. With the same amount of data, harder spectra can be measured with less uncertainty than softer spectra. This effect is shown in Fig. 8.7, where we plot the uncertainty for different injected spectra ($\gamma_{\texttt{astro}} = \{2.3, 2.6, 2.9\}$) that have the same number of expected events in the sample.

8.1.1.1.3 Other samples Plotted in Fig. 8.4 are the confidence regions for other IceCube analyses. The orange contours show the results of a single power-law fit to IceCube's up-going muon neutrino data sample [107], the salmon contours show results from IceCube's 6yr cascade sample [64, 211], the purple contours show results from IceCube's 5yr inelasticity measurement [205], and the blue contour show results from this work. Assuming a continuous single power law across all energies, the large values of $\gamma_{\texttt{astro}}$ in the preferred regions in this analysis are disfavored by the through-going muon and cascade sample results. While these differences may be statistical, other explanations have been explored. A thorough examination of possible detector systematics and physics systematics has not revealed a systematic cause for the differences in single power-law best-fit parameters between samples. However, these samples cover different energies, flavors, regions of the sky, and are susceptible to different systematics and physical effects. Differences due to these factors could help to explain the different spectral measurements and have been tested for within the samples, although presently, we have not found evidence of a primary cause. Tests performed with the cascade sample reveal a preference for spectral softening in the tens to hundreds of TeV energy range [64]. The flux inferred for the overlapping energy range is well consistent with the results reported here. We briefly describe the four samples for the sake of comparison.

- The up-going muon neutrino sample [107], collected over 9.5 years, consists of well-reconstructed muon

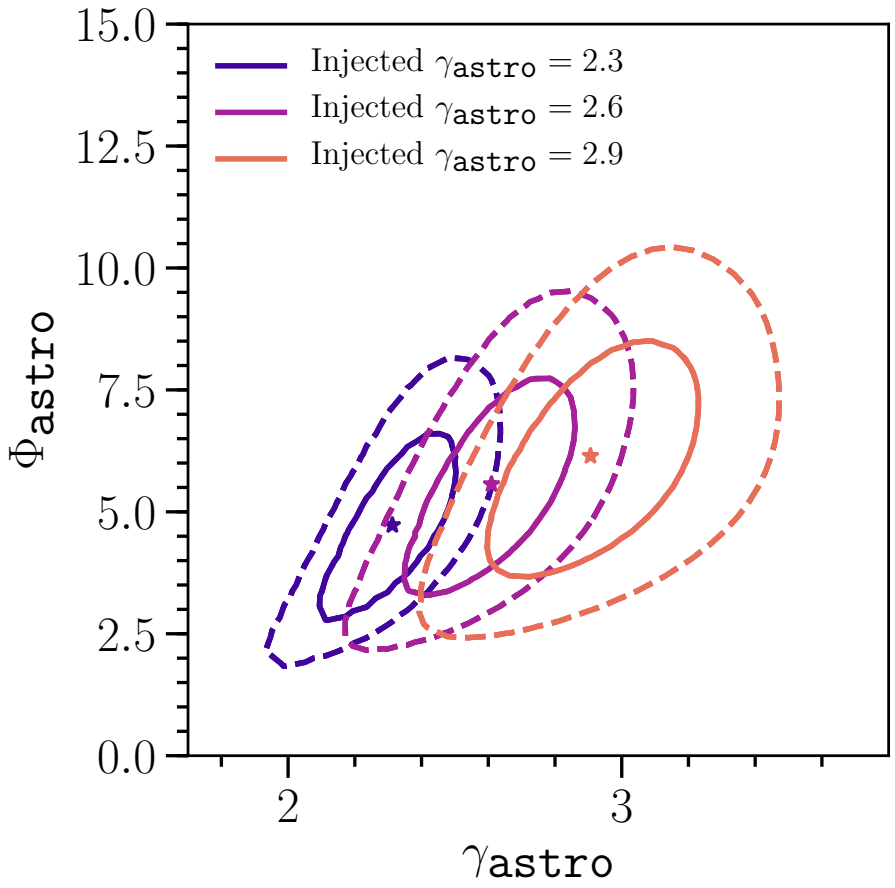

Figure 8.7: ***Astrophysical parameter sensitivity to injected spectra.*** The three colors from blue to salmon show the expected astrophysical parameter 68.3 % (solid lines) and 95.4 % (dashed lines) uncertainty for three injected astrophysical spectra $\gamma_{\texttt{astro}} \in \{2.3, 2.6, 2.9\}$, with corresponding normalizations $\Phi_{\texttt{astro}} \in \{4.59, 5.42, 6.00\}$. In each of these cases the expected number of events for 7.5 years of livetime is injected as data using nominal values for nuisance parameters and holding the number of injected astrophysical events to be equal between the spectra. As expected, the spectral index uncertainty grows with softer spectra from ~ 0.14 to ~ 0.21 when changing $\gamma_{\texttt{astro}}$ from 2.3 to 2.9.

tracks with zenith angle $\theta_z \geq 85^\circ$ that also pass a boosted-decision-tree based cut designed to select for through-going muon neutrino events while removing down-going muon and cascade backgrounds [58]. This sample contains muons of energy between $\sim 100\,\text{GeV}$ and $\sim 10\,\text{PeV}$, with the energy distribution peaked at $\sim 1\,\text{TeV}$. Atmospheric neutrinos dominate the sample, comprising $> 99\,\%$ events in it. The signal of astrophysical events is only apparent at the sample's high-energy range, where the atmospheric spectrum falls below the astrophysical component. At $\sim 20\,\text{TeV}$ in reconstructed muon energy the astrophysical component is $\sim 1/10$th the atmospheric component. The components are equal in flux at $\sim 200\,\text{TeV}$, and the atmospheric component is $\sim 1/10$th of the astrophysical component at $\sim 1\,\text{PeV}$. Events in this sample with neutrino energy between 40 TeV and 3.5 PeV contribute 90 % of the total observed likelihood ratio between the best-fit and the atmospheric-only hypothesis. As a function of the zenith angle, the signal-to-background ratio is lower at the horizon than for up-going events by almost an order of magnitude because of the enhanced atmospheric neutrino production at the horizon in the sensitive energy range. This sample benefits from better control of atmospheric flux systematic parameters due to the large population of atmospheric neutrinos. Additionally, this sample is substantially less affected by uncertainties related to muons from cosmic-ray air-showers than others because of the cuts on the reconstructed zenith angle.

- The cascade neutrino sample, collected over six years, consists of cascade-like events from all directions in the sky that have neutrino energies between $\sim 1\,\text{TeV}$ and $\sim 10\,\text{PeV}$ [64, 211]. Above 60 TeV about 60 % of events in this sample are not contained in the HESE sample. As the sample selects for cascade-like events, it predominantly contains electron and tau neutrinos, but also contains neutral current events from all neutrino flavors and a fraction of misidentified muon neutrinos. The sample has a sensitive energy range from 16 TeV to 2.6 PeV, which is defined as the smallest neutrino energy range for which a non-zero astrophysical flux is consistent with data at the 90 % confidence level. The distribution of the signal-to-background ratio for this sample has additional features compared to the same distribution for the up-going muon neutrino sample. These features are partly due to the rejection of atmospheric neutrino events by accompanying muons, which depends both on the neutrino energy and zenith angle. The signal-to-background ratio ranges from 1:100 at TeV energies to 1:1 at $\sim 20\,\text{TeV}$ to 1000:1 at PeV energies. The sample is least pure near the horizon with a factor of 10 to 100 less signal per background compared to the up-going and down-going regions. Like the up-going muon neutrino sample, this sample also benefits from better control of atmospheric flux systematic parameters due to the large population of atmospheric neutrinos, although to a lesser extent.

- The sample used for the inelasticity measurement, collected over five years, consists of track and cascade

events with their interaction vertex contained within the detector [205]. The sample is optimized to facilitate the measurement of the neutrino interaction inelasticity distribution, using both a veto and a boosted decision tree to select neutrino events while removing atmospheric muons. The sample is sensitive in the 1 TeV to 1 PeV energy range with the bulk of events below 10 TeV. Signal-to-background ratios of 10:1 are achieved for tracks close to 1 PeV and cascades above 100 TeV. Up-going track events in this sample are a factor of 10 to 1000 purer than down-going track events and a factor of 10 to 100 purer for cascades.

- In contrast to these samples, the HESE selection, which this work focuses on, has a similar effective area for all neutrino flavors and a signal-to-background profile with features closer to the cascade sample. Events in this sample with neutrino energy between 69.4 TeV and 1.9 PeV contribute 90 % of the total observed likelihood ratio between the best-fit and the atmospheric-only hypothesis. This signal-to-noise ratio in this sample is comparable to that of the cascade sample above 60 TeV and follows a similar dependence on zenith and energy. Above 60 TeV deposited energy, the sample has 60 events with a signal-to-background ratio greater than 1:10, 59 events with a signal-to-background ratio greater than 1:1, 24 events with a signal-to-background ratio greater than 10:1, and one event with a signal-to-background ratio greater than 10000:1. This variation in signal-to-background ratio stems from the different spectra of the fluxes and the varying rejection power of the veto with respect to the zenith angle.

HESE and the through-going muon neutrino sample have comparable sensitivity to the energy spectrum under the single power-law assumption when one accounts for the parameter-space differences between the best-fit spectral indices as demonstrated in Fig. 8.7 and the difference in sample livetime. The HESE sample suffers from a small sample size but benefits from high astrophysical purity, while the through-going muon neutrino sample benefits from a large sample size but suffers from lower purity and worse energy resolution of tracks. The cascade and inelasticity selections have comparable spectral sensitivity to each other and are more sensitive than the other two samples. Both samples benefit from a large sample size, and compared to the through-going muon sample, the cascade sample and inelasticity sample benefit from the better energy resolution of cascades and starting tracks, respectively. The cascade sample benefits from lower atmospheric neutrino contamination at high energies and improved ν_e effective area with respect to the HESE selection.

8.1.1.1.4 Bayesian analysis A Bayesian analysis was also performed, assuming a single power-law model for the astrophysical spectrum. The marginal posterior distribution is used for statistical inference, as described in Section 7.1. Table 8.2 reports the maximum *a posteriori* (MAP) estimation of each parameter, as

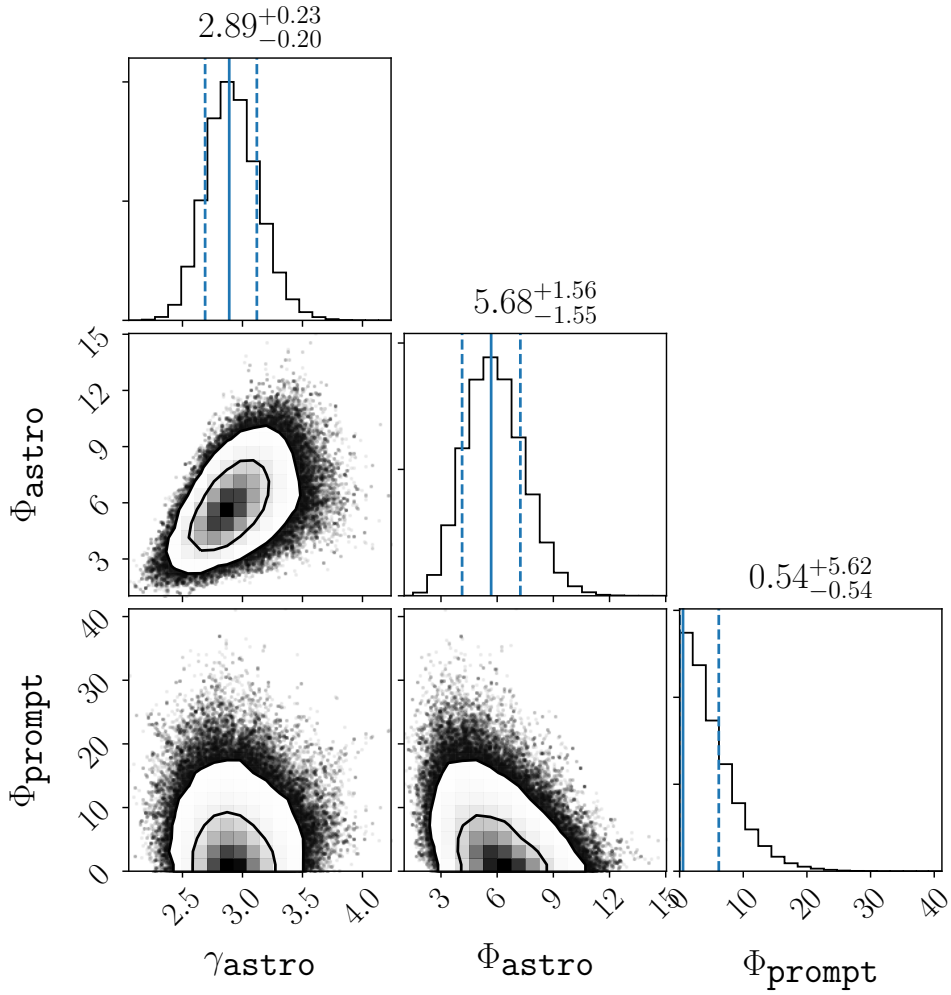

Figure 8.8: ***Single power-law parameters posteriors.*** Diagonal panels show the one-dimensional posterior distribution of the parameters (joint distribution integrated over all other parameters), where the horizontal axis of the panel is the same as the horizontal axis at the bottom of the column, and the vertical axis of the panel is the probability density in arbitrary scale. The solid blue lines denote the MAP estimator for each parameter, and the dashed lines denote the bounds of the one-dimensional 68.3 % HPD region, these numbers are also listed above the diagonal panels. Non-diagonal panels show the two-dimensional posterior distribution of the parameters, where the horizontal and vertical axes of the panel correspond to the horizontal axis at the bottom of the column and the vertical axis at the far left of the row respectively. The innermost contours show the two-dimensional 68.3 % HPD region and the outermost contours show the 95.4 % HPD region. The gray scale of the histogram within the contours shows the probability density in arbitrary scale. The points outside the contours show individual points from the MCMC.

well as the 68.3 % highest posterior density (HPD) region. Changing the prior of the SPL model parameters from linear-uniform to log-uniform has a much smaller effect than the reported errors of these parameters, implying the MAP is dominated by the data rather than the priors of the SPL parameters. The Bayesian analysis finds a most likely spectral index of $\gamma_{\texttt{astro}} = 2.89^{+0.23}_{-0.20}$, which is very similar to the frequentist estimation of these parameters. Figure 8.8 shows one-dimensional projections of the posterior distribution on the diagonal, integrated over the other variables, with the 68.3 % HPD region bounds plotted as vertical lines. The non-diagonal panels of Fig. 8.8 show contours of the 68.3 % and 95.4 % HPD regions of the two-dimensional posterior distribution projection. A histogram of the probability is displayed within the contours, and outside the contours, individual samples from the MCMC are plotted. One can see the correlation between $\gamma_{\texttt{astro}}$ and $\Phi_{\texttt{astro}}$; this correlation arises because the overall normalization of events must be roughly preserved to match the data well. Additionally, the inverse correlation of $\Phi_{\texttt{astro}}$ and $\Phi_{\texttt{prompt}}$ is apparent in this figure, which also arises from a conserved total number of events. Both the Bayesian and frequentist analyses show $\Phi_{\texttt{prompt}}$ to be compatible with zero. The parameter estimations and corresponding errors for the conventional atmospheric normalization are in good agreement with measurements of the atmospheric neutrino flux [212].

Parameter	Frenquentist Analysis		Bayesian Analysis	
	Best-fit value	68.3 % C.L.	Most-likely value	68.3 % H.P.D.
Astrophysical neutrino flux:				
$\Phi_{\texttt{astro}}$	6.37	$[4.75, 7.83]$	5.68	$[4.13, 7.24]$
$\gamma_{\texttt{astro}}$	2.87	$[2.68, 3.08]$	2.89	$[2.69, 3.12]$
Atmospheric neutrino flux:				
$\Phi_{\texttt{conv}}$	1.01	$[0.67, 1.35]$	0.93	$[0.61, 1.29]$
$\Phi_{\texttt{prompt}}$	0.00	$[0.00, 5.34]$	0.54	$[0.00, 6.15]$
$R_{K/\pi}$	1.000	$[0.901, 1.100]$	0.993	$[0.894, 1.095]$
$2\nu/(\nu+\bar{\nu})_{\texttt{atmo}}$	1.002	$[0.902, 1.102]$	0.986	$[0.901, 1.100]$
Cosmic ray flux:				
$\Delta\gamma_{\texttt{CR}}$	-0.053	$[-0.184, -0.005]$	-0.036	$[-0.088, 0.010]$
Φ_{μ}	1.19	$[0.75, 1.64]$	1.20	$[0.73, 1.61]$
Detector:				
$\epsilon_{\texttt{DOM}}$	0.952	$[0.886, 1.045]$	0.935	$[0.848, 1.002]$
$\epsilon_{\texttt{head-on}}$	-0.06	$[-0.54, 0.45]$	-0.07	$[-0.63, 0.39]$
a_s	1.00	$[0.80, 1.20]$	1.01	$[0.80, 1.20]$

Table 8.2: ***Single power-law model parameters.*** The frequentist analysis column shows the best-fit parameters and their corresponding 68.3 % C.L. interval according to Wilks' theorem for the single power-law model. The Bayesian analysis column shows the most-likely values of the parameters as well as the 68.3 % highest probability density (HPD) interval. Parameter name descriptions and priors(constraints) are given in Table 7.1.

8.1.1.2 Double power-law flux

The double power law is an extension to the single power law, which introduces a second power-law component with duplicated free parameters. In this model the astrophysical differential flux is described as

$$\frac{d\Phi_{6\nu}}{dE} = \left(\Phi_{\texttt{astro1}} \left(\frac{E_\nu}{100\,\text{TeV}} \right)^{-\gamma_{\texttt{astro1}}} + \Phi_{\texttt{astro2}} \left(\frac{E_\nu}{100\,\text{TeV}} \right)^{-\gamma_{\texttt{astro2}}} \right) \cdot 10^{-18}\ \text{GeV}^{-1}\text{cm}^{-2}\text{s}^{-1}\text{sr}^{-1}. \tag{8.2}$$

The astrophysical parameters of this model are then $\Phi_{\texttt{astro1}}$, $\Phi_{\texttt{astro2}}$, $\gamma_{\texttt{astro1}}$, and $\gamma_{\texttt{astro2}}$ which adds two more parameters compared to the single power law. The functional form of this model means that it is equivalent to the single power-law model when $\Phi_{\texttt{astro1}}$ or $\Phi_{\texttt{astro2}}$ become zero or when $\gamma_{\texttt{astro1}} = \gamma_{\texttt{astro2}}$. A limitation of this model is that it can only describe a hardening flux and cannot model a transition to a softer spectrum at higher energies. However, the cutoff and log-parabola models can describe this type of behavior and are explored in Sections 8.1.1.3 and 8.1.1.4. We choose improper uniform priors for the astrophysical model where the parameters are restricted to be positive, as was done for the single power law. There is implicitly a degeneracy of the parameters in the model. For the purpose of sampling the posterior distribution, this degeneracy is not an issue. However, it is more informative to construct some parameters that break the degeneracy and examine the posterior distribution in those new variables. We define

$$\begin{aligned} \gamma_{\texttt{hard}} &= \min(\gamma_{\texttt{astro1}}, \gamma_{\texttt{astro2}}), \quad \Phi_{\texttt{hard}} = \begin{cases} \Phi_{\texttt{astro1}} & \gamma_{\texttt{astro1}} \leq \gamma_{\texttt{astro2}} \\ \Phi_{\texttt{astro2}} & \gamma_{\texttt{astro1}} > \gamma_{\texttt{astro2}} \end{cases}, \\ \gamma_{\texttt{soft}} &= \max(\gamma_{\texttt{astro1}}, \gamma_{\texttt{astro2}}), \quad \Phi_{\texttt{soft}} = \begin{cases} \Phi_{\texttt{astro2}} & \gamma_{\texttt{astro1}} \leq \gamma_{\texttt{astro2}} \\ \Phi_{\texttt{astro1}} & \gamma_{\texttt{astro1}} > \gamma_{\texttt{astro2}} \end{cases}, \end{aligned} \tag{8.3}$$

to split the parameters into a hard component and soft component. These parameters are used to display the posterior distribution in Fig. 8.9 in the same manner as we did for the single power-law parameters. In Fig. 8.9 the normalizations of the two components are seen to be anti-correlated as the data require an astrophysical component, but both normalizations are compatible with zero within the two-dimensional 68.3 % HPD region of the bi-normalization posterior distribution. This indicates that two power-law components are not significantly preferred over a single power-law component.

As there are two regions of parameter space where this model is equivalent to a single power law we construct another set of variables to understand the behavior of the model

$$\Delta\gamma_{\texttt{astro}} = \gamma_{\texttt{soft}} - \gamma_{\texttt{hard}}, \quad \Delta\Phi_{\texttt{astro}} = \Phi_{\texttt{soft}} - \Phi_{\texttt{hard}}. \tag{8.4}$$

Fig. 8.10 shows that $\Delta\gamma_{\texttt{astro}}$ is compatible with zero, meaning again that the data is compatible with a single power law in this model.

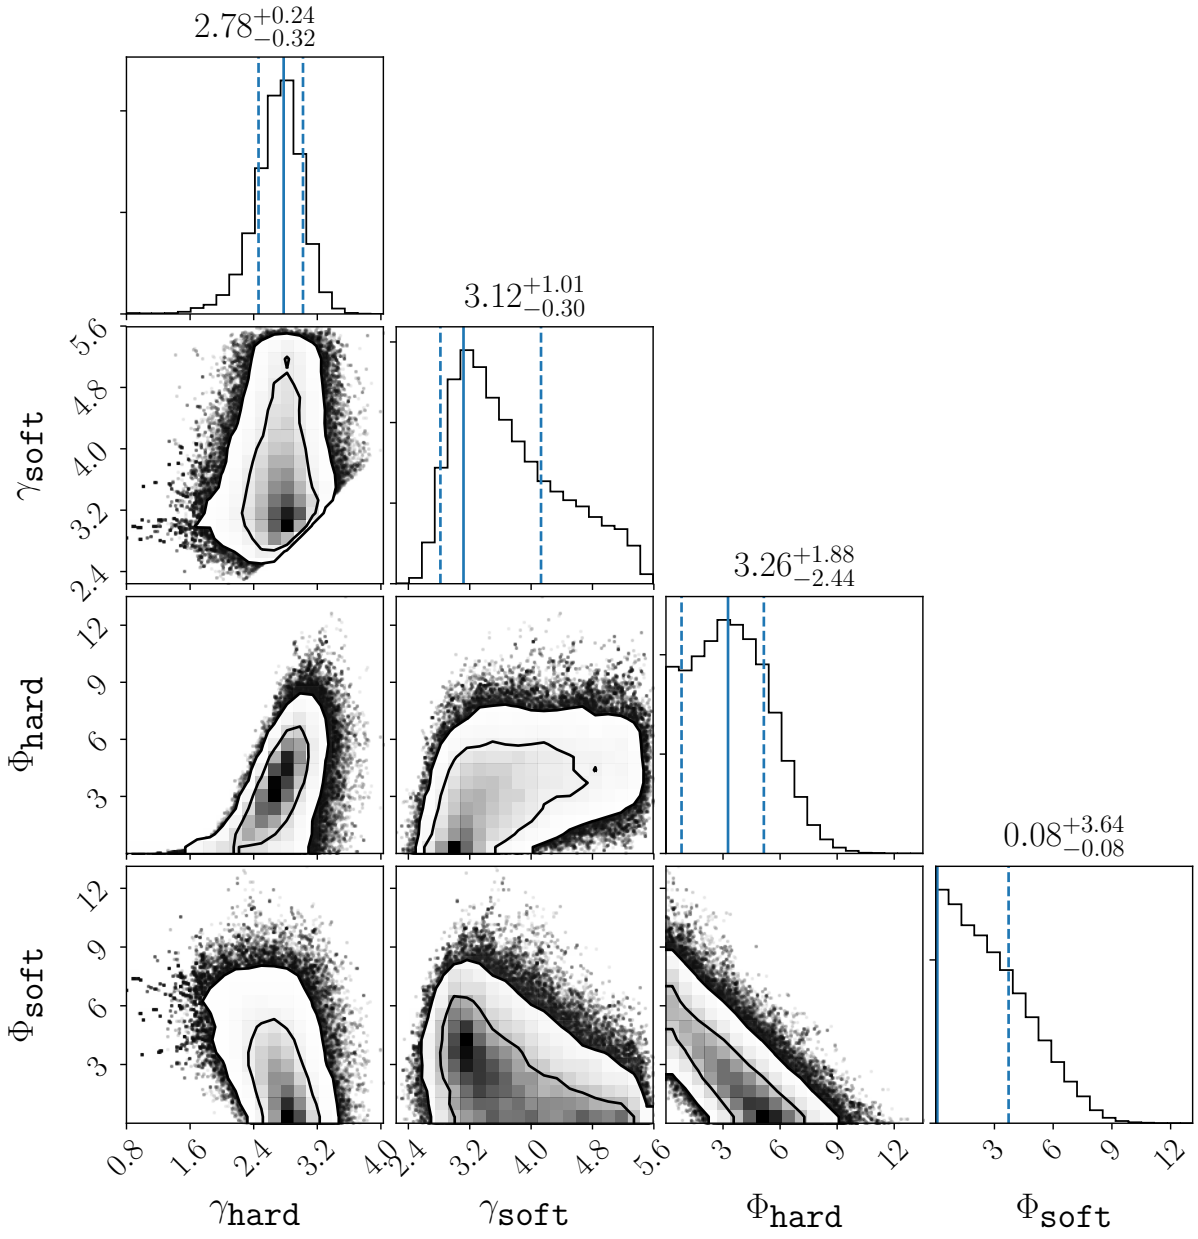

Figure 8.9: ***Double power-law model hard/soft parameters posterior distribution.*** Results derived from the posterior distribution of the model are shown in the same style as Fig. 8.8. The figure shows the one- and two-dimensional posterior distributions for the parameters of the hard and soft components of the astrophysical neutrino flux. The diagonal panels show the one-dimensional posterior with the parameter MAP estimation and 68.3 % HPD region indicated, while the non-diagonal panels show the two-dimensional posterior with 68.3 % and 95.4 % regions indicated.

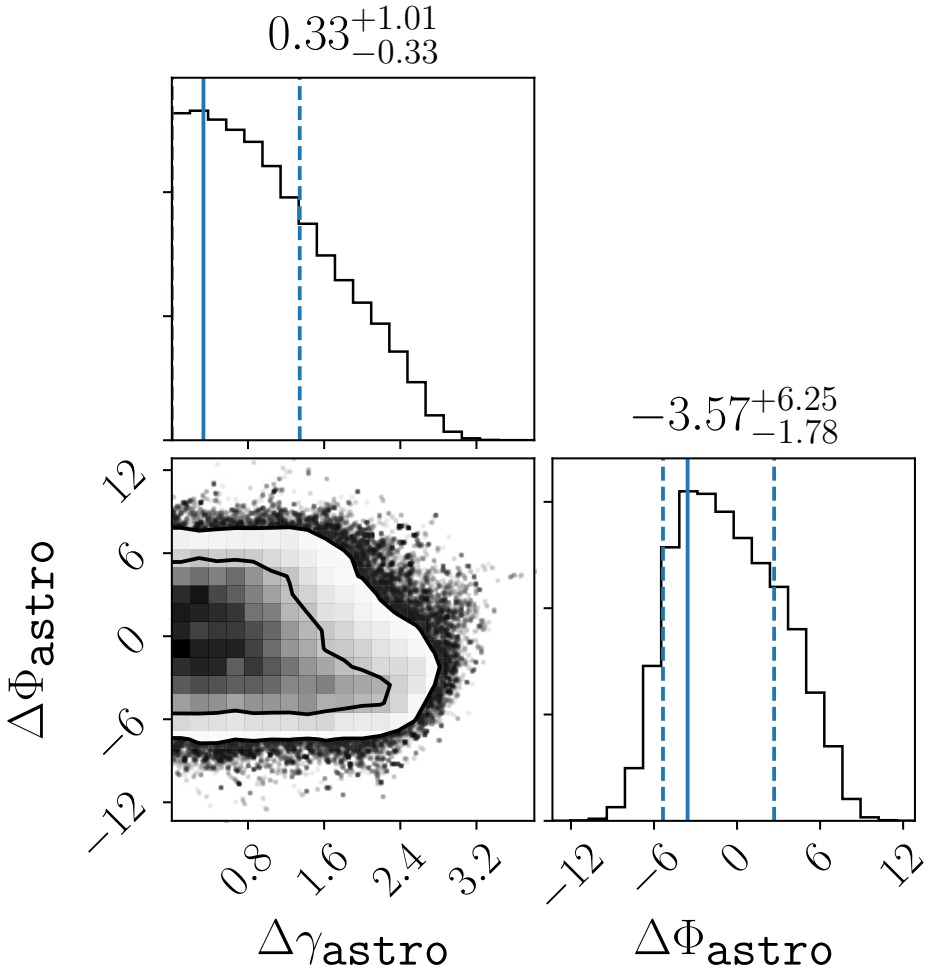

Figure 8.10: ***Double power-law model parameter deltas posterior distribution.*** Results derived from the posterior distribution of the model are shown in the same style as Fig. 8.8. The figure shows the one- and two-dimensional posterior distributions for the difference between the parameters of the hard and soft components of the astrophysical neutrino flux. The diagonal panels show the one-dimensional posterior with the parameter MAP estimation and 68.3 % HPD region indicated, while the non-diagonal panel shows the two-dimensional posterior with 68.3 % and 95.4 % regions indicated.

8.1.1.3 Single power law with spectral cutoff

The flux of the astrophysical component is given as

$$\frac{d\Phi_{6\nu}}{dE} = \Phi_{\texttt{astro}} \left(\frac{E_\nu}{100\,\text{TeV}} \right)^{-\gamma_{\texttt{astro}}} \cdot e^{-\frac{E_\nu}{E_{\texttt{cutoff}}}} \cdot 10^{-18} \text{GeV}^{-1} \text{cm}^{-2} \text{s}^{-1} \text{sr}^{-1}. \tag{8.5}$$

To study the preference between cutoff scenarios, we compute the Bayes factor for many values of the cutoff energy $E_{\texttt{cutoff}}$, where the null hypothesis is the single power-law model, and the alternative hypothesis is the cutoff model. The Bayes factor in this case is defined as

$$\mathcal{B}(E_{\texttt{cutoff}}) = \frac{\int d\vec{\eta}\, \mathcal{L}_{\texttt{cutoff}}(E_{\texttt{cutoff}}, \vec{\eta}) \cdot \Pi(\vec{\eta})}{\int d\vec{\eta}\, \mathcal{L}_{\texttt{SPL}}(\vec{\eta}) \cdot \Pi(\vec{\eta})}, \tag{8.6}$$

where $\mathcal{L}_{\texttt{cutoff}}$ is the likelihood of the cutoff model, $\mathcal{L}_{\texttt{SPL}}$ is the likelihood of the single power-law model, and Π is the set of priors given in Table 7.1. Fig. 8.11 shows the inverse of the Bayes factor $\mathcal{B}$ as a function of $E_{\texttt{cutoff}}$. For most values of $E_{\texttt{cutoff}}$, the Bayes factor is less than one; this implies that the data in this sample favors a model with no cutoff in most cases. For regions where $\mathcal{B} < 1$ we can exclude values of the cutoff with some level of certainty. Fig. 8.11 shows excluded regions of the cutoff chosen according to Jeffreys' scale.

In addition to the Bayes factor treatment described above, we also perform a test using a frequentist test-statistic, defined as

$$\text{TS}(E_{\texttt{cutoff}}) = -2 \log \left(\frac{\max_{\vec{\eta}} \mathcal{L}_{\texttt{cutoff}}(E_{\texttt{cutoff}}, \vec{\eta}) \cdot \Pi(\vec{\eta})}{\max_{\vec{\eta}} \mathcal{L}_{\texttt{SPL}}(\vec{\eta}) \cdot \Pi(\vec{\eta})} \right), \tag{8.7}$$

in order to compare to other IceCube measurements of a spectral cutoff. The null hypothesis as before is the single power-law model, and the alternative hypothesis is the spectral cutoff model with the cutoff energy as a free parameter. This model-comparison test obtains a p-value of 0.71 and a best-fit cutoff energy of 5.0 PeV. To further visualize the cutoff energy parameter space favored or disfavored, we plot the test statistic as a function of the cutoff energy in Fig. 8.12. Cutoff energies below 1.2 PeV are disfavored at more than the 68.3 % confidence level while cutoff energies above this, including the no cutoff scenario, are compatible within the 68.3 % confidence level.

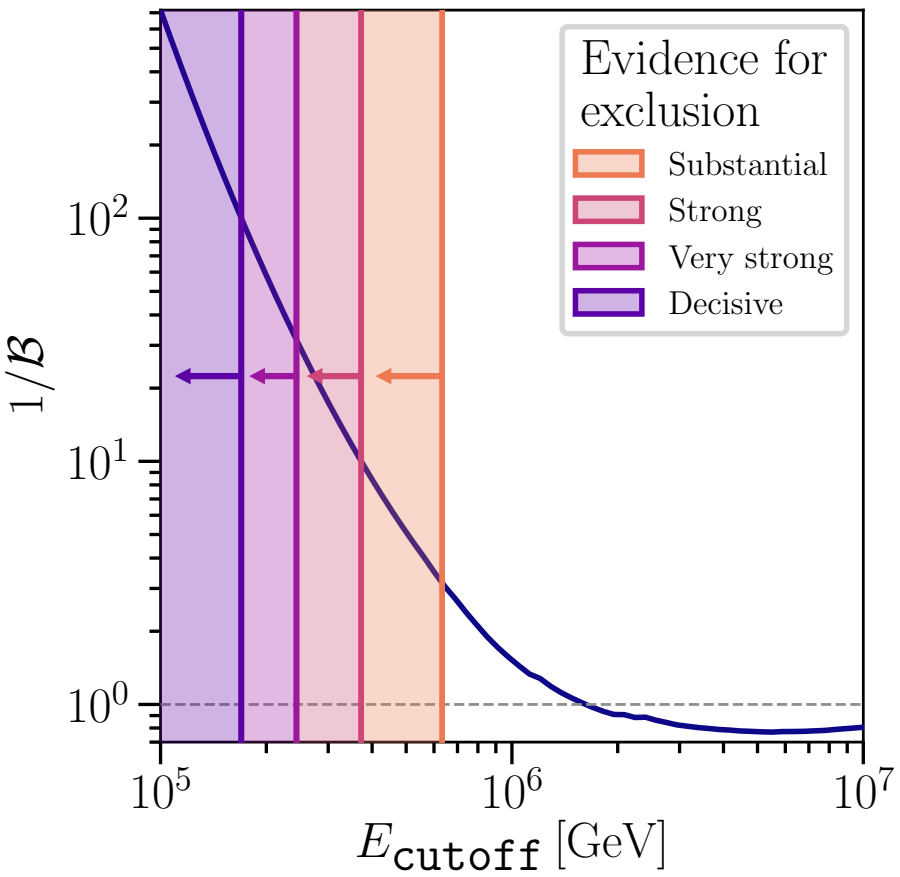

Figure 8.11: ***Spectral cutoff Bayes factors and regions of exclusion.*** The inverse of the Bayes factor, $\mathcal{B}$, is plotted in the blue curve as a function of the cutoff energy assumed in the alternative hypothesis. Regions of the cutoff energy shown by the shaded regions are disfavored with respect to the null hypothesis with varying degrees of certainty according to Jeffreys' scale. The Bayes factor is computed for each value of the cutoff energy with the single power-law as a null hypothesis. The gray dashed line indicates where the evidence of the cutoff model and the single power-law model are equal. Values above this line favor the single power-law model, while values below this line favor the cutoff model.

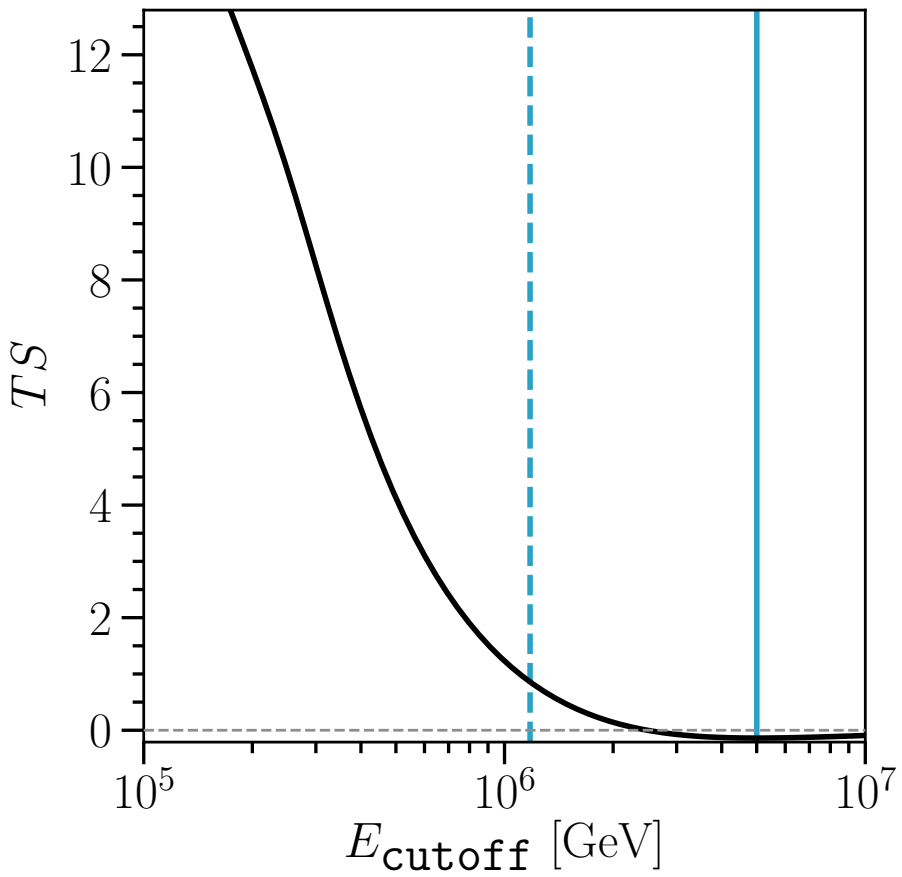

Figure 8.12: ***Cutoff model comparison using frequentist test-statistic.*** The test-statistic comparing the cutoff hypothesis and the single power-law hypothesis is plotted as the black curve. The gray dashed line indicates where the test-statistic of the cutoff model and the single power-law model are equal; there is a slight preference for the cutoff scenario in the several PeV region, although this is not statistically significant. We also report the results of parameter estimation for the cutoff model in this plot. The best-fit value for the cutoff energy, $E_{\texttt{cutoff}}$, is shown as the solid blue line and the dashed blue line indicates the boundary of the 68.3 % confidence interval.

8.1.1.4 Log-parabola flux

In log-energy log-flux space, the single power law can be represented as a line. A simple functional extension is to add curvature to this line. This gives the log-parabola model which has the form

$$\frac{d\Phi_{6\nu}}{dE} = \Phi_{\texttt{astro}} \left(\frac{E_\nu}{100\,\mathrm{TeV}} \right)^{-\left(\alpha + \beta \log_{10}\left(\frac{E_\nu}{100\,\mathrm{TeV}}\right)\right)} \cdot 10^{-18} \mathrm{GeV}^{-1} \mathrm{cm}^{-2} \mathrm{s}^{-1} \mathrm{sr}^{-1}, \tag{8.8}$$

where α is the spectral index at 100 TeV, and β governs how the effective spectral index changes with energy. In the Bayesian analysis of this model, we have chosen improper uniform priors for both α and β. At 100 TeV the most-likely spectral index ($\alpha = 2.78$) is still soft, and compatible with the most-likely SPL spectral index ($\gamma_{\texttt{astro}} = 2.89$) within the 68.3 % HPD region of α. There is one region of the parameter space where the log-parabola model becomes the same as a single power law, when $\beta = 0$. This region of parameter space is within the 68.3 % HPD region of β, informing us that the data is most compatible with a model that is close to a single power law rather than a model with larger curvature.

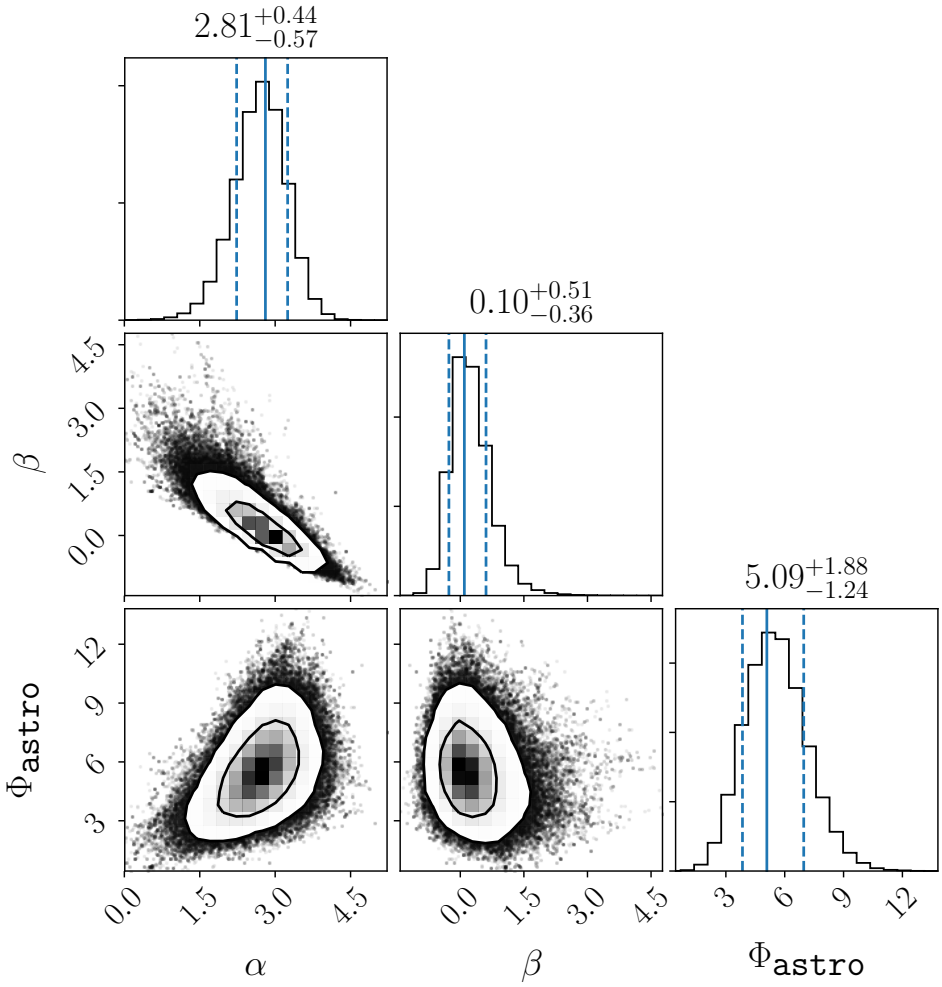

Figure 8.13: ***Log-parabola astrophysical model parameters posterior distribution.*** Results derived from the posterior distribution of the model are shown in the same style as Fig. 8.8. The figure shows the one- and two-dimensional posterior distributions for the astrophysical flux normalization, $\Phi_{\mathtt{astro}}$; the spectral index at 100 TeV, α; and the change in spectral index, β. The diagonal panels show the one-dimensional posterior with the parameter MAP estimation and 68.3 % HPD region indicated, while the non-diagonal panels show the two-dimensional posterior with 68.3 % and 95.4 % regions indicated. As we can see from the posterior distribution, β is compatible with zero which implies that an unbroken power law is a good fit to the data under these model constraints.

8.1.1.5 Segmented power-law flux

The models explored in previous sections restrict the spectrum to be described by an unbroken power-law-like model across the entire energy range. In this section, a more general parameterization of the astrophysical flux is introduced. The neutrino energy spectrum is split into segments equally spaced in $\log E_\nu$, assumed to behave as E^{-2} within each segment, and then the normalizations of each segment are allowed to vary independently. While not entirely general, this model can describe a wide variety of fluxes with the current detector energy resolution. The astrophysical neutrino flux within each segment is given by

$$\frac{d\Phi_{6\nu}}{dE} = \Phi_i \left(\frac{E_\nu}{E_{c,i}} \right)^{-2} \cdot 10^{-18} \mathrm{GeV}^{-1} \mathrm{cm}^{-2} \mathrm{s}^{-1} \mathrm{sr}^{-1}, \tag{8.9}$$

where Φ_i is the normalization constant for each bin, $E_{c,i}$ is the log-center of each bin.

This model is analyzed in the same way as previous analyses [60, 61, 62]. Namely, the best-fit point for the normalizations and their errors are obtained, which are plotted for seven energy segments (enumerated in Table 8.3) in the left panel of Fig. 8.14. Other energy segments are profiled over and not shown as they are poorly constrained by the data, and do not provide significant information. It is computationally prohibitive in such a high-dimensional parameter space to find the one-dimensional confidence intervals, which requires computing the profile likelihood at many points in the parameter space. This computational difficulty stems from the presence of many local minima. Instead, the errors are estimated by fixing all parameters except one normalization, and finding the range of this normalization for which $\Delta \max_{\vec{\eta}} \mathcal{L} \cdot \Pi \leq 0.5$. This corresponds to an approximate one-dimensional 68.3 % confidence interval, with the caveat that this procedure may underestimate the errors on each segment compared to the profile likelihood technique.

As a complement to the aforementioned frequentist approach, this model is also analyzed in a Bayesian way. Assuming improper positive uniform priors for the normalizations of the power-law segments, we sample the model's posterior distribution. The right panel of Fig. 8.14 shows the one-dimensional MAP estimation of each normalization independently for the same seven energy segments as before; the remaining energy segments have been marginalized over as the data does not significantly constrain them. Errors of each normalization are constructed by integrating the joined distribution over all other parameters and then computing the 68.3 % HPD region of that segment. These errors account for correlations between normalization segments, but the one-dimensional MAP estimators cannot be considered simultaneously as a multi-dimensional MAP. Although a multi-dimensional MAP would be a more natural complement to the frequentist results presented, such a calculation is fraught with computational challenges and so is not used here. The one-dimensional posterior density is also plotted as a turquoise band to demonstrate the shape of the distribution, although the relative scale between bands is arbitrary. Finally, in Table 8.3, the segments' normalizations are reported for both the frequentist and Bayesian analysis.

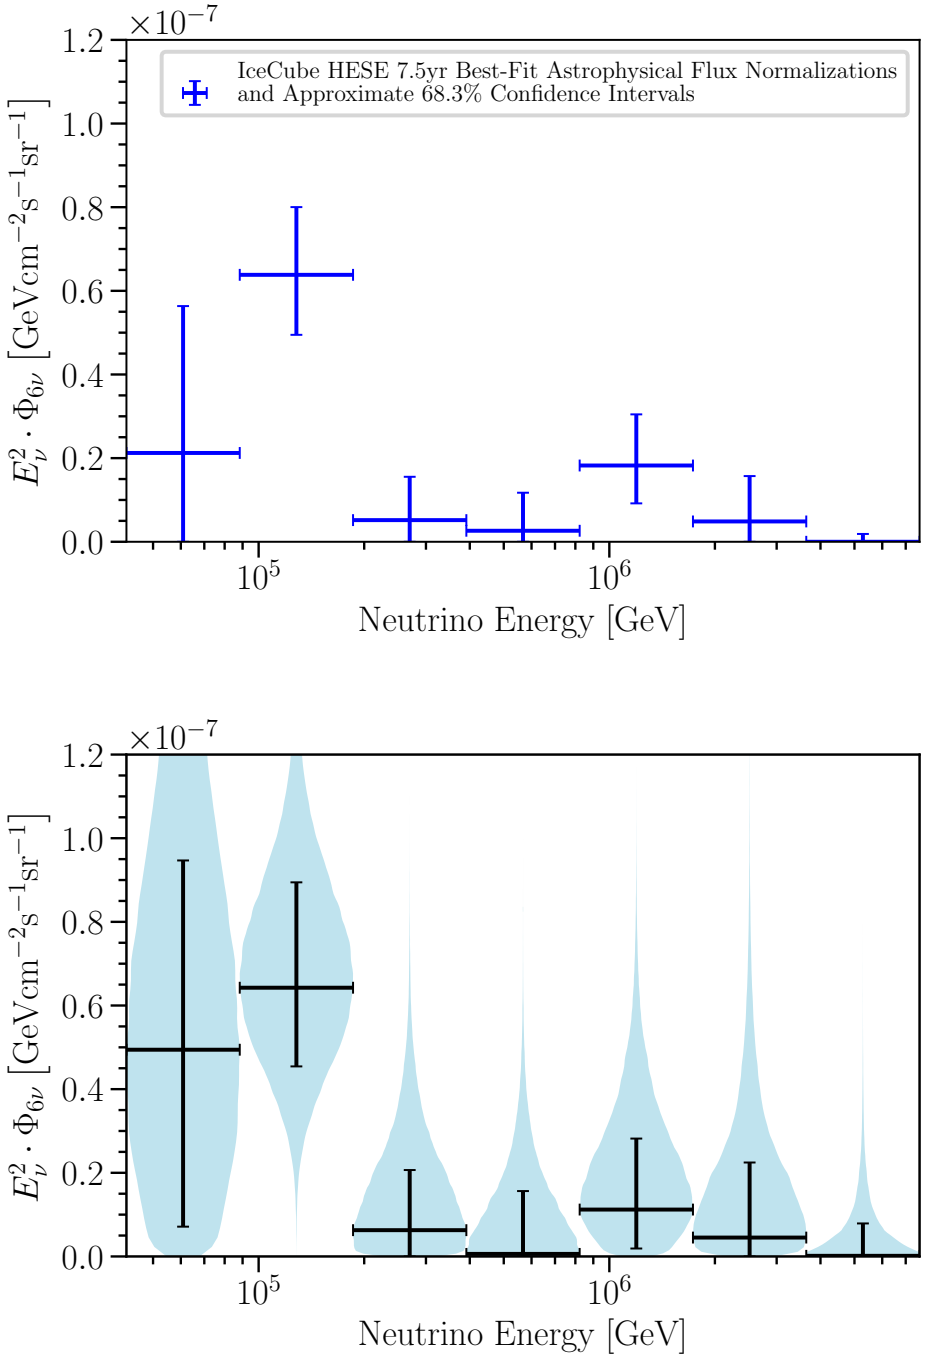

Figure 8.14: ***Segmented power-law fit.*** The differential flux as obtained by fitting the normalizations of independent E^{-2} segments defined in true neutrino energy. Left: in this plot each error bar shows the region in which $\Delta\mathcal{L} \leq 0.5$ while holding all other parameters fixed, providing an approximate 68.3 % confidence level interval for the astrophysical normalization in that particular segment. Right: in this plot each error bar shows the 68.3 % highest probability density credible interval for the astrophysical normalization in that particular segment, assuming a uniform prior on the normalizations. The width of the turquoise bands is proportional to the posterior density of the normalization in that segment, and the horizontal scales are arbitrary.

Energy Range	Frequentist Analysis	Bayesian Analysis
	Best-fit value and 68.3 % C.L.	Most-likely value and 68.3 % H.P.D.
$[4.20 \cdot 10^4, 8.83 \cdot 10^4]$	$5.7^{+9.5}_{-5.7} \times 10^{-18}$	$1.3^{+1.2}_{-1.1} \times 10^{-17}$
$[8.83 \cdot 10^4, 1.86 \cdot 10^5]$	$3.89^{+0.99}_{-0.87} \times 10^{-18}$	$3.9^{+1.5}_{-1.1} \times 10^{-18}$
$[1.86 \cdot 10^5, 3.91 \cdot 10^5]$	$7.1^{+1.4}_{-7.1} \times 10^{-20}$	$8.6^{+20}_{-8.6} \times 10^{-20}$
$[3.91 \cdot 10^5, 8.23 \cdot 10^5]$	$8.1^{+28}_{-8.1} \times 10^{-21}$	$1.9^{+47}_{-1.9} \times 10^{-21}$
$[8.23 \cdot 10^5, 1.73 \cdot 10^6]$	$1.3^{+0.86}_{-0.64} \times 10^{-20}$	$7.9^{+12}_{-6.5} \times 10^{-21}$
$[1.73 \cdot 10^6, 3.64 \cdot 10^6]$	$7.7^{+17}_{-7.7} \times 10^{-22}$	$7.2^{+28}_{-7.2} \times 10^{-22}$
$[3.64 \cdot 10^6, 7.67 \cdot 10^6]$	$0.0^{+6.7}_{-0.0} \times 10^{-23}$	$5.4^{+280}_{-5.4} \times 10^{-24}$

Table 8.3: ***Segmented power-law model normalizations.*** The left-most column shows the energy range in GeV of each segment, while the other columns show the six-neutrino flux at the center of each bin in units of $[\mathrm{GeV}^{-1}\mathrm{s}^{-1}\mathrm{sr}^{-1}\mathrm{cm}^{-2}]$ for the frequentist and Bayesian analyses. The frequentist analysis column shows the best-fit parameters and their approximate 68.3 % confidence interval. The Bayesian analysis column shows the most-likely values of the parameters, as well as the 68.3 % highest probability density interval (HPD).

The most notable feature of the segmented power-law fit, reported in this section, is the large neutrino flux at the lower energy range: between 60 TeV and 200 TeV. This lower-energy contribution to the observed flux is what drives the soft spectral index reported in Section 8.1.1.1 for the single power-law scenario. Under the assumption that the astrophysical flux sources are transparent to gamma rays, we would expect a correspondingly large gamma-ray flux to be observed. The flux measured in the second-lowest bin around $\sim$ 100 TeV is not incompatible with current gamma-ray measurements on its own, but may not be compatible depending on the underlying neutrino spectrum. For this data to remain compatible with a gamma-ray transparent source model, a break in the spectrum is needed above $\sim$ 50 TeV, such that the spectrum is harder below the break. However, the cascade sample has measured a normalization of $E_\nu^2 \Phi_{6\nu} = \sim 1.3 \times 10^{-7}\,\mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}$ around $\sim$ 30 TeV, implying that the required spectral break does not occur above $\sim$ 50 TeV. This would suggest the existence of gamma-ray opaque sources that dominate the neutrino flux at the lowest energies in this analysis [213, 214, 215, 216, 217].

Pionic gamma rays accompany high-energy neutrinos at the site of production, in fact their emission rates are intimately related by [218]

$$\frac{1}{3}\sum_\alpha E_\nu^2 Q_{\nu_\alpha}(E_\nu) \simeq \frac{K_\pi}{4}\left[E_\gamma^2 Q_\gamma(E_\gamma)\right]_{E_\gamma = 2E_\nu}, \tag{8.10}$$

where $Q_\nu(\gamma)$ is the source rate function, K_π accounts for the ratio of charged-to-neutral pion production via proton-gas interactions (pp) or photo-hadronic interactions ($p\gamma$). However, these pionic gamma rays interact with the extragalactic background light (EBL) and cascade to lower energies [219], contributing to the isotropic gamma-ray background (IGRB). The high intensity of the neutrino flux at energies below 100 TeV compared to the IGRB flux measured by the Fermi satellite may indicate that sources responsible for their production cannot be transparent to very-high-energy gamma rays. In this scenario, the gamma

rays produced interact with lower energy photons in the source, initiating electro-magnetic cascades that could be visible at lower frequencies, and therefore, would not overshoot the measured IGRB flux. We should also note that photo-hadronic interactions should be the dominant channel of neutrino production in sources responsible for the high-intensity flux, as hadro-nuclear interactions would result in an overwhelming gamma-ray flux at lower energies that cannot be tolerated by the measured IGRB [220, 215, 217].

8.1.2 Atmospheric flux from charmed hadrons

Although most neutrinos in cosmic-ray air showers are produced from the decay of muons, pions, and kaons, at high energies, the decay of charmed hadrons can produce neutrinos as well. Due to the short decay length of charmed hadrons compared to their interaction length, they mostly decay without losing energy, yielding a harder spectrum of neutrinos than the conventional component [206], that is more uniform across the sky. Two ingredients are needed to compute the flux from charmed hadrons: the production cross section of charmed hadrons and the cosmic-ray flux. For this work, the relevant part of the production cross section has not been measured at collider experiments as it is only accessible in the far-forward region very close to the beam-line where there is little to no instrumentation. The charmed hadron production cross section can be computed by means of perturbative QCD [221, 222, 223, 224, 225] and non-perturbative techniques such as dipole-model interactions [226, 227, 43]. In the region of interest, from 10 TeV to 10 PeV, the expected flux has an uncertainty of at least a factor of two in normalization due to renormalization and factorization scale uncertainties, as well as uncertainties in the cosmic-ray composition, and charm mass uncertainties according to [228, 225, 229]. Although these uncertainty estimations do not include the possibility of additional non-perturbative contributions, *e.g.* intrinsic charm [230]. The prompt flux shape variation in this region of interest arises primarily from changes in the cosmic-ray models. In this analysis, we use the BERSS calculation [117] with passing fractions from [65] to predict the baseline prompt contribution to the data sample.

This sample is not expected to be sensitive to the prompt neutrino flux as the contribution to the overall flux is small in the sample's energy range, and much of the flux is suppressed in the downgoing region. An analysis of the HESE sample can be performed considering only the atmospheric muon and neutrino components. This results in a best-fit prompt normalization of 21.56 times the baseline model and is shown in Fig. 8.15. As can be seen from Fig. 8.15, the predicted angular distribution in this background-only fit fails to explain the Southern-sky event rate. The atmospheric only hypothesis is disfavored by approximately 8σ in comparison to the single power-law astrophysical plus atmospheric flux hypothesis.

Compared to constraints on the prompt normalization from other IceCube samples, the best-fit prompt

	Frequentist upper limit (90 % C.L.)	Bayesian model rejection (strong)
Northern sky muons IC59 [231]	$3.80 \times \phi_{ERS}$	–
Northern sky muons IC86 [58]	$1.06 \times \phi_{ERS}$	–
All-sky medium-energy starting cascades [63]	$1.52 \times \phi_{ERS}$	–
HESE 7.5 years (this work)	$9.82 \times \phi_{BERSS}$	$13.29 \times \phi_{BERSS}$

Table 8.4: ***Summary of constraints on the flux of charmed mesons.*** The benchmark ϕ_{ERS} [227] and ϕ_{BERSS} [117], are such that $\phi_{ERS}(100\,\text{TeV}) \approx 2.5 \cdot \phi_{BERSS}(100\,\text{TeV})$.

normalization obtained in the background-only fit is in tension with these results [231, 58, 63]. Some of these constraints have been obtained considering a single-power-law astrophysical component, and are thus dependent on this model assumption. However, the constraints from [231] predate the observation of high-energy extraterrestrial neutrinos and are conservative because this scenario is equivalent to zero contribution from the astrophysical flux. The latter results in a constraint of the prompt normalization of 3.80 times the ERS calculation [227] at 90 % C.L.; a model which is approximately 2.5 times larger than the benchmark model used in this analysis.

When we allow the existence of a single power-law astrophysical component, the best-fit prompt component normalization is zero. In this same scenario, using the frequentist statistical construction assuming Wilks' theorem with one degree of freedom, a 68.3 % C.L. prompt normalization upper bound of 5.34 and a 90 % upper limit of 9.82 is obtained. This result is in agreement with the results summarized in Table 8.4.

Additionally, in the Bayesian framework, the most-likely value of the prompt normalization is $0.54^{+5.62}_{-0.54}$ when assuming an improper uniform prior for the prompt normalization. In this case, the prompt normalization posterior distribution strongly depends on the prior choice. For this reason, we report our Bayesian prompt normalization results in terms of the Bayes factor between the no-prompt hypothesis and a given prompt normalization; see Section 7.1 for details. In Fig. 8.16 we show the Bayes factor obtained assuming a uniform prior on the astrophysical neutrino normalization and spectral index. We find that prompt normalizations greater than 13.29 are disfavored at the strong level, according to Jeffreys' scale, compared to the no-prompt scenario.

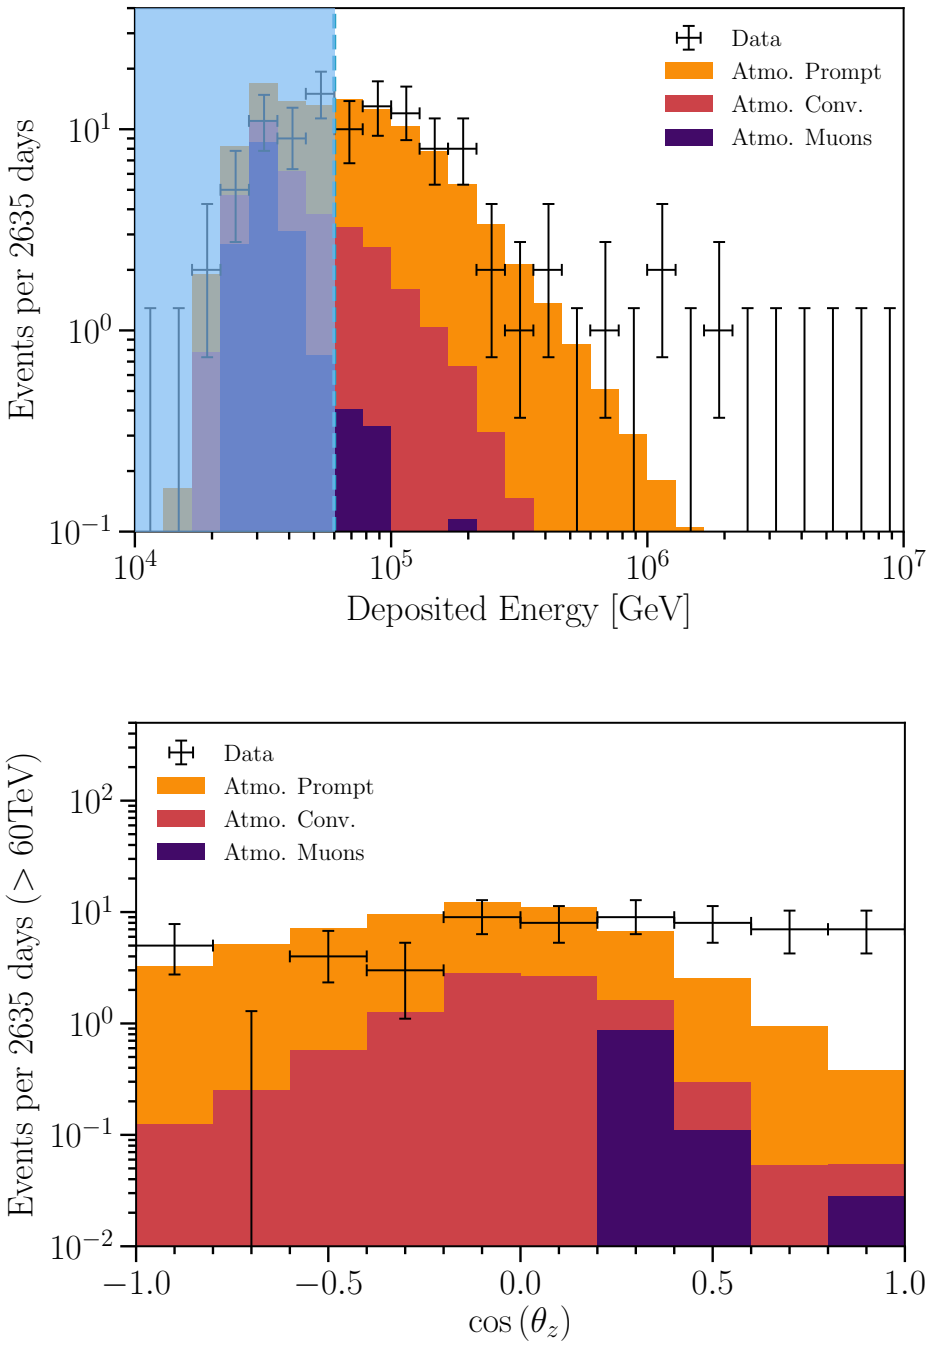

Figure 8.15: ***Atmospheric-background-only fit to the data.*** In these figures we present the best fit in the absence of an astrophysical component. The left panel shows the deposited energy distribution and the right panel the angular distribution. As can be seen in the right panel, the angular distribution is in tension with the expectation in several bins. This amounts to an approximately 8σ difference with respect to the best-fit astrophysical model. The colors are the same as in Fig. 8.1.

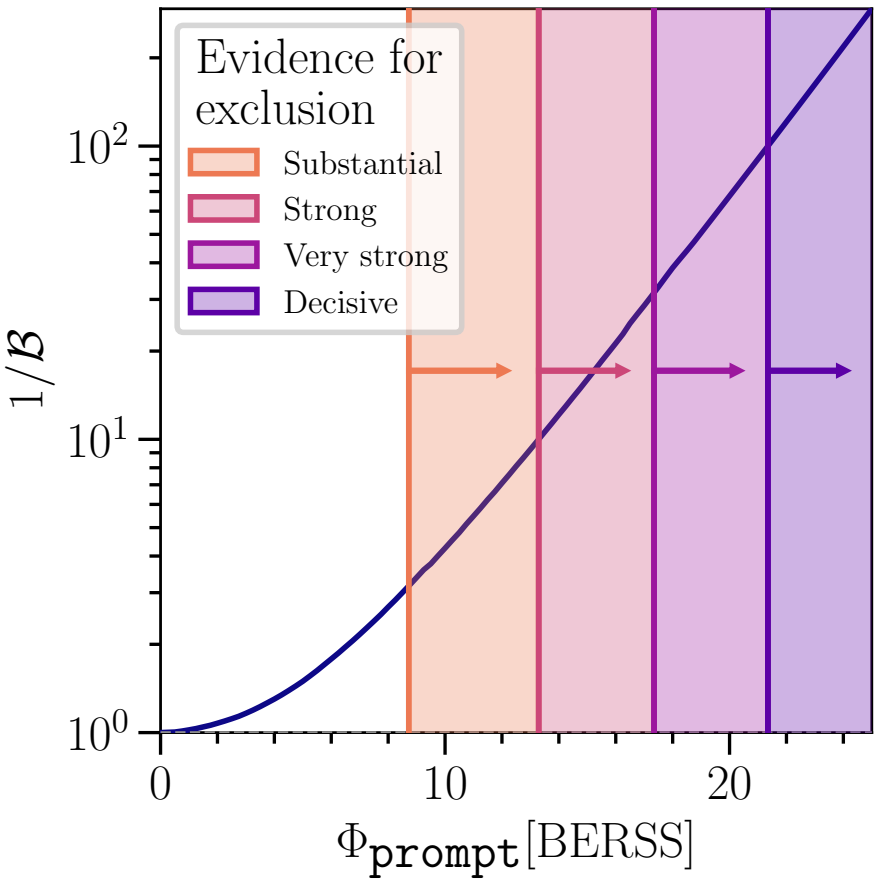

Figure 8.16: ***Prompt neutrino flux normalization constraints.*** The horizontal axis shows the size of the prompt normalization with respect to the baseline BERSS model discussed in Chapter 6 with passing fractions given in [65]. The vertical axis gives the reciprocal of the Bayes factor; decreasing Bayes factors imply a more disfavored prompt normalization. The shaded regions denote exclusions of the prompt normalization according to Jeffreys' scale from substantial to decisive.

8.1.3 Source-specific models

Section 8.1.1 characterizes the observed astrophysical neutrino events using generic models. These studies show that a single power law is a good fit to the data. Nevertheless, in this section, we study the compatibility of the observed events with specific source predictions of the astrophysical neutrino flux proposed in the literature. The specific source fluxes used in this analysis, together with the result of the segmented power-law fit, can be seen in Fig. 8.18. These models were chosen because they have a significant flux contribution in the energy range that this analysis is sensitive to. Thus, *e.g.* we do not test cosmogenic [232] neutrino flux models, which predict neutrinos from cosmic rays interacting with the cosmic microwave background since they are expected to contribute at higher energies where dedicated IceCube searches exist [233]; see [234] for

a recent discussion on the expected rate of cosmogenic flux in this analysis energy range.

Astrophysical neutrino flux predictions have in principle many parameters that may modify the expected flux. For simplicity, this analysis does not study the models' internal parameters but is limited to some nominal values of their parameters, fixing the source model to a single flux. This analysis means that the tests do not cover these models in their entirety, and so the results presented here should not be interpreted as categorically excluding any of the models tested. The analysis takes the form of a model selection test with two non-nested alternatives, which we assume to be equally likely in our interpretation of the results.

In the primary analysis, the two alternatives are a single power law and a specific source model. The Bayes factor of these two scenarios is used as a criterion for model selection; see Section 7.1 for details. Thus, the single power law serves as a benchmark model to compare against. In the case of the single power-law model the evidence is calculated by marginalizing over the two model parameters, normalization and spectral index, assuming uniform priors in a compact region defined as $(\Phi_{\texttt{astro}}, \gamma_{\texttt{astro}}) \in [0, 25] \times [2, 4]$, as well as the analysis nuisance parameters with priors given in Table 7.1. Boundaries of this uniform box prior are chosen to encompass recent measurements of the astrophysical flux and the bulk of the posterior distribution mass. To first order, expansion of this box prior changes the evidence by a factor proportional to the parameter space's size. The alternative scenario has no free parameters, and thus the posterior integral is only over the nuisance parameters. Figure 8.18 shows the specific source fluxes with a color scale that orders scenarios by their evidence. In this study, the single power law is penalized due to additional model complexity with respect to the specific source scenario.

Since some of these models are not intended to explain the whole astrophysical neutrino spectrum, a secondary analysis is also performed. In this analysis, two models are considered: a single power law on its own and a single power law together with a source model. In this case, when comparing the null and alternative hypotheses, the constant parameter-space factor that results from any choice of the single power-law flux parameter boundaries will cancel. So we are able to use improper uniform priors without introducing an arbitrary scaling factor. The same prior dependence in the other Bayesian analyses remains here, as the two models' likelihoods can peak in different regions of parameter space.

The results of both analyses are shown in Table 8.5. For each model, the "Model only Bayes factor" is reported as the result of the primary analysis, and for the secondary analysis, the "Model plus single power-law Bayes factor" is reported together with the most-likely spectral index and normalization and their errors.

Given the obtained Bayes factors models can be organized into two categories:

- Models with Bayes factors much less than one for both analyses. In this case, the single power law is a

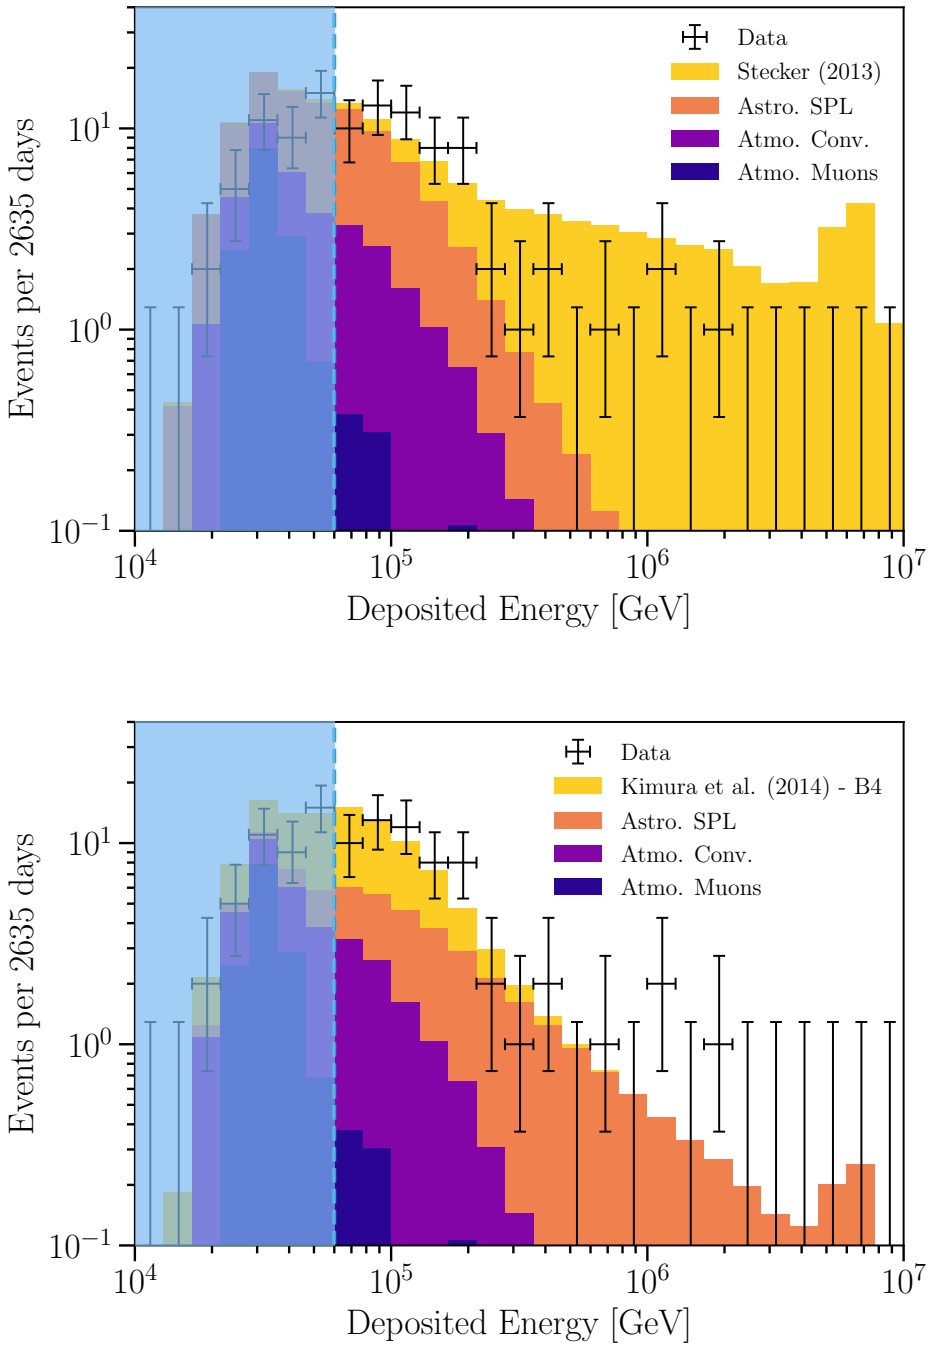

Figure 8.17: ***Source-specific model energy distributions.*** Different panels show the predicted energy distribution compare to the data for a subset of the models considered in Table 8.5. In the cases shown above there are two astrophysical components: a specific source model and an additional single power-law component. All the components are shown as a stacked histogram at the best-fit value of the components normalizations. Left: The Stecker [235] model is shown as an example of a case where the Bayes factor assuming equally likely alternatives indicates significant preference for the single power-law model over the two component scenario. Right: The Kimura et al. (B4) [237] model is shown as an example of a case where an additional single power law is needed to explain the distribution of events.

Source class	Model	Model only Bayes factor	Model + SPL Bayes factor	Most-likely SPL $\gamma_{\texttt{astro}}$	Most-likely SPL $\Phi_{\texttt{astro}}$
AGN core	Stecker [235]	7.26×10^{-13}	8.27×10^{-11}	$3.9^{+0.58}_{-0.45}$	$4.24^{+1.59}_{-1.29}$
AGN	Fang et al. [236]	0.379	0.153	$4.07^{+0.64}_{-0.69}$	$2.1^{+1.5}_{-1.16}$
LLAGN	Kimura et al. (B1) [237]	7.12×10^{-6}	5.20×10^{-7}	$4.74^{+0.26}_{-0.86}$	$1.03^{+0.89}_{-1.03}$
LLAGN	Kimura et al. (B4) [237]	3.74×10^{-4}	0.219	$2.43^{+0.3}_{-0.27}$	$1.57^{+0.98}_{-0.95}$
LLAGN	Kimura et al. (two component) [237]	1.87×10^{-4}	3.63×10^{-6}	$4.25^{+0.75}_{-0.86}$	$0.0^{+0.62}_{-0}$
BLLac	Padovani et al. [238]	1.08×10^{-10}	1.75×10^{-7}	$3.66^{+0.51}_{-0.4}$	$4.92^{+1.64}_{-1.44}$
GRB choked jet	Senno et al. [239]	0.353	2.31	$3.7^{+0.54}_{-0.64}$	$3.16^{+1.81}_{-1.19}$
SBG	Bartos et al. [240]	1.16×10^{-14}	1.62×10^{-16}	$4.43^{+0.56}_{-1.05}$	$0.0^{+0.51}_{-0}$
LLBLLac	Tavecchio et al. [241]	0.104	0.703	$3.79^{+0.74}_{-0.43}$	$3.63^{+1.41}_{-1.47}$
GRB	Biehl et al. [242]	1.31×10^{-6}	0.145	$3.32^{+0.43}_{-0.35}$	$5.36^{+1.72}_{-1.39}$

Table 8.5: ***Astrophysical neutrino flux model comparison test results.*** Each row shows the source-specific scenario tested, the Bayes factor of the model on its own, the Bayes factor of the model in conjunction with a power-law component, the most likely spectral index of the accompanying power-law component with corresponding 68.3 % HPD region, and the most likely normalization of the accompanying power-law component with corresponding 68.3 % HPD region. Assuming the two alternatives are equally likely, we can interpret small Bayes factors (those less than one) as indicating a preference for the single power-law model and large Bayes factors (those greater than one) as indicating a preference for the alternative scenario.

better description of the data, and the addition of a single power law to the source model does not alter this conclusion. Figure 8.17 (left) provides an example of this category.

- Models with a Bayes factor much less than one compared to the single power law, but that improve when introducing an additional single power-law component. Models in this category can only describe part of the flux and would require the existence of a second component to be compatible with the data. Figure 8.17 (right) provides an example of this category.

To conclude, in this section, we have studied models of astrophysical neutrinos proposed in the literature with nominal parameters and compare them to our baseline parameterization, the single power-law spectrum. No tested scenario with nominal parameters is substantially preferred over the baseline model, and some scenarios – those with Bayes factors much smaller than unity – are disfavored; see Table 8.5. More extensive tests of models in the literature using the information provided in [202] are encouraged.

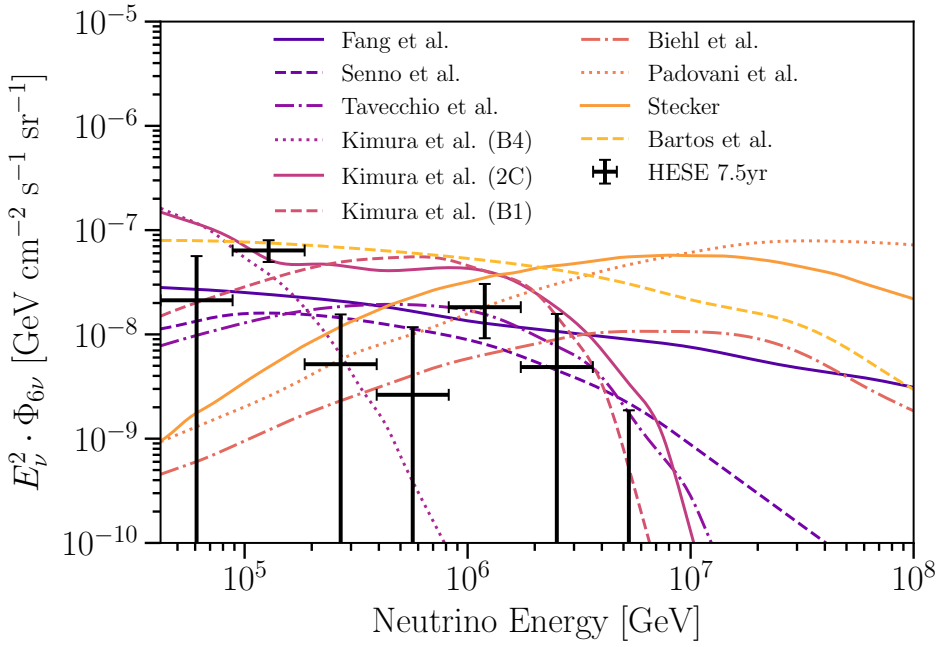

Figure 8.18: ***Source-specific models tested in this work together with the segmented-fit outcome.*** Models discussed in Section 8.1.3 and listed in Table 8.5 are shown as lines. The HESE segmented power-law model, described in Section 8.1.1.5, best-fit normalizations are shown as black crosses. Models are ordered in the color scale from largest (darkest color) to smallest (lightest color) evidence as reported in Table 8.5. Above 10 PeV no upper limits are shown as simulation is not available above this energy.

Chapter 9

Conclusions and next steps

9.1 Analysis conclusions

This work's immediate result is an updated analysis of the high-energy starting event (HESE) sample, which helped discover astrophysical neutrinos. This updated analysis improves the description of atmospheric backgrounds (Chapter 6) and incorporates an additional 4.5 years of data taking. For this analysis, the largest change to the atmospheric background description is the update of the atmospheric passing fractions to use the calculation from [65] and the change in the baseline neutrino flux from charmed hadrons to the flux calculated in [117]. The detector uncertainties and descriptions have also been improved. The detector single-photo-electron distributions have been re-calibrated, resulting in a $\sim 4\,\%$ change in the inferred charge on average. Uncertainties on the absolute detector efficiency and angular acceptance are now included. Finally, reconstruction uncertainties are now accounted for in a more detailed manner.

The atmospheric-only scenario requires a prompt neutrino normalization $\sim$ 20 times larger than the baseline model (Section 8.1.2). Which is in disagreement with previous prompt normalization upper limits from complimentary IceCube analyses; the strongest constraining the normalization to be less than ~ 2.65 times the BERSS flux at 100 TeV [58] at the 90 % C.L, and the weakest constraining the normalization to be less than ~ 9.5 times the BERSS flux at 100 TeV [231] at the 90 % C.L. which is conservative because of the zero observed astrophysical flux. The atmospheric-only solution is disfavored with a decisive criterion according to Jeffreys' scale, using only this sample, when compared to a model that incorporates a single power-law astrophysical component due to the inability of an atmospheric-only model to reproduce the observed zenith distribution. These results reinforce the conclusions of previous analyses [59, 60, 61] regarding the astrophysical neutrino flux. Although the prompt component has a distinct angular signature in the HESE sample, the component's normalization is far too small for this analysis to be sensitive without orders

of magnitude more data.

The observed data is compared to generic astrophysical neutrino model assumptions using frequentist and Bayesian statistical prescriptions (Section 7.1). We report the preferred parameter space for a simple single power-law model $2.87^{+0.20}_{-0.19}$ (Section 8.1.1.1) and find this is in agreement with previous results after accounting for the differences in analysis design and the possibility of statistical fluctuations. The shift in the spectral index comes partially from the extension of the analysis energy range, and as a result of many additional cascade events in the low-energy region of the analysis observed in the latter 4.5 years. The significance and uncertainties of this astrophysical measurement remain robust in the face of additional systematics. Results from other IceCube samples differ in their best-fit parameters, but overlap in their 95.4 % C.L. regions.

Comparisons of the data to generic astrophysical neutrino models have also been presented, such as, a double power law (Section 8.1.1.2), a spectral cutoff (Section 8.1.1.3), and log-parabola (Section 8.1.1.4). No preference is found for any of these models compared to the single power law, in particular, the existence of a cutoff is constrained to be at energies greater than 370 TeV with a strong criterion according to Jeffreys' scale and similarly excluded below 670 TeV at the 90 % C.L.

As a new part of this analysis, the preference for some specific source scenarios proposed in the literature [235, 236, 237, 238, 239, 240, 241, 242] is quantified. These tests are performed assuming the models' nominal parameters and compared to the single power law using the Bayes factor. Some models provide a compatible description of the data, while others fail to explain the observations even with the addition of a power-law component. No scenario tested provides a substantially improved description of the data compared to the single power-law, which describes the data well. Although, models that do not describe the data well with nominal parameters may do so with other model parameters, and other models not considered may provide better descriptions of the data. On this point, readers are encouraged to use the data release [202] to perform more detailed tests.

Other measurements have been performed with this sample in a manner consistent with the analysis techniques presented here. These are the measurement of the neutrino flavor composition [108], searches for additional neutrino interactions [159] and dark matter in the galactic core [160], and a measurement of the neutrino cross-section [161].

We conclude that, given the available data in this sample's sensitive energy range, the astrophysical neutrino flux is well described by a single power law and there is no evidence for additional spectral structure in this sample. Despite this, many models remain compatible with the data, and larger samples will be required to differentiate between the different proposed spectra. Uncertainty of the spectral index is reduced after adding 4.5 years of data, however, it has also shifted to a softer spectrum away from other measurements.

These other measurements also have reduced uncertainties with the addition of more data. Although a purely statistical explanation cannot be ruled out, differences between these spectral measurements remain unexplained and will require further study to resolve.

9.2 The next generation

The analysis presented here updates previous studies of the HESE sample and is only a small part of the picture for astrophysical neutrinos and the mystery of cosmic ray origin. However, the techniques and solutions developed here remain broadly applicable to future analyses and will enable more precise measurements as we accumulate more data, and develop additional samples. This section explores what the next generation of analyses may look like and the possible avenues to explore beyond the work presented here.

The next objectives for analyzing the astrophysical neutrino flux are to expand the energy range and flavor information of the analysis and simultaneously examine these different channels for additional structure or differences between them. Some analyses have already examined these channels, for instance, samples that look at muon neutrinos from the Northern hemisphere between 100 GeV and 10 PeV or cascade focused samples that extend down to ~ 1 TeV. However, the separate analyses of these samples have either lacked extended systematics treatments, lacked enough data to make significant conclusions, or are too disparate for meaningful comparison between each other. Further investigation of the astrophysical neutrino flux will require the simultaneous analysis of all these data samples and expand upon the techniques developed for this work. Many of these samples were developed for the initial measurement of the astrophysical neutrino. However, as more data becomes available, the precision of these analyses will no longer rely on maximizing the selection's acceptance. Instead, we will be able to focus on the analysis accuracy in an attempt to understand the differences we have observed between the various astrophysical neutrino measurement channels, even if this comes at the cost of some precision.

9.2.1 Data samples

The cascade data channel can achieve low background rates in all directions because it is unlikely for muons to produce cascade-like detector signatures, even at lower energies. The HESE selection as a cascade-dominated sample takes advantage of this in part but cannot cover lower energies. Additional selection techniques can obtain a high purity sample of cascade-like neutrino events that extends down to 1 TeV. Two IceCube event selections have achieved this, namely, the "medium-energy starting event selection" (MESE) and the "cascade selection". These cascade dominated selections are largely comprised of NC, ν_e CC, and a sub-population of ν_τ CC events. In the MESE selection, some ν_μ CC events are also present in the sample. These samples

cover the entire sky, enabling the combined ν_e and ν_τ flux to be accurately measured in both hemispheres, although finer grained angular studies are limited by poor directional resolution. A major advantage of these samples is the higher energy resolution of cascade events, enabling more detailed studies of the neutrino energy spectrum. A sample of partially contained cascades can be included to extend the analysis to even higher energies since a partially contained selection has a much larger effective area. Although, the energy resolution for these partially contained events will not as good as contained events.

Beyond cascades, we hope to incorporate similar information from ν_μ CC interactions. This is available through the different track dominated event selections. Two IceCube selections use a similar approach to obtain a sample of high-quality through-going muon neutrino tracks, namely by restricting selected events to tracks known to be up-going. These are the selection used for the Northern hemisphere astrophysical neutrino flux measurement, and the selection used in the search for eV scale sterile neutrinos. Between these two selections, the primary difference lies in the method of selection cuts; the astrophysical measurement sample uses a boosted decision tree, while the sterile search sample uses only hand tunes straight cuts. This choice stems from a difference in motivation for the samples. The astrophysical sample aims to maximize effective area and angular range, while the sterile search sample aims for the high-quality reconstruction and maximum data-simulation agreement. These samples only cover the Northern hemisphere but enable constraints on the energy and angular distribution of muon neutrinos within that range. One disadvantage of these samples is their limited energy resolution that stems from the challenges of inferring muon energies. Even assuming an isotropic 1:1:1 astrophysical flux, there are benefits to including one of these samples in the analysis. These samples extend between 100 GeV and 10 PeV, and below a reconstructed energy of ~ 20 TeV atmospheric neutrinos dominate the flux. This large population of atmospheric neutrinos can help to pin down the unknown values of the atmospheric systematic parameters, thereby reducing uncertainty in the astrophysical measurements even for other samples.

To examine ν_μ CC events in the Southern hemisphere at lower energies (< 100 TeV) a severe suppression of the overwhelming muon background is required. The enhanced starting track event selection (ESTES) fills this role by using a dynamic veto that takes advantage of the reconstructed track direction to compute the probability of a muon producing a similar signature. Even neglecting the muon background, a through-going track sample will inevitably be overwhelmed by atmospheric neutrino backgrounds below 200 TeV. Only a sample that actively vetos atmospheric events can have sensitivity to the astrophysical flux below 200 TeV. In addition to the high-purity Southern hemisphere track information, the selection also has better energy resolution than other track samples. This improved energy resolution gives an analysis sensitivity to smaller-scale features in the energy spectrum that the Northern sky track samples cannot discern.

In addition to the event selections described above, many identifiers for tau neutrino events have been

developed to find tau neutrino events in low-level data. However, there are no known properties of ν_τ CC events that allow them to be readily separated from muon backgrounds and other types of neutrino interactions, leading to only a handful of candidate tau events. Even without clear cut tau events, these identifiers can still differentiate ν_τ CC events from others on a statistical basis if the muon background is first reduced. Thus, the most effective approach is to apply these identifiers to existing selections and statistically infer the tau neutrino flux properties, rather than trying to separate the events outright. Such identifiers use different methodologies, but all hinge on the properties of a particular class of ν_τ CC events. There are two physically distinct particle showers in this class of events, one from the initial neutrino DIS and the other from the tau decay. By combining these event selections and identifiers, an analysis can test for structure and discrepancies in the different channels while also obtaining more precise measurements of the astrophysical flux that are less limited by systematic uncertainties.

9.2.1.1 Cascade selection

The cascade selection focuses on achieving a large effective area for cascade-like neutrino events while reducing muon backgrounds. This is achieved through a combination of cascade identification, veto-like cuts, quality cuts, containment cuts, and machine learning based classification. In an analysis of this high purity selection of cascade-like neutrino events the combined astrophysical flux of electron neutrinos and tau neutrinos can be precisely measured to lower energies. The only confounding factors in this measurement are the proportion of neutral-current muon neutrino events, which can be constrained with other samples, and systematic variation to the contribution of atmospheric neutrinos to the sample. Unfortunately we do not have the computational capacity or theoretical framework to accurately model muons from atmospheric showers in this type of event selection, so a number of cuts are made to remove sources of data-simulation disagreement that may arise from muons in the sample. Similarly to the calculation techniques used for the analysis in this work, the cascade selection analysis computes the effect of accompanying muons with the older passing fraction calculation described in [131]. In this case the calculation of the passing fractions relies only on a single parameter, the median muon energy that triggers the detector. The value of this parameter is chosen as a best-guess to be 100 GeV, which is verified by comparison to dedicated cosmic ray air-shower simulations. While this treatment is sufficient for the current analysis of the cascade sample, more accurate calculations of the atmospheric suppression are desired for a global analysis so that possible differences between the different observation channels can be examined.

Absent the ability to compartmentalize the detector response to muons and neutrinos, the estimation of atmospheric backgrounds must rely on other methods than the calculation presented in Section 6.1.2 or a more complicated extension of this. Other available methods of modeling the interplay between

related atmospheric muons and neutrinos and their combined effect on the event selection rely on air-shower simulations. Air-shower simulations to this end have several downsides: they tend to be computationally expensive by comparison to signal simulation, must often be specialized for a particular event selection, are only available for a small subset of hadronic interaction models, and must be rerun for any hadronic interaction model changes. For this reason it may prove difficult to obtain more accurate atmospheric neutrino predictions and add additional atmospheric systematic uncertainties to an analysis of the cascade sample. An easier first step may be to use the less sensitive medium-energy starting event sample, for which accurate background estimations are more straightforward to obtain.

9.2.1.2 Medium-energy starting events

Much like the cascade selection, the medium-energy starting event selection also produces a cascade dominated sample that extends down to 1 TeV. However, MESE achieves this by modifying the veto definitions of the HESE selection and adding additional cuts so that the charge threshold may be lowered. The foremost cut is a charge dependent fiducial volume cut, where the fiducial volume size is reduced for lower charge events. This smaller fiducial volume removes most of the lower energy muons, which are more likely to have a reconstructed vertex near the detector edge than in the center.

The MESE sample has similar properties to high-energy starting event selection in that only events with a contained interaction vertex are selected; it also retains the physical separation between veto and fiducial volume. A similar set of calculations can be applied to estimate the MESE sample's backgrounds, and the selection is susceptible to the same systematics. Primarily the addition of this selection allows us to explore the lower energies of the astrophysical spectrum between 10 TeV and 100 TeV. The lower energy range comes with its own challenges, though, these are primarily a larger data sample and higher atmospheric backgrounds. To ensure accurate measurements with this increased sample size and background, we must ensure a higher degree of accuracy in the background estimates and potentially allow additional systematic variations. In expanding our systematics considerations, the next logical step is the set of uncertainties related to cosmic rays and their production of neutrinos. Thankfully, MCEq allows arbitrary cosmic ray model fluxes and compositions as input and can select from a range of hadronic models. With these tools, we can directly evaluate the effect of these systematics on the atmospheric neutrino flux. At the same time, the improved passing fraction calculation leverages this same flexibility so that these systematics may be taken into account for accompanying muons as well.

9.2.1.3 Northern sky through-going tracks and eV scale sterile neutrino searches

The Northern sky through-going track selection is used to measure the astrophysical neutrino flux chooses well-reconstructed up-going track-like events; a boosted decision tree differentiates between neutrinos and misreconstructed muons near the horizon. The boosted decision tree identification allows other quality cuts to be loosened, increasing the selection's effective area. The larger effective area is optimal for an initial measurement of the astrophysical neutrino flux as much of the sensitivity comes from the high-energy tail where there are few events overall. One difficulty with analyzing this type of sample is that the astrophysical contribution is buried beneath the atmospheric flux at lower energies. Combined with the limited energy resolution of tracks, this makes it difficult to discern features in the energy spectrum beyond power-law-like behavior, especially below 100 TeV. To make an accurate measurement of the astrophysical spectrum with this sample at lower energies, good data-simulation agreement and a detailed systematic treatment of the atmospheric background are necessary. A reasonable way to achieve this is through more stringent quality cuts, which will reduce the effective area but potentially remove events that are more poorly modeled by simulation. Similar steps have been taken in the sample used for eV scale sterile neutrino searches, but at lower energies.

Another similar sample, which does not have a particular name, is used to search for eV scale sterile neutrinos. This particular search is a shape analysis that relies on percent-level differences in the atmospheric data compared to the standard model expectation. Additionally, it is necessary to account for many potential confounding systematics. As a result, the sample has sub-percent-level agreement between data and simulation, and a detailed treatment of the relevant systematics. This precision is achieved by strict quality cuts for the event reconstructions and limiting the zenith range to avoid the horizon. These same features are desirable in a global analysis of the astrophysical neutrino flux, so it may be easier to use this sample instead of the track sample currently used for astrophysical measurements. The main caveat is that these optimizations made for lower energies may not necessarily hold at higher energies where the sample is sensitive to the astrophysical flux.

9.2.1.4 Enhanced starting track event selection

The two muon neutrino samples described above cover only the Northern hemisphere, so we would like to include a track sample with sensitivity in the Southern sky. The enhanced starting track event selection (ESTES) works by taking the reconstructed event directions, computing the probability that a muon could sneak through the detector region behind the event vertex, and rejecting events where this probability is high. This method is similar to veto techniques as a sufficiently detailed veto definition that depends on the

event's properties could produce the same event selection. With ESTES, the main challenge for background estimation is accounting for the interplay between muon and neutrino event properties and how they may conspire to accept or reject the event. A sufficiently detailed P_{light} that depends on both muon and neutrino properties can account for this and allow us to use the calculation in Section 6.1.2 for background estimation. Although ESTES's effective area is smaller than other selections when all directions are considered, it is unique in its high astrophysical purity for down-going TeV track events.

In addition to adding track information in the Northern sky, ESTES also has better energy resolution than other track samples. For the starting tracks that ESTES selects the initial hadronic cascade is observed, this additional information constrains the possible neutrino energies more than through-going track observations. The improved energy resolution will allow an analysis ESTES to be more sensitive to features in the energy spectrum of astrophysical neutrinos.

9.2.2 Analysis improvements

Following Fig. 1.2, we can examine what components of the analysis may be improved. In the previous subsection, we have already added additional samples, improving the "event selection". Event reconstruction can be improved, but this topic has not been the focus of this work, so we will ignore it in this discussion. The other interesting aspects that can be improved lie in the MC generation, likelihood function, and reweighting; in fact, the three are intimately related. Signal estimation is relatively optimal at this stage of IceCube's operation apart from some unmodelled detector and ice systematics. On the other hand, background estimation remains challenging and is susceptible to a wider array of systematic effects. Background estimation can be broken down into two categories, the direct contribution of muons to the event selection and the effect of related muons on the acceptance of atmospheric neutrinos.

9.2.2.1 Atmospheric muons

For muons that directly contribute to observations, the main problem is the large number of muons in proportion to neutrinos. We expect the detection of approximately 10^6 more muons than neutrinos, based purely on the interaction cross section. Assuming the simulation of a single muon or neutrino event incurs a similar computational cost, the simulation of background muons is vastly more expensive than the simulation of neutrinos for equivalent periods of time. The event selections must also be extremely efficient at rejecting muons to measure the neutrino flux, so for all the expensive background estimation we can muster, there is often very little information regarding the distribution of muon backgrounds in the sample. This shortcoming is exemplified in Fig. 6.19, which shows the expected distribution of muons in the HESE selection. Most

of the bins in this figure are empty, and bins with simulated events only contain a handful. This situation indicates inadequate simulation, and the aforementioned plot is a projection that reduces the number of bins by approximately a factor of three. Clearly, this situation can be improved. The two main ways to improve upon this without orders of magnitude more computing power are to improve simulation efficiency and to make better use of available information.

On simulation efficiency, the solution lies in generating events that are more likely to pass the event selection. A technique to improve simulation efficiency has previously been employed in lower energy IceCube neutrino samples. The goal is to preferentially generate events more likely to pass the event selection or equivalently reject events unlikely to pass the event selection before computational costs are incurred. The distribution of event properties at an intermediate level of the event selection is estimated using a small simulation sample and a kernel density estimator. Once the distribution of events in energy, direction, and impact position is determined, a larger sample is generated with this as the new generation distribution. This approach lowers the computational resources dedicated to events that do not pass the event selection and leads to a larger sample at the final selection level. This technique works well if the event selection has multiple levels and a reasonable distribution of muons can be obtained at some point.

However, veto techniques are too efficient at removing these backgrounds with the first cut for this to work. For these cases, we can still attempt to bias the simulation generation distributions, but we must think more carefully about the kinds of events that pass the veto. For a neutrino event, the relevant portion of the veto is a subspace of a cylindrical region behind the interaction vertex, effectively the region near where the event enters the detector. Light detected in this region depends on the accompanying muon's stochastic energy losses and how close these energy losses are to the optical modules. These energy losses are determined after the muon propagation stage, which happens to be computationally inexpensive. Little computational cost has been incurred at this stage, so it is effective to bias the simulation chain here. Knowing the energy losses, the expected number of photons to be observed in the veto region can be estimated. From this estimation, the events are probabilistically rejected. This probabilistic rejection is then accounted for in the weighting scheme.

Once we have improved estimates of the muon distributions from biased simulation, the task remains to better use the information in the analysis. With the current treatment, empty bins in the muon distribution are poorly modeled from the likelihood perspective as the possibility of a non-zero expectation is not entertained. The development of methods to treat these empty bins will improve our statistical description of the muon background information we have available, and remove existing bias from our results. Some suggestions on this topic have been made, but so far, none are simultaneously satisfactory in their modeling of the per bin expectation and the corresponding uncertainty. Possible solutions include nested models of

the binned parameter space, kernel density estimation of the distribution with corresponding uncertainty estimation, and hierarchical modeling of the physical distributions and selection efficiencies. Better treatment of this problem will be vital as we attempt to measure features in more background dominated regions. However, more work on this topic is needed before we can arrive at a satisfactory solution.

9.2.2.2 Atmospheric neutrinos

As shown in the reweighting definitions, the atmospheric neutrino expectations depend on the neutrino flux and passing fraction calculations. Both the flux model and the passing fractions are susceptible to systematic effects in the cosmic ray flux and hadronic interactions. Thankfully it is possible to model these through a combination of the software packages MCEq and nuveto for the flux and passing fractions, respectively. Modeling of the hadronic interaction uncertainties can be rather complex, with modeling and calculation methods varying wildly.

So far, a rough scheme has been developed to describe the hadronic model uncertainties by Barr et. al. [113], which breaks the parameter space into rectangular box regions and assigns uncertainties to each region based on experimental measurements. The chosen box region parameters can then be varied independently according to their prescribed uncertainties. It is possible to develop better uncertainty treatments with the available data, but these have yet to be constructed by analyzers. The Barr scheme for hadronic model modifications has been implemented in MCEq so that these modifications can be translated to variations in the neutrino flux and passing fractions with our available tools. Ideally, we would account for hadronic model variations by allowing these box scaling parameters to vary according to known uncertainties, thereby modifying the background neutrino flux. This method is implemented in IceCube's search for eV scale sterile neutrinos as a first-order approximation, where the gradient of the neutrino flux is computed with respect to the Barr parameters, and the atmospheric flux is allowed to vary linearly. We would hope to apply this same technique to the atmospheric flux and the passing fractions simultaneously in the next generation of analyses. Updating to better treatments of the uncertainty as they become available, and considering higher-order variations may be needed as we incorporate more data into the analysis.

A wide variety of cosmic ray flux models currently exist based on a wealth of measurements. Unfortunately, most cosmic ray experiments cannot differentiate between different nuclei on the level of a single nucleon. For this reason, many models of the cosmic ray flux will group various nuclei into between three and five mass groups, amongst which the spectrum may differ.

Fortunately, these models mostly follow a similar generic prescription that is so far capable of describing the observed data. The cosmic ray spectrum is approximately a power law for any narrow energy range, with changes in the spectral index at the "knee", "second knee", and "ankle". This spectral structure motivates a

multi-component functional form common to many CR flux models, where each component is power law with a cutoff,

$$\frac{d^2\Phi_m}{dEd\Omega}(E_m) = \sum_i^{n_m} a_{i,m} E_m^{-b_{i,m}} e^{-E/c_{i,m}}. \tag{9.1}$$

In this parameterization, the flux for each mass group m depends on the normalizations $a_{i,m}$, spectral indices $b_{i,m}$, and cutoff energies $c_{i,m}$ of the different components. Most models are constructed using similar basis functions, the parameters of which are then fit to some subset of the available cosmic ray data. The same exercise can be performed where a likelihood connects the cosmic ray model parameters and the available data, but with MCEq we can predict the neutrino flux for any arbitrary cosmic ray model. With this connection, we can construct a combined analysis of the neutrino and cosmic ray data, thereby accounting for cosmic ray uncertainties in the most direct manner possible. This setup has some computational limitations as MCEq is slower than other steps in the likelihood evaluation for current analyses, but some possible workarounds limit the extra computation needed for each evaluation.

9.2.2.3 Detector systematics

One of the largest sources of systematic uncertainty in IceCube directional reconstructions is the ice properties. Measurements of the ice properties that use different sources of information tend to be in tension with one another, and so far, no clear explanation has surfaced. A newer approach that models the ice as a birefringent material has shown promise in relieving some of these tensions, but this work is still in progress. Currently, the bulk ice systematics are modeled in analyses as overall modifications to the ice properties. However, shape uncertainties in the depth-dependent ice properties are neglected. Uncertainties on individual DOM properties, including the local ice effects, are either neglected or treated with a global modification. It is possible to account for all of these uncertainties, but many additional parameters make the current treatment with discrete simulation sets for each parameter variation untenable. An alternate approach is given in [243], where a single simulation set is produced in which the detector systematic parameters are randomly chosen according to prior distributions for small bunches of events. After the generation of this simulation set, it is possible to estimate each parameter's effect on the per bin expectation linearly. This treatment has performed well for IceCube's search for eV scale sterile neutrinos and should be applied to future analyses of these samples to more completely account for the detector uncertainties.

9.3 Final remarks

IceCube has opened the door to a new era of astrophysical measurements with its observations of astrophysical neutrinos. Initial measurements of the astrophysical neutrino flux have reaffirmed our understanding of

neutrino physics while simultaneously challenging our expectations. There is still much to learn from astrophysical neutrinos and much to be understood about even our current data. The work presented here focuses on improving the analysis of IceCube's astrophysical data, with several new techniques developed in the process. Such new techniques will be vital as we move into an era of precision measurements, in addition to the suggested analysis improvements described above. Beyond these improvements, next-generation detectors will also be vital to answering the questions raised in recent years. A low energy extension to IceCube will help pin down the ice properties and neutrino oscillation parameters. The planned high-energy extension to IceCube will increase the effective area of IceCube above 10 TeV by a factor of 10, allowing even more precise astrophysical neutrino measurements. Finally, a partner radio array will be capable of observing neutrino interactions at even higher energies, complimenting the IceCube measurements.

Through the combination of these efforts, I have high hopes for the future of neutrino astrophysics.

Bibliography

[1] G. G. Raffelt, "Particle physics from stars," *Ann. Rev. Nucl. Part. Sci.*, vol. 49, pp. 163–216, 1999. DOI: 10.1146/annurev.nucl.49.1.163. arXiv: hep-ph/9903472 [hep-ph].

[2] J. N. Bahcall, A. M. Serenelli, and S. Basu, "New solar opacities, abundances, helioseismology, and neutrino fluxes," *Astrophys. J.*, vol. 621, pp. L85–L88, 2005. DOI: 10.1086/428929. arXiv: astro-ph/0412440 [astro-ph].

[3] A. B. McDonald, "Nobel Lecture: The Sudbury Neutrino Observatory: Observation of flavor change for solar neutrinos," *Rev. Mod. Phys.*, vol. 88, no. 3, p. 030 502, 2016. DOI: 10.1103/RevModPhys.88.030502.

[4] K. Hirata *et al.*, "Observation of a Neutrino Burst from the Supernova SN 1987a," *Phys. Rev. Lett.*, vol. 58, pp. 1490–1493, 1987, [,727(1987)]. DOI: 10.1103/PhysRevLett.58.1490.

[5] C. B. Bratton *et al.*, "Angular Distribution of Events From Sn1987a," *Phys. Rev.*, vol. D37, p. 3361, 1988. DOI: 10.1103/PhysRevD.37.3361.

[6] H. A. Bethe, "Supernova mechanisms," *Rev. Mod. Phys.*, vol. 62, pp. 801–866, 1990. DOI: 10.1103/RevModPhys.62.801.

[7] W. D. Arnett and J. L. Rosner, "Neutrino Mass Limits From Sn1987a," *Phys. Rev. Lett.*, vol. 58, p. 1906, 1987. DOI: 10.1103/PhysRevLett.58.1906.

[8] I. Goldman, Y. Aharonov, G. Alexander, and S. Nussinov, "Implications of the Supernova Sn1987a Neutrino Signals," *Phys. Rev. Lett.*, vol. 60, p. 1789, 1988. DOI: 10.1103/PhysRevLett.60.1789.

[9] A. Manohar, "A Limit on the Neutrino-neutrino Scattering Cross-section From the Supernova," *Phys. Lett.*, vol. B192, p. 217, 1987. DOI: 10.1016/0370-2693(87)91171-3.

[10] J. M. Lattimer and J. Cooperstein, "Limits on the Neutrino Magnetic Moment from SN 1987a," *Phys. Rev. Lett.*, vol. 61, pp. 23–26, 1988, [Erratum: Phys. Rev. Lett.61,2633(1988)]. DOI: 10.1103/PhysRevLett.61.23,10.1103/PhysRevLett.61.2633.

[11] G. Raffelt and D. Seckel, "Bounds on Exotic Particle Interactions from SN 1987a," *Phys. Rev. Lett.*, vol. 60, p. 1793, 1988. DOI: 10.1103/PhysRevLett.60.1793.

[12] M. S. Turner, "Axions from SN 1987a," *Phys. Rev. Lett.*, vol. 60, p. 1797, 1988. DOI: 10.1103/PhysRevLett.60.1797.

[13] T. K. Gaisser, F. Halzen, and T. Stanev, "Particle astrophysics with high-energy neutrinos," *Phys. Rept.*, vol. 258, pp. 173–236, 1995, [Erratum: Phys. Rept.271,355(1996)]. DOI: 10.1016/0370-1573(95)00003-Y. arXiv: hep-ph/9410384 [hep-ph].

[14] T. Gaisser, *Cosmic Rays and Particle Physics*. Cambridge University Press, 1990, ISBN: 9780521339315. [Online]. Available: https://books.google.com/books?id=qJ7Z6oIMqeUC.

[15] A. Letessier-Selvon and T. Stanev, "Ultrahigh Energy Cosmic Rays," *Rev. Mod. Phys.*, vol. 83, pp. 907–942, 2011. DOI: 10.1103/RevModPhys.83.907. arXiv: 1103.0031 [astro-ph.HE].

[16] K. Kotera and A. V. Olinto, "The Astrophysics of Ultrahigh Energy Cosmic Rays," *Ann. Rev. Astron. Astrophys.*, vol. 49, pp. 119–153, 2011. DOI: 10.1146/annurev-astro-081710-102620. arXiv: 1101.4256 [astro-ph.HE].

[17] M. Ackermann *et al.*, "The spectrum of isotropic diffuse gamma-ray emission between 100 MeV and 820 GeV," *Astrophys. J.*, vol. 799, p. 86, 2015. DOI: 10.1088/0004-637X/799/1/86. arXiv: 1410.3696 [astro-ph.HE].

[18] F. Fenu, "The cosmic ray energy spectrum measured using the Pierre Auger Observatory," pp. 9–16, 2017, [PoSICRC2017,486(2018)]. DOI: 10.22323/1.301.0486.

[19] M. Ahlers, K. Helbing, and C. Pérez de los Heros, "Probing Particle Physics with IceCube," *Eur. Phys. J.*, vol. C78, no. 11, p. 924, 2018. DOI: 10.1140/epjc/s10052-018-6369-9. arXiv: 1806.05696 [astro-ph.HE].

[20] M. Ackermann *et al.*, "Fundamental Physics with High-Energy Cosmic Neutrinos," *Bull. Am. Astron. Soc.*, vol. 51, p. 215, 2019. arXiv: 1903.04333 [astro-ph.HE].

[21] C. A. Argüelles, M. Bustamante, A. Kheirandish, S. Palomares-Ruiz, J. Salvado, and A. C. Vincent, "Fundamental physics with high-energy cosmic neutrinos today and in the future," in *36th International Cosmic Ray Conference (ICRC 2019) Madison, Wisconsin, USA, July 24-August 1, 2019*, 2019. arXiv: 1907.08690 [astro-ph.HE].

[22] J. S. Diaz, A. Kostelecky, and M. Mewes, "Testing Relativity with High-Energy Astrophysical Neutrinos," *Phys. Rev.*, vol. D89, no. 4, p. 043 005, 2014. DOI: 10.1103/PhysRevD.89.043005. arXiv: 1308.6344 [astro-ph.HE].

[23] C. A. Argüelles, T. Katori, and J. Salvado, "New Physics in Astrophysical Neutrino Flavor," *Phys. Rev. Lett.*, vol. 115, p. 161 303, 2015. DOI: 10.1103/PhysRevLett.115.161303. arXiv: 1506.02043 [hep-ph].

[24] J. H. Davis and J. Silk, "Spectral and Spatial Distortions of PeV Neutrinos from Scattering with Dark Matter," 2015. arXiv: 1505.01843 [hep-ph].

[25] J. F. Cherry, A. Friedland, and I. M. Shoemaker, "Short-baseline neutrino oscillations, Planck, and IceCube," 2016. arXiv: 1605.06506 [hep-ph].

[26] C. A. Argüelles, A. Kheirandish, and A. C. Vincent, "Imaging Galactic Dark Matter with High-Energy Cosmic Neutrinos," *Phys. Rev. Lett.*, vol. 119, no. 20, p. 201 801, 2017. DOI: 10.1103/PhysRevLett.119.201801. arXiv: 1703.00451 [hep-ph].

[27] K. J. Kelly and P. A. N. Machado, "Multimessenger Astronomy and New Neutrino Physics," *JCAP*, vol. 1810, p. 048, 2018. DOI: 10.1088/1475-7516/2018/10/048. arXiv: 1808.02889 [hep-ph].

[28] F. Capozzi, I. M. Shoemaker, and L. Vecchi, "Neutrino Oscillations in Dark Backgrounds," *JCAP*, vol. 1807, no. 07, p. 004, 2018. DOI: 10.1088/1475-7516/2018/07/004. arXiv: 1804.05117 [hep-ph].

[29] Y. Farzan and S. Palomares-Ruiz, "Flavor of cosmic neutrinos preserved by ultralight dark matter," *Phys. Rev.*, vol. D99, no. 5, p. 051 702, 2019. DOI: 10.1103/PhysRevD.99.051702. arXiv: 1810.00892 [hep-ph].

[30] K.-Y. Choi, E. J. Chun, and J. Kim, "Neutrino Oscillations in Dark Matter," 2019. arXiv: 1909.10478 [hep-ph].

[31] R. W. Rasmussen, L. Lechner, M. Ackermann, M. Kowalski, and W. Winter, "Astrophysical neutrinos flavored with Beyond the Standard Model physics," *Phys. Rev.*, vol. D96, no. 8, p. 083 018, 2017. DOI: 10.1103/PhysRevD.96.083018. arXiv: 1707.07684 [hep-ph].

[32] J. Ellis, N. E. Mavromatos, A. S. Sakharov, and E. K. Sarkisyan-Grinbaum, "Limits on Neutrino Lorentz Violation from Multimessenger Observations of TXS 0506+056," *Phys. Lett.*, vol. B789, pp. 352–355, 2019. DOI: 10.1016/j.physletb.2018.11.062. arXiv: 1807.05155 [astro-ph.HE].

[33] J. Liao and D. Marfatia, "IceCube's astrophysical neutrino energy spectrum from CPT violation," *Phys. Rev.*, vol. D97, no. 4, p. 041 302, 2018. DOI: 10.1103/PhysRevD.97.041302. arXiv: 1711.09266 [hep-ph].

[34] M. G. Aartsen *et al.*, "Search for neutrinos from decaying dark matter with IceCube," *Eur. Phys. J.*, vol. C78, no. 10, p. 831, 2018. DOI: 10.1140/epjc/s10052-018-6273-3. arXiv: 1804.03848 [astro-ph.HE].

[35] M. Aartsen *et al.*, "Measurement of the multi-TeV neutrino cross section with IceCube using Earth absorption," *Nature*, vol. 551, pp. 596–600, 2017. DOI: 10.1038/nature24459. arXiv: 1711.08119 [hep-ex].

[36] Y. Xu, "TeV-PeV neutrino-nucleon cross section measurement with 5 years of IceCube data," *PoS*, vol. DIS2018, p. 019, 2018. DOI: 10.22323/1.316.0019. arXiv: 1809.06782 [hep-ex].

[37] V. Bertone, R. Gauld, and J. Rojo, "Neutrino Telescopes as QCD Microscopes," *JHEP*, vol. 01, p. 217, 2019. DOI: 10.1007/JHEP01(2019)217. arXiv: 1808.02034 [hep-ph].

[38] L. A. Anchordoqui, C. García Canal, and J. F. Soriano, "Probing strong dynamics with cosmic neutrinos," *Phys. Rev.*, vol. D100, no. 10, p. 103 001, 2019. DOI: 10.1103/PhysRevD.100.103001. arXiv: 1902.10134 [hep-ph].

[39] Y. Xu, "Measurement of high energy neutrino interaction cross section using neutrino-induced electromagnetic and hadronic showers observed in five years of icecube data," PhD thesis, Stony Brook University, 2019.

[40] E. M. Henley and J. Jalilian-Marian, "Ultra-high energy neutrino-nucleon scattering and parton distributions at small x," *Phys. Rev.*, vol. D73, p. 094 004, 2006. DOI: 10.1103/PhysRevD.73.094004. arXiv: hep-ph/0512220 [hep-ph].

[41] M. M. Block, L. Durand, and P. Ha, "Connection of the virtual $\gamma^* p$ cross section of ep deep inelastic scattering to real γp scattering, and the implications for νN and ep total cross sections," *Phys. Rev.*, vol. D89, no. 9, p. 094 027, 2014. DOI: 10.1103/PhysRevD.89.094027. arXiv: 1404.4530 [hep-ph].

[42] V. P. Goncalves and D. R. Gratieri, "Investigating the effects of the QCD dynamics in the neutrino absorption by the Earth's interior at ultrahigh energies," *Phys. Rev.*, vol. D92, no. 11, p. 113 007, 2015. DOI: 10.1103/PhysRevD.92.113007. arXiv: 1510.03186 [hep-ph].

[43] Argüelles, C. A. and Halzen, F. and Wille, L. and Kroll, M. and Reno, M. H., "High-energy behavior of photon, neutrino, and proton cross sections," *Phys. Rev.*, vol. D92, no. 7, p. 074 040, 2015. DOI: 10.1103/PhysRevD.92.074040. arXiv: 1504.06639 [hep-ph].

[44] J. Alvarez-Muniz, F. Halzen, T. Han, and D. Hooper, "Phenomenology of high-energy neutrinos in low scale quantum gravity models," *Phys. Rev. Lett.*, vol. 88, p. 021 301, 2002. DOI: 10.1103/PhysRevLett.88.021301. arXiv: hep-ph/0107057 [hep-ph].

[45] L. A. Anchordoqui, C. A. Garcia Canal, H. Goldberg, D. Gomez Dumm, and F. Halzen, "Probing leptoquark production at IceCube," *Phys. Rev. D*, vol. 74, p. 125 021, 2006. DOI: 10.1103/PhysRevD.74.125021. arXiv: hep-ph/0609214.

[46] U. K. Dey and S. Mohanty, "Constraints on Leptoquark Models from IceCube Data," *JHEP*, vol. 04, p. 187, 2016. DOI: 10.1007/JHEP04(2016)187. arXiv: 1505.01037 [hep-ph].

[47] B. Dutta, Y. Gao, T. Li, C. Rott, and L. E. Strigari, "Leptoquark implication from the CMS and IceCube experiments," *Phys. Rev. D*, vol. 91, p. 125 015, 2015. DOI: 10.1103/PhysRevD.91.125015. arXiv: 1505.00028 [hep-ph].

[48] U. K. Dey, S. Mohanty, and G. Tomar, "Leptoquarks: 750 GeV Diphoton Resonance and IceCube Events," Jun. 2016. arXiv: 1606.07903 [hep-ph].

[49] U. K. Dey, D. Kar, M. Mitra, M. Spannowsky, and A. C. Vincent, "Searching for Leptoquarks at IceCube and the LHC," *Phys. Rev. D*, vol. 98, no. 3, p. 035 014, 2018. DOI: 10.1103/PhysRevD.98.035014. arXiv: 1709.02009 [hep-ph].

[50] B. Chauhan, B. Kindra, and A. Narang, "Discrepancies in simultaneous explanation of flavor anomalies and IceCube PeV events using leptoquarks," *Phys. Rev. D*, vol. 97, no. 9, p. 095 007, 2018. DOI: 10.1103/PhysRevD.97.095007. arXiv: 1706.04598 [hep-ph].

[51] D. Bečirević, B. Panes, O. Sumensari, and R. Zukanovich Funchal, "Seeking leptoquarks in IceCube," *JHEP*, vol. 06, p. 032, 2018. DOI: 10.1007/JHEP06(2018)032. arXiv: 1803.10112 [hep-ph].

[52] A. G. Cohen, D. Kaplan, and A. Nelson, "Progress in electroweak baryogenesis," *Ann. Rev. Nucl. Part. Sci.*, vol. 43, pp. 27–70, 1993. DOI: 10.1146/annurev.ns.43.120193.000331. arXiv: hep-ph/9302210.

[53] J. Ellis, K. Sakurai, and M. Spannowsky, "Search for Sphalerons: IceCube vs. LHC," *JHEP*, vol. 05, p. 085, 2016. DOI: 10.1007/JHEP05(2016)085. arXiv: 1603.06573 [hep-ph].

[54] J. Alvarez-Muniz, J. L. Feng, F. Halzen, T. Han, and D. Hooper, "Detecting microscopic black holes with neutrino telescopes," *Phys. Rev.*, vol. D65, p. 124 015, 2002. DOI: 10.1103/PhysRevD.65.124015. arXiv: hep-ph/0202081 [hep-ph].

[55] K. J. Mack, N. Song, and A. C. Vincent, "Signatures of microscopic black holes and extra dimensions at future neutrino telescopes," 2019. arXiv: 1912.06656 [hep-ph].

[56] S. R. Klein, "Probing high-energy interactions of atmospheric and astrophysical neutrinos," in. 2020, pp. 75–107. DOI: 10.1142/9789813275027_0004. arXiv: 1906.02221 [astro-ph.HE].

[57] M. G. Aartsen *et al.*, "Evidence for Astrophysical Muon Neutrinos from the Northern Sky with IceCube," *Phys. Rev. Lett.*, vol. 115, no. 8, p. 081 102, 2015. DOI: 10.1103/PhysRevLett.115.081102. arXiv: 1507.04005 [astro-ph.HE].

[58] M. G. Aartsen *et al.*, "Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data," *Astrophys. J.*, vol. 833, no. 1, p. 3, 2016. DOI: 10.3847/0004-637X/833/1/3. arXiv: 1607.08006 [astro-ph.HE].

[59] M. G. Aartsen *et al.*, "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector," *Science*, vol. 342, p. 1 242 856, 2013. DOI: 10.1126/science.1242856. arXiv: 1311.5238 [astro-ph.HE].

[60] M. G. Aartsen *et al.*, "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data," *Phys. Rev. Lett.*, vol. 113, p. 101 101, 2014. DOI: 10.1103/PhysRevLett.113.101101. arXiv: 1405.5303 [astro-ph.HE].

[61] M. G. Aartsen *et al.*, "The IceCube Neutrino Observatory - Contributions to ICRC 2015 Part II: Atmospheric and Astrophysical Diffuse Neutrino Searches of All Flavors," in *Proceedings, 34th International Cosmic Ray Conference (ICRC 2015): The Hague, The Netherlands, July 30-August 6, 2015*, 2015. arXiv: 1510.05223 [astro-ph.HE]. [Online]. Available: https://inspirehep.net/record/1398539/files/arXiv:1510.05223.pdf.

[62] M. G. Aartsen *et al.*, "The IceCube Neutrino Observatory - Contributions to ICRC 2017 Part II: Properties of the Atmospheric and Astrophysical Neutrino Flux," in *Proceedings, 35th International Cosmic Ray Conference (ICRC 2017): Bexco, Busan, Korea, July 12-20, 2017*, vol. ICRC2017, 2017. arXiv: 1710.01191 [astro-ph.HE].

[63] M. G. Aartsen *et al.*, "Atmospheric and astrophysical neutrinos above 1 TeV interacting in IceCube," *Phys. Rev.*, vol. D91, no. 2, p. 022 001, 2015. DOI: 10.1103/PhysRevD.91.022001. arXiv: 1410.1749 [astro-ph.HE].

[64] M. G. Aartsen *et al.*, "Characteristics of the diffuse astrophysical electron and tau neutrino flux with six years of IceCube high energy cascade data," 2020. arXiv: 2001.09520 [astro-ph.HE].

[65] C. A. Argüelles, S. Palomares-Ruiz, A. Schneider, L. Wille, and T. Yuan, "Unified atmospheric neutrino passing fractions for large-scale neutrino telescopes," *JCAP*, vol. 1807, no. 07, p. 047, 2018. DOI: 10.1088/1475-7516/2018/07/047. arXiv: 1805.11003 [hep-ph].

[66] Argüelles, Carlos A. and Schneider, Austin and Yuan, Tianlu, "A binned likelihood for stochastic models," *JHEP*, vol. 06, p. 030, 2019. DOI: 10.1007/JHEP06(2019)030. arXiv: 1901.04645 [physics.data-an].

[67] T. K. Gaisser, R. Engel, and E. Resconi, *Cosmic Rays and Particle Physics: 2nd Edition*. Cambridge University Press, Jun. 2016, ISBN: 978-0-521-01646-9.

[68] A. Abdo *et al.*, "Fermi Large Area Telescope Measurements of the Diffuse Gamma-Ray Emission at Intermediate Galactic Latitudes," *Phys. Rev. Lett.*, vol. 103, p. 251 101, 2009. DOI: 10.1103/PhysRevLett.103.251101. arXiv: 0912.0973 [astro-ph.HE].

[69] C. Dermer *et al.*, "On the Physics Connecting Cosmic Rays and Gamma Rays: Towards Determining the Interstellar Cosmic Ray Spectrum," Mar. 2013. arXiv: 1303.6482 [astro-ph.HE].

[70] K. V. Ptitsyna and S. V. Troitsky, "Physical conditions in potential sources of ultra-high-energy cosmic rays. I. Updated Hillas plot and radiation-loss constraints," *Phys. Usp.*, vol. 53, pp. 691–701, 2010. DOI: 10.3367/UFNe.0180.201007c.0723. arXiv: 0808.0367 [astro-ph].

[71] V. F. Hess, "Uber beobachtungen der durchdringenden strahlung bei sieben freiballonfahrten," *Phys. Zeits.*, vol. 13, pp. 1084–1091, 1912.

[72] M. Tanabashi *et al.*, "Review of particle physics," *Phys. Rev. D*, vol. 98, p. 030 001, 3 Aug. 2018. DOI: 10.1103/PhysRevD.98.030001. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevD.98.030001.

[73] F. Vannucci, "Interactions of neutrinos with matter," *Prog. Part. Nucl. Phys.*, vol. 95, pp. 1–47, 2017. DOI: 10.1016/j.ppnp.2017.03.003.

[74] D. Akimov *et al.*, "Observation of Coherent Elastic Neutrino-Nucleus Scattering," *Science*, vol. 357, no. 6356, pp. 1123–1126, 2017. DOI: 10.1126/science.aao0990. arXiv: 1708.01294 [nucl-ex].

[75] S. L. Glashow, "Resonant Scattering of Antineutrinos," *Phys. Rev.*, vol. 118, pp. 316–317, 1960. DOI: 10.1103/PhysRev.118.316.

[76] A. Loewy, S. Nussinov, and S. L. Glashow, "The Effect of Doppler Broadening on the 6.3 PeV W^- Resonance in $\bar{\nu}_e e^-$ Collisions," 2014. arXiv: 1407.4415 [hep-ph].

[77] A. Cooper-Sarkar, P. Mertsch, and S. Sarkar, "The high energy neutrino cross-section in the Standard Model and its uncertainty," *JHEP*, vol. 08, p. 042, 2011. DOI: 10.1007/JHEP08(2011)042. arXiv: 1106.3723 [hep-ph].

[78] R. Gandhi, C. Quigg, M. H. Reno, and I. Sarcevic, "Ultrahigh-energy neutrino interactions," *Astropart. Phys.*, vol. 5, pp. 81–110, 1996. DOI: 10.1016/0927-6505(96)00008-4. arXiv: hep-ph/9512364 [hep-ph].

[79] J. H. Koehne, K. Frantzen, M. Schmitz, T. Fuchs, W. Rhode, D. Chirkin, and J. Becker Tjus, "PROPOSAL: A tool for propagation of charged leptons," *Comput. Phys. Commun.*, vol. 184, pp. 2070–2090, 2013. DOI: 10.1016/j.cpc.2013.04.001.

[80] M. G. Aartsen *et al.*, "Measurement of South Pole ice transparency with the IceCube LED calibration system," *Nucl. Instrum. Meth.*, vol. A711, pp. 73–89, 2013. DOI: 10.1016/j.nima.2013.01.054. arXiv: 1301.5361 [astro-ph.IM].

[81] M. G. Aartsen *et al.*, "The IceCube Neutrino Observatory: Instrumentation and Online Systems," *JINST*, vol. 12, no. 03, P03012, 2017. DOI: 10.1088/1748-0221/12/03/P03012. arXiv: 1612.05093 [astro-ph.IM].

[82] R. Abbasi *et al.*, "The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping," *Nucl. Instrum. Meth.*, vol. A601, pp. 294–316, 2009. DOI: 10.1016/j.nima.2009.01.001. arXiv: 0810.4930 [physics.ins-det].

[83] R. Abbasi *et al.*, "The Design and Performance of IceCube DeepCore," *Astropart. Phys.*, vol. 35, pp. 615–624, 2012. DOI: 10.1016/j.astropartphys.2012.01.004. arXiv: 1109.6096 [astro-ph.IM].

[84] F. Halzen and D. Hooper, "High-energy neutrino astronomy: The Cosmic ray connection," *Rept. Prog. Phys.*, vol. 65, pp. 1025–1078, 2002. DOI: 10.1088/0034-4885/65/7/201. arXiv: astro-ph/0204527.

[85] J. Learned and K. Mannheim, "High-energy neutrino astrophysics," *Ann. Rev. Nucl. Part. Sci.*, vol. 50, pp. 679–749, 2000. DOI: 10.1146/annurev.nucl.50.1.679.

[86] J. G. Learned and S. Pakvasa, "Detecting tau-neutrino oscillations at PeV energies," *Astropart. Phys.*, vol. 3, pp. 267–274, 1995. DOI: 10.1016/0927-6505(94)00043-3. arXiv: hep-ph/9405296 [hep-ph].

[87] D. F. Cowen, "Tau neutrinos in IceCube," *J. Phys. Conf. Ser.*, vol. 60, pp. 227–230, 2007. DOI: 10.1088/1742-6596/60/1/048.

[88] C. Patrignani *et al.*, "Review of Particle Physics," *Chin. Phys.*, vol. C40, no. 10, p. 100 001, 2016. DOI: 10.1088/1674-1137/40/10/100001.

[89] Y. Farzan and A. Yu. Smirnov, "Leptonic unitarity triangle and CP violation," *Phys. Rev.*, vol. D65, p. 113 001, 2002. DOI: 10.1103/PhysRevD.65.113001. arXiv: hep-ph/0201105 [hep-ph].

[90] A. Palladino and F. Vissani, "The natural parameterization of cosmic neutrino oscillations," *Eur. Phys. J.*, vol. C75, p. 433, 2015. DOI: 10.1140/epjc/s10052-015-3664-6. arXiv: 1504.05238 [hep-ph].

[91] M. Bustamante, J. F. Beacom, and W. Winter, "Theoretically palatable flavor combinations of astrophysical neutrinos," *Phys. Rev. Lett.*, vol. 115, no. 16, p. 161 302, 2015. DOI: 10.1103/PhysRevLett.115.161302. arXiv: 1506.02645 [astro-ph.HE].

[92] M. Ahlers, M. Bustamante, and S. Mu, "Unitarity Bounds of Astrophysical Neutrinos," *Phys. Rev.*, vol. D98, no. 12, p. 123 023, 2018. DOI: 10.1103/PhysRevD.98.123023. arXiv: 1810.00893 [astro-ph.HE].

[93] C. A. Argüelles, K. Farrag, T. Katori, R. Khandelwal, S. Mandalia, and J. Salvado, "Probe of Sterile Neutrinos Using Astrophysical Neutrino Flavor," 2019. arXiv: 1909.05341 [hep-ph].

[94] M. G. Aartsen *et al.*, "Flavor Ratio of Astrophysical Neutrinos above 35 TeV in IceCube," *Phys. Rev. Lett.*, vol. 114, no. 17, p. 171 102, 2015. DOI: 10.1103/PhysRevLett.114.171102. arXiv: 1502.03376 [astro-ph.HE].

[95] M. G. Aartsen *et al.*, "A combined maximum-likelihood analysis of the high-energy astrophysical neutrino flux measured with IceCube," *Astrophys. J.*, vol. 809, no. 1, p. 98, 2015. DOI: 10.1088/0004-637X/809/1/98. arXiv: 1507.03991 [astro-ph.HE].

[96] A. Palladino, G. Pagliaroli, F. L. Villante, and F. Vissani, "What is the Flavor of the Cosmic Neutrinos Seen by IceCube?" *Phys. Rev. Lett.*, vol. 114, no. 17, p. 171 101, 2015. DOI: 10.1103/PhysRevLett.114.171101. arXiv: 1502.02923 [astro-ph.HE].

[97] O. Mena, S. Palomares-Ruiz, and A. C. Vincent, "Flavor Composition of the High-Energy Neutrino Events in IceCube," *Phys. Rev. Lett.*, vol. 113, p. 091 103, 2014. DOI: 10.1103/PhysRevLett.113.091103. arXiv: 1404.0017 [astro-ph.HE].

[98] S. Palomares-Ruiz, A. C. Vincent, and O. Mena, "Spectral analysis of the high-energy IceCube neutrinos," *Phys. Rev.*, vol. D91, no. 10, p. 103 008, 2015. DOI: 10.1103/PhysRevD.91.103008. arXiv: 1502.02649 [astro-ph.HE].

[99] A. Karle, "Entwicklung eines neuartigen atmosphärischen Tscherenkovdetektors und Messungen an hochenergetischer Kosmischer Strahlung zwischen 15 und 1000 TeV," PhD thesis, München Univ., 1994.

[100] M. G. Aartsen *et al.*, "Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data," *Phys. Rev.*, vol. D91, no. 7, p. 072 004, 2015. DOI: 10.1103/PhysRevD.91.072004. arXiv: 1410.7227 [hep-ex].

[101] M. G. Aartsen *et al.*, "Measurement of Atmospheric Neutrino Oscillations at 6-56 GeV with IceCube DeepCore," 2017. arXiv: 1707.07081 [hep-ex].

[102] N. Whitehorn, J. van Santen, and S. Lafebre, "Penalized Splines for Smooth Representation of High-dimensional Monte Carlo Datasets," *Comput. Phys. Commun.*, vol. 184, pp. 2214–2220, 2013. DOI: 10.1016/j.cpc.2013.04.008. arXiv: 1301.2184 [physics.data-an].

[103] IceCube, ***photospline***, https://github.com/IceCubeOpenSource/photospline, 2019.

[104] D. Williams, "Light propagation in the South Pole ice," *AIP Conf. Proc.*, vol. 1630, no. 1, C. Finley *et al.*, Eds., pp. 146–149, 2015. DOI: 10.1063/1.4902793.

[105] D. Chirkin and M. Rongen, "Light diffusion in birefringent polycrystals and the IceCube ice anisotropy," *PoS*, vol. ICRC2019, p. 854, 2020. DOI: 10.22323/1.358.0854. arXiv: 1908.07608 [astro-ph.HE].

[106] A. Achterberg *et al.*, "First Year Performance of The IceCube Neutrino Telescope," *Astropart. Phys.*, vol. 26, pp. 155–173, 2006. DOI: 10.1016/j.astropartphys.2006.06.007. arXiv: astro-ph/0604450 [astro-ph].

[107] J. Stettner, "Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data," in *The IceCube Neutrino Observatory – Contributions to the 36th International Cosmic Ray Conference (ICRC2019)*, 2019. arXiv: 1908.09551 [astro-ph.HE].

[108] M. G. Aartsen *et al.*, "HESE Tau discovery paper," *PRL*, 201X.

[109] M. G. Aartsen *et al.*, "In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes," 2020. arXiv: 2002.00997 [physics.ins-det].

[110] M. G. Aartsen *et al.*, "Energy Reconstruction Methods in the IceCube Neutrino Telescope," *JINST*, vol. 9, P03009, 2014. DOI: 10.1088/1748-0221/9/03/P03009. arXiv: 1311.4767 [physics.ins-det].

[111] M. Usner, "Search for astrophysical tau-neutrinos in six years of high-energy starting events in the icecube detector," PhD thesis, Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät, 2018. DOI: http://dx.doi.org/10.18452/19458.

[112] T. K. Gaisser and M. Honda, "Flux of atmospheric neutrinos," *Ann. Rev. Nucl. Part. Sci.*, vol. 52, pp. 153–199, 2002. DOI: 10.1146/annurev.nucl.52.050102.090645. arXiv: hep-ph/0203272 [hep-ph].

[113] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev, "A three-dimensional calculation of atmospheric neutrinos," *Phys. Rev.*, vol. D70, p. 023 006, 2004. DOI: 10.1103/PhysRevD.70.023006. arXiv: astro-ph/0403630 [astro-ph].

[114] M. Honda, T. Kajita, K. Kasahara, S. Midorikawa, and T. Sanuki, "Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data," *Phys. Rev.*, vol. D75, p. 043 006, 2007. DOI: 10.1103/PhysRevD.75.043006. arXiv: astro-ph/0611418 [astro-ph].

[115] O. N. Petrova, T. S. Sinegovskaya, and S. I. Sinegovsky, "High-energy spectra of atmospheric neutrinos," *Phys. Part. Nucl. Lett.*, vol. 9, pp. 766–768, 2012. DOI: 10.1134/S1547477112070138.

[116] T. Montaruli and F. Ronga, "Comparison of muon and neutrino times from decays of mesons in the atmosphere," 2011. arXiv: 1109.6238 [hep-ex].

[117] A. Bhattacharya, R. Enberg, M. H. Reno, I. Sarcevic, and A. Stasto, "Perturbative charm production and the prompt atmospheric neutrino flux in light of RHIC and LHC," *JHEP*, vol. 06, p. 110, 2015. DOI: 10.1007/JHEP06(2015)110. arXiv: 1502.01076 [hep-ph].

[118] A. Gazizov and M. P. Kowalski, "ANIS: High energy neutrino generator for neutrino telescopes," *Comput. Phys. Commun.*, vol. 172, pp. 203–213, 2005. DOI: 10.1016/j.cpc.2005.03.113. arXiv: astro-ph/0406439 [astro-ph].

[119] A. M. Dziewonski and D. L. Anderson, "Preliminary reference earth model," *Phys. Earth Planet. Interiors*, vol. 25, pp. 297–356, 1981. DOI: 10.1016/0031-9201(81)90046-7.

[120] A. C. Vincent, C. A. Argüelles, and A. Kheirandish, "High-energy neutrino attenuation in the Earth and its associated uncertainties," *JCAP*, vol. 1711, no. 11, p. 012, 2017, [JCAP1711,012(2017)]. DOI: 10.1088/1475-7516/2017/11/012. arXiv: 1706.09895 [hep-ph].

[121] T. K. Gaisser and S. R. Klein, "A new contribution to the conventional atmospheric neutrino flux," *Astropart. Phys.*, vol. 64, pp. 13–17, 2015. DOI: 10.1016/j.astropartphys.2014.10.006. arXiv: 1409.4924 [astro-ph.HE].

[122] H. P. Dembinski, R. Engel, A. Fedynitch, T. Gaisser, F. Riehn, and T. Stanev, "Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to 10^{11} GeV," *PoS*, vol. ICRC2017, p. 533, 2017. arXiv: 1711.11432 [astro-ph.HE].

[123] A. Fedynitch, J. Becker Tjus, and P. Desiati, "Influence of hadronic interaction models and the cosmic ray spectrum on the high energy atmospheric muon and neutrino flux," *Phys. Rev.*, vol. D86, p. 114 024, 2012. DOI: 10.1103/PhysRevD.86.114024. arXiv: 1206.6710 [astro-ph.HE].

[124] A. V. Karelin, S. V. Borisov, A. M. Galper, and S. A. Voronov, "The proton and helium cosmic ray spectra from 50-GeV to 15-TeV," *Astrophys. Space Sci. Trans.*, vol. 7, pp. 235–238, 2011. DOI: 10.5194/astra-7-235-2011.

[125] B. Bartoli *et al.*, "Cosmic ray proton plus helium energy spectrum measured by the ARGO-YBJ experiment in the energy range 3–300 TeV," *Phys. Rev.*, vol. D91, no. 11, p. 112 017, 2015. DOI: 10.1103/PhysRevD.91.112017. arXiv: 1503.07136 [hep-ex].

[126] Y. S. Yoon *et al.*, "Proton and Helium Spectra from the CREAM-III Flight," *Astrophys. J.*, vol. 839, no. 1, p. 5, 2017. DOI: 10.3847/1538-4357/aa68e4. arXiv: 1704.02512 [astro-ph.HE].

[127] R. Alfaro *et al.*, "All-particle cosmic ray energy spectrum measured by the HAWC experiment from 10 to 500 TeV," *Phys. Rev.*, vol. D96, no. 12, p. 122 001, 2017. DOI: 10.1103/PhysRevD.96.122001. arXiv: 1710.00890 [astro-ph.HE].

[128] G. Collin and J. Conrad, *An estimation of systematics for up-going atmospheric muon neutrino flux at the south pole*, http://hdl.handle.net/1721.1/98078, 2015.

[129] B. J. P. Jones, "Sterile Neutrinos in Cold Climates," PhD thesis, MIT, 2015. DOI: 10.2172/1221354. [Online]. Available: http://lss.fnal.gov/archive/thesis/2000/fermilab-thesis-2015-17.pdf.

[130] S. Schonert, T. K. Gaisser, E. Resconi, and O. Schulz, "Vetoing atmospheric neutrinos in a high energy neutrino telescope," *Phys. Rev.*, vol. D79, p. 043 009, 2009. DOI: 10.1103/PhysRevD.79.043009. arXiv: 0812.4308 [astro-ph].

[131] T. K. Gaisser, K. Jero, A. Karle, and J. van Santen, "Generalized self-veto probability for atmospheric neutrinos," *Phys. Rev.*, vol. D90, no. 2, p. 023 009, 2014. DOI: 10.1103/PhysRevD.90.023009. arXiv: 1405.0525 [astro-ph.HE].

[132] T. K. Gaisser, "Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio," *Astropart. Phys.*, vol. 35, pp. 801–806, 2012. DOI: 10.1016/j.astropartphys.2012.02.010. arXiv: 1111.6675 [astro-ph.HE].

[133] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, "Charm production in SIBYLL," *EPJ Web Conf.*, vol. 99, p. 12 001, 2015. DOI: 10.1051/epjconf/20159912001. arXiv: 1502.06353 [hep-ph].

[134] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, "A new version of the event generator Sibyll," *PoS*, vol. ICRC2015, p. 558, 2016. arXiv: 1510.00568 [hep-ph].

[135] K. Labitzke, J. J. Barnett, and B. Edwards, Eds., *Handbook MAP 16*. University of Illinois, Urbana: SCOSTEP, 1985.

[136] A. E. Hedin, "Extension of the MSIS thermospheric model into the middle and lower atmosphere," *J. Geophys. Res.*, vol. 96, p. 1159, 1991. DOI: 10.1029/2002JA009430.

[137] A. Fedynitch, R. Engel, T. K. Gaisser, F. Riehn, and T. Stanev, "Calculation of conventional and prompt lepton fluxes at very high energy," *EPJ Web Conf.*, vol. 99, p. 08 001, 2015. DOI: 10.1051/epjconf/20159908001. arXiv: 1503.00544 [hep-ph].

[138] F. Riehn, H. P. Dembinski, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, "The hadronic interaction model SIBYLL 2.3c and Feynman scaling," *PoS*, vol. ICRC2017, p. 301, 2017. arXiv: 1709.07227 [hep-ph].

[139] D. Chirkin and W. Rhode, "Muon Monte Carlo: A high-precision tool for muon propagation through matter," 2004. arXiv: hep-ph/0407075 [hep-ph].

[140] H. Abramowicz, E. M. Levin, A. Levy, and U. Maor, "A parametrization of σ_{T} (γ^* p) above the resonance region $Q^2 \geq 0$," *Phys. Lett.*, vol. B269, pp. 465–476, 1991. DOI: 10.1016/0370-2693(91)90202-2.

[141] H. Abramowicz and A. Levy, "The ALLM parameterization of σ_{tot} ($\gamma^* p$): an update," 1997. arXiv: hep-ph/9712415 [hep-ph].

[142] S. I. Dutta, M. H. Reno, I. Sarcevic, and D. Seckel, "Propagation of muons and taus at high-energies," *Phys. Rev.*, vol. D63, p. 094 020, 2001. DOI: 10.1103/PhysRevD.63.094020. arXiv: hep-ph/0012350 [hep-ph].

[143] L. B. Bezrukov and E. V. Bugaev, "Nucleon shadowing effects in photon nucleus interaction," *Yad. Fiz.*, vol. 33, pp. 1195–1207, 1981, [Sov. J. Nucl. Phys.33,635(1981)].

[144] E. V. Bugaev and Yu. V. Shlepin, "Photonuclear interaction of high-energy muons and tau leptons," *Phys. Rev.*, vol. D67, p. 034 027, 2003. DOI: 10.1103/PhysRevD.67.034027. arXiv: hep-ph/0203096 [hep-ph].

[145] E. Bugaev, T. Montaruli, Y. Shlepin, and I. A. Sokalski, "Propagation of tau neutrinos and tau leptons through the earth and their detection in underwater / ice neutrino telescopes," *Astropart. Phys.*, vol. 21, pp. 491–509, 2004. DOI: 10.1016/j.astropartphys.2004.03.002. arXiv: hep-ph/0312295 [hep-ph].

[146] T. K. Gaisser, T. Stanev, and S. Tilav, "Cosmic ray energy spectrum from measurements of air showers," *Front. Phys.(Beijing)*, vol. 8, pp. 748–758, 2013. DOI: 10.1007/s11467-013-0319-7. arXiv: 1303.3565 [astro-ph.HE].

[147] J. Abraham *et al.*, "Measurement of the depth of maximum of extensive air showers above 10^{18} eV," *Phys. Rev. Lett.*, vol. 104, p. 091 101, 2010. DOI: 10.1103/PhysRevLett.104.091101. arXiv: 1002.0699 [astro-ph.HE].

[148] A. Aab *et al.*, "Depth of maximum of air-shower profiles at the Pierre Auger Observatory. II. Composition implications," *Phys. Rev.*, vol. D90, no. 12, p. 122 006, 2014. DOI: 10.1103/PhysRevD.90.122006. arXiv: 1409.5083 [astro-ph.HE].

[149] V. I. Zatsepin and N. V. Sokolskaya, "Three component model of cosmic ray spectra from 100-gev up to 100-pev," *Astron. Astrophys.*, vol. 458, pp. 1–5, 2006. DOI: 10.1051/0004-6361:20065108. arXiv: astro-ph/0601475 [astro-ph].

[150] S. Ostapchenko, "Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model," *Phys. Rev.*, vol. D83, p. 014 018, 2011. DOI: 10.1103/PhysRevD.83.014018. arXiv: 1010.1869 [hep-ph].

[151] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner, "EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider," *Phys. Rev.*, vol. C92, no. 3, p. 034 906, 2015. DOI: 10.1103/PhysRevC.92.034906. arXiv: 1306.0121 [hep-ph].

[152] S. Roesler, R. Engel, and J. Ranft, "The Monte Carlo event generator DPMJET-III," in *Advanced Monte Carlo for radiation physics, particle transport simulation and applications. Proceedings, Conference, MC2000, Lisbon, Portugal, October 23-26, 2000*, 2000, pp. 1033–1038. DOI: 10.1007/978-3-642-18211-2_166. arXiv: hep-ph/0012252 [hep-ph]. [Online]. Available: http://www-public.slac.stanford.edu/sciDoc/docMeta.aspx?slacPubNumber=SLAC-PUB-8740.

[153] A. M. Hillas, "Cosmic Rays: Recent Progress and some Current Questions," in *Conference on Cosmology, Galaxy Formation and Astro-Particle Physics on the Pathway to the SKA Oxford, England, April 10-12, 2006*, 2006. arXiv: astro-ph/0607109 [astro-ph].

[154] T. Pierog, "Open issues in hadronic interactions for air showers," *EPJ Web Conf.*, vol. 145, p. 18 002, 2017. DOI: 10.1051/epjconf/201614518002.

[155] J. Blumer, R. Engel, and J. R. Horandel, "Cosmic Rays from the Knee to the Highest Energies," *Prog. Part. Nucl. Phys.*, vol. 63, pp. 293–338, 2009. DOI: 10.1016/j.ppnp.2009.05.002. arXiv: 0904.0725 [astro-ph.HE].

[156] D. Heck, G. Schatz, T. Thouw, J. Knapp, and J. N. Capdevielle, "CORSIKA: A Monte Carlo code to simulate extensive air showers," *FZKA-6019*, 1998. [Online]. Available: https://web.ikp.kit.edu/corsika/physics_description/corsika_phys.pdf.

[157] E.-J. Ahn, R. Engel, T. K. Gaisser, P. Lipari, and T. Stanev, "Cosmic ray interaction event generator SIBYLL 2.1," *Phys. Rev.*, vol. D80, p. 094 003, 2009. DOI: 10.1103/PhysRevD.80.094003. arXiv: 0906.4113 [hep-ph].

[158] J. van Santen, "Neutrino interactions in icecube above 1 tev constraints on atmospheric charmed-meson production and investigation of the astrophysical neutrino flux with 2 years of icecube data taken 2010–2012," Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2016-06-04, PhD thesis, University of Wisconsin at Madison, 2014, p. 188, ISBN: 9781321369502. [Online]. Available: http://search.proquest.com.ezproxy.library.wisc.edu/docview/1637727437?accountid=465.

[159] M. G. Aartsen *et al.*, "First Taste of Quantum Gravity with Neutrino Flavour in IceCube," *Nature Physics*, 201X.

[160] M. G. Aartsen *et al.*, "Searches For Connections Between Dark Matter And High Energy Neutrinos With IceCube," *Phys. Rev.*, vol. D, 201X.

[161] M. G. Aartsen *et al.*, "Measurement of the high-energy all-flavor neutrino-nucleon cross section with IceCube," *Phys. Rev.*, vol. D, no. 6, 201X.

[162] A. Stuart, M. G. Kendall, *et al.*, *The advanced theory of statistics*. Griffin, 1963.

[163] F. James, L. Lyons, Y. Perrin, E. O. for Nuclear Research, and CERN., *Proceedings*, ser. CERN (Series). CERN, European Organization for Nuclear Research, 2000, ISBN: 9789290831655. [Online]. Available: https://books.google.com/books?id=La5JxgEACAAJ.

[164] S. D. Biller and S. M. Oser, "Another Look at Confidence Intervals: Proposal for a More Relevant and Transparent Approach," *Nucl. Instrum. Meth.*, vol. A774, pp. 103–119, 2015. DOI: 10.1016/j.nima.2014.11.081. arXiv: 1405.5010 [physics.data-an].

[165] R. D. Cousins, "Lectures on Statistics in Theory: Prelude to Statistics in Practice," Jul. 2018. arXiv: 1807.05996 [physics.data-an].

[166] R. J. Barlow, "Practical Statistics for Particle Physics," May 2019. arXiv: 1905.12362 [physics.data-an].

[167] A. Diaz, C. A. Argüelles, G. H. Collin, J. M. Conrad, and M. H. Shaevitz, "Where Are We With Light Sterile Neutrinos?," 2019. arXiv: 1906.00045 [hep-ex].

[168] R. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM Journal on Scientific Computing*, vol. 16, no. 5, pp. 1190–1208, 1995. DOI: 10.1137/0916069. eprint: https://doi.org/10.1137/0916069. [Online]. Available: https://doi.org/10.1137/0916069.

[169] S. S. Wilks, "The large-sample distribution of the likelihood ratio for testing composite hypotheses," *Ann. Math. Statist.*, vol. 9, no. 1, pp. 60–62, Mar. 1938. DOI: 10.1214/aoms/1177732360. [Online]. Available: https://doi.org/10.1214/aoms/1177732360.

[170] S. Algeri, J. Aalbers, K. Dundas Morå, and J. Conrad, "Searching for new physics with profile likelihoods: Wilks and beyond," 2019. arXiv: 1911.10237 [physics.data-an].

[171] U. von Toussaint, "Bayesian inference in physics," *Rev. Mod. Phys.*, vol. 83, pp. 943–999, 3 Sep. 2011. DOI: 10.1103/RevModPhys.83.943. [Online]. Available: https://link.aps.org/doi/10.1103/RevModPhys.83.943.

[172] P. S. Laplace, *Théorie analytique des probabilités*. Courcier, 1820.

[173] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, "emcee: The MCMC Hammer," *Publ. Astron. Soc. Pac.*, vol. 125, pp. 306–312, 2013. DOI: 10.1086/670067. arXiv: 1202.3665 [astro-ph.IM].

[174] R. Trotta, "Bayesian Methods in Cosmology," 2017. arXiv: 1701.01467 [astro-ph.CO].

[175] S. Gariazzo, "Constraining power of open likelihoods, made prior-independent," 2019. arXiv: 1910.06646 [astro-ph.CO].

[176] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, "Importance Nested Sampling and the MultiNest Algorithm," 2013. arXiv: 1306.2144 [astro-ph.IM].

[177] H. Jeffreys, *The Theory of Probability*, ser. Oxford Classic Texts in the Physical Sciences. OUP Oxford, 1998, ISBN: 9780191589676. [Online]. Available: https://books.google.com/books?id=vh9Act9rtzQC.

[178] D. Chirkin, "Likelihood description for comparing data with simulation of limited statistics," 2013. arXiv: 1304.0735 [astro-ph.IM].

[179] T. Glüsenkamp, "Probabilistic treatment of the uncertainty from the finite size of weighted Monte Carlo data," *Eur. Phys. J. Plus*, vol. 133, no. 6, p. 218, 2018. DOI: 10.1140/epjp/i2018-12042-x. arXiv: 1712.01293 [physics.data-an].

[180] T. Glüsenkamp, "A unified perspective on modified Poisson likelihoods for limited Monte Carlo data," 2019. arXiv: 1902.08831 [astro-ph.IM].

[181] J. S. Gainer, J. Lykken, K. T. Matchev, S. Mrenna, and M. Park, "Exploring Theory Space with Monte Carlo Reweighting," *JHEP*, vol. 10, p. 078, 2014. DOI: 10.1007/JHEP10(2014)078. arXiv: 1404.7129 [hep-ph].

[182] L. Lyons, *STATISTICS FOR NUCLEAR AND PARTICLE PHYSICISTS*. 1986, ISBN: 9780521379342. [Online]. Available: http://www.cambridge.org/uk/catalogue/catalogue.asp?isbn=0521255406.

[183] S. D. Poisson, *Recherches sur la probabilité des jugements en matière criminelle et en matière civile précédées des règles générales du calcul des probabilités*. Bachelier, 1837.

[184] R. J. Barlow and C. Beeston, "Fitting using finite Monte Carlo samples," *Comput. Phys. Commun.*, vol. 77, pp. 219–228, 1993. DOI: 10.1016/0010-4655(93)90005-W.

[185] K. Cranmer, G. Lewis, L. Moneta, A. Shibata, and W. Verkerke, "HistFactory: A tool for creating statistical models for use with RooFit and RooStats," 2012.

[186] C. Argüelles, A. Schneider, and T. Yuan, *MCLLH*, https://github.com/austinschneider/MCLLH, 2019.

[187] K. Pearson, "On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 50, no. 302, pp. 157–175, 1900.

[188] B. K. Cogswell, D. J. Ernst, K. T. L. Ufheil, J. T. Gaglione, and J. M. Malave, "Neutrino oscillations: ILL experiment revisited," 2018. arXiv: 1802.07763 [hep-ph].

[189] G. Cowan, *Statistical data analysis*. Oxford university press, 1998.

[190] R. D. Cousins and V. L. Highland, "Incorporating systematic uncertainties into an upper limit," *Nucl. Instrum. Meth.*, vol. A320, pp. 331–335, 1992. DOI: 10.1016/0168-9002(92)90794-5.

[191] K. Abe *et al.*, "Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector," *Phys. Rev.*, vol. D96, no. 9, p. 092 006, 2017, [Erratum: Phys. Rev.D98,no.1,019902(2018)]. DOI: 10.1103/PhysRevD.96.092006,10.1103/PhysRevD.98.019902. arXiv: 1707.01048 [hep-ex].

[192] K. Abe *et al.*, "Search for CP Violation in Neutrino and Antineutrino Oscillations by the T2K Experiment with 2.2×10^{21} Protons on Target," *Phys. Rev. Lett.*, vol. 121, no. 17, p. 171 802, 2018. DOI: 10.1103/PhysRevLett.121.171802. arXiv: 1807.07891 [hep-ex].

[193] G. Bohm and G. Zech, "Statistics of weighted Poisson events and its applications," *Nucl. Instrum. Meth.*, vol. A748, pp. 1–6, 2014. DOI: 10.1016/j.nima.2014.02.021. arXiv: 1309.1287 [physics.data-an].

[194] D. Fink, *A compendium of conjugate priors*, 1997.

[195] J. Bernardo and A. Smith, *Bayesian Theory*, ser. Wiley Series in Probability and Statistics. Wiley, 2009, ISBN: 9780470317716. [Online]. Available: https://books.google.com.pe/books?id=11nSgIcd7xQC.

[196] M. G. Aartsen *et al.*, "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A," *Science*, vol. 361, no. 6398, eaat1378, 2018. DOI: 10.1126/science.aat1378. arXiv: 1807.08816 [astro-ph.HE].

[197] M. G. Aartsen *et al.*, "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert," *Science*, vol. 361, no. 6398, pp. 147–151, 2018. DOI: 10.1126/science.aat2890. arXiv: 1807.08794 [astro-ph.HE].

[198] M. G. Aartsen *et al.*, "Time-integrated Neutrino Source Searches with 10 years of IceCube Data," *Phys. Rev. Lett.*, vol. 124, p. 051 103, 2020. DOI: 10.1103/PhysRevLett.124.051103. arXiv: 1910.08488 [astro-ph.HE].

[199] A. Albert *et al.*, "First all-flavor neutrino pointlike source search with the ANTARES neutrino telescope," *Phys. Rev.*, vol. D96, no. 8, p. 082 001, 2017. DOI: 10.1103/PhysRevD.96.082001. arXiv: 1706.01857 [astro-ph.HE].

[200] M. G. Aartsen *et al.*, "Search for steady point-like sources in the astrophysical muon neutrino flux with 8 years of IceCube data," *Eur. Phys. J.*, vol. C79, no. 3, p. 234, 2019. DOI: 10.1140/epjc/s10052-019-6680-0. arXiv: 1811.07979 [hep-ph].

[201] M. G. Aartsen *et al.*, "Search for Sources of Astrophysical Neutrinos Using Seven Years of IceCube Cascade Events," 2019. arXiv: 1907.06714 [astro-ph.HE].

[202] IceCube, ***HESE Data Release***, https://github.com/IceCubeOpenSource/HESE_data_release, 2020.

[203] V. Barger, L. Fu, J. G. Learned, D. Marfatia, S. Pakvasa, and T. J. Weiler, "Glashow resonance as a window into cosmic neutrino sources," *Phys. Rev.*, vol. D90, p. 121 301, 2014. DOI: 10.1103/PhysRevD.90.121301. arXiv: 1407.3255 [astro-ph.HE].

[204] A. Garcia, R. Gauld, A. Heijboer, and J. Rojo, "Complete predictions for high-energy neutrino propagation in matter," Apr. 2020. arXiv: 2004.04756 [hep-ph].

[205] M. G. Aartsen *et al.*, "Measurements using the inelasticity distribution of multi-TeV neutrino interactions in IceCube," *Phys. Rev.*, vol. D99, no. 3, p. 032 004, 2019. DOI: 10.1103/PhysRevD.99.032004. arXiv: 1808.07629 [hep-ex].

[206] T. Gaisser, R. Engel, and E. Resconi, *Cosmic Rays and Particle Physics*. Cambridge University Press, 2016, ISBN: 9781316598436. [Online]. Available: https://books.google.com/books?id=nHFNDAAAQBAJ.

[207] E. Fermi, "On the Origin of the Cosmic Radiation," *Phys. Rev.*, vol. 75, pp. 1169–1174, 1949. DOI: 10.1103/PhysRev.75.1169.

[208] H. Nunokawa, B. Panes, and R. Zukanovich Funchal, "How Unequal Fluxes of High Energy Astrophysical Neutrinos and Antineutrinos can Fake New Physics," *JCAP*, vol. 1610, no. 10, p. 036, 2016. DOI: 10.1088/1475-7516/2016/10/036. arXiv: 1604.08595 [hep-ph].

[209] I. M. Shoemaker and K. Murase, "Probing BSM Neutrino Physics with Flavor and Spectral Distortions: Prospects for Future High-Energy Neutrino Telescopes," *Phys. Rev.*, vol. D93, no. 8, p. 085 004, 2016. DOI: 10.1103/PhysRevD.93.085004. arXiv: 1512.07228 [astro-ph.HE].

[210] M. Ajello *et al.*, "3FHL: The Third Catalog of Hard Fermi-LAT Sources," *Astrophys. J. Suppl.*, vol. 232, no. 2, p. 18, 2017. DOI: 10.3847/1538-4365/aa8221. arXiv: 1702.00664 [astro-ph.HE].

[211] H. Niederhausen, "Measurement of the high energy astrophysical neutrino flux using electron and tau neutrinos observed in four years of icecube data," PhD thesis, Stony Brook University, 2018.

[212] M. Aartsen *et al.*, "Measurement of the Atmospheric ν_e flux in IceCube," *Phys. Rev. Lett.*, vol. 110, no. 15, p. 151 105, 2013. DOI: 10.1103/PhysRevLett.110.151105. arXiv: 1212.4760 [hep-ex].

[213] K. Murase, D. Guetta, and M. Ahlers, "Hidden Cosmic-Ray Accelerators as an Origin of TeV-PeV Cosmic Neutrinos," *Phys. Rev. Lett.*, vol. 116, no. 7, p. 071 101, 2016. DOI: 10.1103/PhysRevLett.116.071101. arXiv: 1509.00805 [astro-ph.HE].

[214] S. Ando, I. Tamborra, and F. Zandanel, "Tomographic Constraints on High-Energy Neutrinos of Hadronuclear Origin," *Phys. Rev. Lett.*, vol. 115, no. 22, p. 221 101, 2015. DOI: 10.1103/PhysRevLett.115.221101. arXiv: 1509.02444 [astro-ph.HE].

[215] K. Bechtol, M. Ahlers, M. Di Mauro, M. Ajello, and J. Vandenbroucke, "Evidence against star-forming galaxies as the dominant source of IceCube neutrinos," *Astrophys. J.*, vol. 836, no. 1, p. 47, 2017. DOI: 10.3847/1538-4357/836/1/47. arXiv: 1511.00688 [astro-ph.HE].

[216] P. Meszaros, "Astrophysical Sources of High Energy Neutrinos in the IceCube Era," *Ann. Rev. Nucl. Part. Sci.*, vol. 67, pp. 45–67, 2017. DOI: 10.1146/annurev-nucl-101916-123304. arXiv: 1708.03577 [astro-ph.HE].

[217] A. Capanema, A. Esmaili, and K. Murase, "New constraints on the origin of medium-energy neutrinos observed by IceCube," *Phys. Rev. D*, vol. 101, no. 10, p. 103 012, 2020. DOI: 10.1103/PhysRevD.101.103012. arXiv: 2002.07192 [hep-ph].

[218] F. Halzen and A. Kheirandish, "Multimessenger search for the sources of cosmic rays using cosmic neutrinos," *Frontiers in Astronomy and Space Sciences*, vol. 6, p. 32, 2019, ISSN: 2296-987X. DOI: 10.3389/fspas.2019.00032. [Online]. Available: https://www.frontiersin.org/article/10.3389/fspas.2019.00032.

[219] V. S. Berezinsky and A. Yu. Smirnov, "Cosmic neutrinos of ultra-high energies and detection possibility," *Astrophys. Space Sci.*, vol. 32, pp. 461–482, 1975. DOI: 10.1007/BF00643157.

[220] K. Murase, M. Ahlers, and B. C. Lacki, "Testing the Hadronuclear Origin of PeV Neutrinos Observed with IceCube," *Phys. Rev.*, vol. D88, no. 12, p. 121 301, 2013. DOI: 10.1103/PhysRevD.88.121301. arXiv: 1306.3417 [astro-ph.HE].

[221] M. V. Garzelli, S. Moch, and G. Sigl, "Lepton fluxes from atmospheric charm revisited," *JHEP*, vol. 10, p. 115, 2015. DOI: 10.1007/JHEP10(2015)115. arXiv: 1507.01570 [hep-ph].

[222] R. Gauld, J. Rojo, L. Rottoli, S. Sarkar, and J. Talbert, "The prompt atmospheric neutrino flux in the light of LHCb," *JHEP*, vol. 02, p. 130, 2016. DOI: 10.1007/JHEP02(2016)130. arXiv: 1511.06346 [hep-ph].

[223] R. Gauld, J. Rojo, L. Rottoli, and J. Talbert, "Charm production in the forward region: constraints on the small-x gluon and backgrounds for neutrino astronomy," *JHEP*, vol. 11, p. 009, 2015. DOI: 10.1007/JHEP11(2015)009. arXiv: 1506.08025 [hep-ph].

[224] A. Bhattacharya, R. Enberg, Y. S. Jeong, C. S. Kim, M. H. Reno, I. Sarcevic, and A. Stasto, "Prompt atmospheric neutrino fluxes: perturbative QCD models and nuclear effects," *JHEP*, vol. 11, p. 167, 2016. DOI: 10.1007/JHEP11(2016)167. arXiv: 1607.00193 [hep-ph].

[225] M. V. Garzelli, S. Moch, O. Zenaiev, A. Cooper-Sarkar, A. Geiser, K. Lipka, R. Placakyte, and G. Sigl, "Prompt neutrino fluxes in the atmosphere with PROSA parton distribution functions," *JHEP*, vol. 05, p. 004, 2017. DOI: 10.1007/JHEP05(2017)004. arXiv: 1611.03815 [hep-ph].

[226] V. P. Goncalves and M. V. T. Machado, "Saturation Physics in Ultra High Energy Cosmic Rays: Heavy Quark Production," *JHEP*, vol. 04, p. 028, 2007. DOI: 10.1088/1126-6708/2007/04/028. arXiv: hep-ph/0607125 [hep-ph].

[227] R. Enberg, M. H. Reno, and I. Sarcevic, "Prompt neutrino fluxes from atmospheric charm," *Phys. Rev.*, vol. D78, p. 043 005, 2008. DOI: 10.1103/PhysRevD.78.043005. arXiv: 0806.0418 [hep-ph].

[228] B. Abelev *et al.*, "Measurement of charm production at central rapidity in proton-proton collisions at $\sqrt{s} = 2.76$ TeV," *JHEP*, vol. 07, p. 191, 2012. DOI: 10.1007/JHEP07(2012)191. arXiv: 1205.4007 [hep-ex].

[229] O. Zenaiev, M. Garzelli, K. Lipka, S. Moch, A. Cooper-Sarkar, F. Olness, A. Geiser, and G. Sigl, "Improved constraints on parton distributions using LHCb, ALICE and HERA heavy-flavour measurements and implications for the predictions for prompt atmospheric-neutrino fluxes," *JHEP*, vol. 04, p. 118, 2020. DOI: 10.1007/JHEP04(2020)118. arXiv: 1911.13164 [hep-ph].

[230] R. Laha and S. J. Brodsky, "IceCube can constrain the intrinsic charm of the proton," *Phys. Rev.*, vol. D96, no. 12, p. 123 002, 2017. DOI: 10.1103/PhysRevD.96.123002. arXiv: 1607.08240 [hep-ph].

[231] M. G. Aartsen *et al.*, "Search for a diffuse flux of astrophysical muon neutrinos with the IceCube 59-string configuration," *Phys. Rev.*, vol. D89, no. 6, p. 062 007, 2014. DOI: 10.1103/PhysRevD.89.062007. arXiv: 1311.7048 [astro-ph.HE].

[232] F. Halzen, "Tev to Eev diffuse gamma-rays and neutrinos," in *First Palaiseau Workshop: Towards a Major Atmospheric Cerenkov Detector for TeV Astro/particle Physics Palaiseau, France, June 11-12, 1992*, 1992, pp. 85–102. [Online]. Available: http://lss.fnal.gov/cgi-bin/find_paper.pl?other/mad-ph-714.pdf.

[233] M. G. Aartsen *et al.*, "Differential limit on the extremely-high-energy cosmic neutrino flux in the presence of astrophysical background from nine years of IceCube data," *Phys. Rev.*, vol. D98, no. 6, p. 062 003, 2018. DOI: 10.1103/PhysRevD.98.062003. arXiv: 1807.01820 [astro-ph.HE].

[234] I. Safa, A. Pizzuto, C. A. Argüelles, F. Halzen, R. Hussain, A. Kheirandish, and J. Vandenbroucke, "Observing EeV neutrinos through the Earth: GZK and the anomalous ANITA events," 2019. arXiv: 1909.10487 [hep-ph].

[235] F. W. Stecker, "PeV neutrinos observed by IceCube from cores of active galactic nuclei," *Phys. Rev.*, vol. D88, no. 4, p. 047 301, 2013. DOI: 10.1103/PhysRevD.88.047301. arXiv: 1305.7404 [astro-ph.HE].

[236] K. Fang and K. Murase, "Linking High-Energy Cosmic Particles by Black Hole Jets Embedded in Large-Scale Structures," *Phys. Lett.*, vol. 14, p. 396, 2018, [Nature Phys.14,no.4,396(2018)]. DOI: 10.1038/s41567-017-0025-4. arXiv: 1704.00015 [astro-ph.HE].

[237] S. S. Kimura, K. Murase, and K. Toma, "Neutrino and Cosmic-Ray Emission and Cumulative Background from Radiatively Inefficient Accretion Flows in Low-Luminosity Active Galactic Nuclei," *Astrophys. J.*, vol. 806, p. 159, 2015. DOI: 10.1088/0004-637X/806/2/159. arXiv: 1411.3588 [astro-ph.HE].

[238] P. Padovani, M. Petropoulou, P. Giommi, and E. Resconi, "A simplified view of blazars: the neutrino background," *Mon. Not. Roy. Astron. Soc.*, vol. 452, no. 2, pp. 1877–1887, 2015. DOI: 10.1093/mnras/stv1467. arXiv: 1506.09135 [astro-ph.HE].

[239] N. Senno, K. Murase, and P. Meszaros, "Choked Jets and Low-Luminosity Gamma-Ray Bursts as Hidden Neutrino Sources," *Phys. Rev.*, vol. D93, no. 8, p. 083 003, 2016. DOI: 10.1103/PhysRevD.93.083003. arXiv: 1512.08513 [astro-ph.HE].

[240] I. Bartos and S. Marka, "Spectral Decline of PeV Neutrinos from Starburst Galaxies," 2015. arXiv: 1509.00983 [astro-ph.HE].

[241] F. Tavecchio and G. Ghisellini, "High-energy cosmic neutrinos from spine-sheath BL Lac jets," *Mon. Not. Roy. Astron. Soc.*, vol. 451, no. 2, pp. 1502–1510, 2015. DOI: 10.1093/mnras/stv1023. arXiv: 1411.2783 [astro-ph.HE].

[242] D. Biehl, D. Boncioli, C. Lunardini, and W. Winter, "Tidally disrupted stars as a possible origin of both cosmic rays and neutrinos at the highest energies," *Sci. Rep.*, vol. 8, no. 1, p. 10 828, 2018. DOI: 10.1038/s41598-018-29022-4. arXiv: 1711.03555 [astro-ph.HE].

[243] M. G. Aartsen *et al.*, "Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube," *JCAP*, vol. 1910, no. 10, p. 048, 2019. DOI: 10.1088/1475-7516/2019/10/048. arXiv: 1909.01530 [hep-ex].

[244] D. Chirkin, "Event reconstruction in IceCube based on direct event re-simulation," in *Proceedings, 33rd International Cosmic Ray Conference (ICRC2013): Rio de Janeiro, Brazil, July 2-9, 2013*, p. 0581. [Online]. Available: http://www.cbpf.br/%7Eicrc2013/papers/icrc2013-0581.pdf.

[245] P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré, "Markov chain monte carlo without likelihoods," *Proceedings of the National Academy of Sciences*, vol. 100, no. 26, pp. 15 324–15 328, 2003, ISSN: 0027-8424. DOI: 10.1073/pnas.0306899100. eprint: https://www.pnas.org/content/100/26/15324.full.pdf. [Online]. Available: https://www.pnas.org/content/100/26/15324.

[246] T. Yuan, "The 8-parameter Fisher-Bingham distribution on the sphere," 2019. arXiv: 1906.08247 [physics.data-an].

[247] J. Braun, J. Dumm, F. De Palma, C. Finley, A. Karle, and T. Montaruli, "Methods for point source analysis in high energy neutrino telescopes," *Astropart. Phys.*, vol. 29, pp. 299–305, 2008. DOI: 10.1016/j.astropartphys.2008.02.007. arXiv: 0801.1604 [astro-ph].

[248] M. G. Aartsen *et al.*, "All-sky Search for Time-integrated Neutrino Emission from Astrophysical Sources with 7 yr of IceCube Data," *Astrophys. J.*, vol. 835, no. 2, p. 151, 2017. DOI: 10.3847/1538-4357/835/2/151. arXiv: 1609.04981 [astro-ph.HE].

[249] M. Ackermann, M. Ajello, W. B. Atwood, L. Baldini, J. Ballet, G. Barbiellini, D. Bastieri, K. Bechtol, R. Bellazzini, B. Berenji, R. D. Blandford, E. D. Bloom, E. Bonamente, A. W. Borgland, T. J. Brandt, J. Bregeon, *et al.*, "Fermi-LAT Observations of the Diffuse γ-Ray Emission: Implications for Cosmic Rays and the Interstellar Medium," *ApJ*, vol. 750, 3, p. 3, May 2012. DOI: 10.1088/0004-637X/750/1/3. arXiv: 1202.4039 [astro-ph.HE].

[250] M. G. Aartsen *et al.*, "Constraints on Galactic Neutrino Emission with Seven Years of IceCube Data," *Astrophys. J.*, vol. 849, no. 1, p. 67, 2017. DOI: 10.3847/1538-4357/aa8dfb. arXiv: 1707.03416 [astro-ph.HE].

Appendix A

HESE source searches

A.1 High-energy astrophysical neutrino source searches

A.1.1 Directional reconstruction for neutrino source searches

In the search for neutrino sources, a conservative choice is made to model the background with data-derived distributions in an effort to avoid potential bias from mismodeling of the backgrounds. As Monte-Carlo (MC) simulation is not used in this approach, we are able to use a more accurate but computationally more expensive directional reconstruction for these analyses, that would otherwise be prohibitive to use on an MC sample. Events in data are reconstructed by fully re-simulating cascades or tracks with in-ice photon propagation [244]. A re-simulation of the light propagation is performed for each considered combination of direction and interaction vertex. This re-simulated information is re-used when considering different times and energies for the event. We first perform a localized random search in the reconstruction parameter space to find the minimum of the test-statistic described in [178], comparing re-simulated waveforms to data. This provides a localized starting point in the high-dimensional parameter space for efficiently estimating the directional posterior distribution. An "approximate Bayesian computation" (ABC) [245] is performed around the estimated minimum using the same test-statistic in order to sample the posterior distribution of the reconstruction parameters; assuming a uniform prior in direction and vertex position. This method allows us to sample the test-statistic posterior distribution despite the non-deterministic nature of the re-simulation step. After the ABC step, we marginalize over other parameters to obtain the posterior distribution of the event direction. We then parameterize the directional posterior distribution with an eight-parameter Fisher-Bingham distribution [246], which is used in the analyses described later in this section.

The errors reported for these events with this reconstruction method are nearly twice the angular size

on average than those quoted previously for the same events. This stems from the different treatment of uncertainties. Previously, re-simulations of the events used variations of the scattering and absorption for individual ice layers. This previously used method accounts for statistical variations and uncertainties in the ice layers, but neglects other known detector uncertainties. The reconstruction now used for the source searches, on the other hand, explicitly includes calibration uncertainty on the simulated distribution of observed charge in its test-statistic. This new method of introducing the detector uncertainties covers the known uncertainties in the photon arrival distribution from calibration and is more conservative than the method previously used. This change in the treatment of uncertainty results in larger angular errors which are consistent with other more complete treatments of the calibration uncertainty [243]. Therefore, the previously reported errors were likely underestimated. This ice model used in the new reconstruction also incorporates more calibration information than the model used previously, a change that improves the accuracy of the reconstructed direction for cascades. The inferred directional probability distribution of some events has changed substantially with respect to what was inferred with the original reconstruction. The extent of these changes is expected based on studies of simulated events, given that the previous errors undercover, and the modeling of the ice properties used for reconstruction has been improved significantly. We note that the directional information inferred from cascade events, particularly in the azimuthal dimension of the detector, is highly dependent on the modeling of the ice which continues to improve. The ice model used for the source searches reconstruction in this work adds information about the azimuthal anisotropy of the scattering and absorption of photons based on [104]. Ice models in development will eventually add information about the birefringent crystalline structure of the ice [105], changing again what we infer about the cascade directional probability distributions.

Except four events classified as double cascades, data are reconstructed based on the classification described above. Since algorithms to estimate the double cascade angular uncertainties were not available, the four double cascades are reconstructed assuming a cascade hypothesis for the source searches.

A.1.2 Neutrino source searches

In a simple, spatial-only, analysis we searched for clustering consistent with either point-like or spatially-extended galactic emission. In each case, we used an unbinned likelihood function, described in [247], given by

$$\mathcal{L}(n_s; \vec{\psi}_s) = \prod_i^N \left[\frac{n_s}{N} \cdot S(\vec{x}_i) + \left(1 - \frac{n_s}{N}\right) \cdot B \right], \tag{A.1}$$

where N is the total number of observed events, n_s is the expected number of signal events, $\vec{\psi}_s$ is the source position in the sky, $\vec{x}_i$ represents the properties of event i, $B = 1/(4\pi)$ is the background spatial

distribution which we take to be uniform, and $S(\vec{x}_i) = P_i(\vec{\psi}_s)$ is the spatial distribution expected from the signal population which is taken to be equal to the posterior density of the event direction $P_i(\vec{x}_s)$ at the source position. We neglect the energy information in this search. For point-like sources, $S(\vec{x}_i)$ is determined entirely by the posterior distributions described in Chapter 5; for spatially extended emission, those must be convolved with the assumed spatial distribution of the source. In order to quantify preference for the alternative (presence of a neutrino source) versus the null (isotropic scenario) hypothesis, we use the following test statistic

$$\mathrm{TS} = -2\ln[\mathcal{L}(n_s = 0)/\mathcal{L}(\hat{n}_s)], \tag{A.2}$$

where $\hat{n}_s$ is the value that maximizes the likelihood. The null-hypothesis p-value is determined by comparing the observed test statistical value, $\mathrm{TS}_{\mathtt{obs}}$, with the TS distribution expected given background only, and correcting for the number of source hypotheses considered. This background TS distribution is determined by repeating the experiment using modified versions of the dataset where we have randomized the right ascension of all events.

The results of testing for point-like emission in the sky from many directions are shown in Fig. A.1. The maximum test statistic is found at equatorial coordinates $(\alpha, \delta) = (342.1^\circ, 1.3^\circ)$, with a null-hypothesis post-trials p-value of 0.092. The hottest spot no longer correlates with the galactic plane as was the case in previous analyses. In an unweighted hottest-spot test restricted to a predefined source list consisting of 74 source candidates – which has been studied in previous iterations of this analysis [60], with through-going tracks [248], and with contained cascades [201] – no significant emission was found; the null-hypothesis post-trials p-value is 0.76. This list of 74 sources is enumerated in the supplementary material of [60].

We performed two searches for spatially-extended Galactic emission. First, as in past work [60, 61, 62], we tested for uniform emission within a Galactic plane region expressed in terms of Galactic latitude as $|b| < b_{\mathrm{max}}$, scanning over possible values $b_{\mathrm{max}} \in [2.5^\circ, 30^\circ]$. Here, the greatest pre-trial significance ($\mathrm{p} = 0.02$) was found for $b_{\mathrm{max}} = 5^\circ$. After accounting for multiple, partially correlated tests over the range of the b_{max} scan, we find a post-trials p-value of 0.062.

Fermi-LAT has inferred a diffuse flux of gamma rays originating in decays of neutral pions produced by interactions between cosmic rays and Galactic gas and dust [249]. Although this inference is based on measurements made at energies far below the threshold of this analysis ($E_\gamma \lesssim 100\,\mathrm{GeV}$), it is possible that this gamma ray flux extends to higher energies and is accompanied by high energy neutrinos produced in decays of charged pions. Therefore, in a second, more model-dependent search similar to that performed in [201, 250], we test for emission following the spatial profile of the *Fermi*-LAT best-fit result on Galactic hadronic emission below 100 GeV, using the method described in [250]. Here we find a p-value of 0.089.

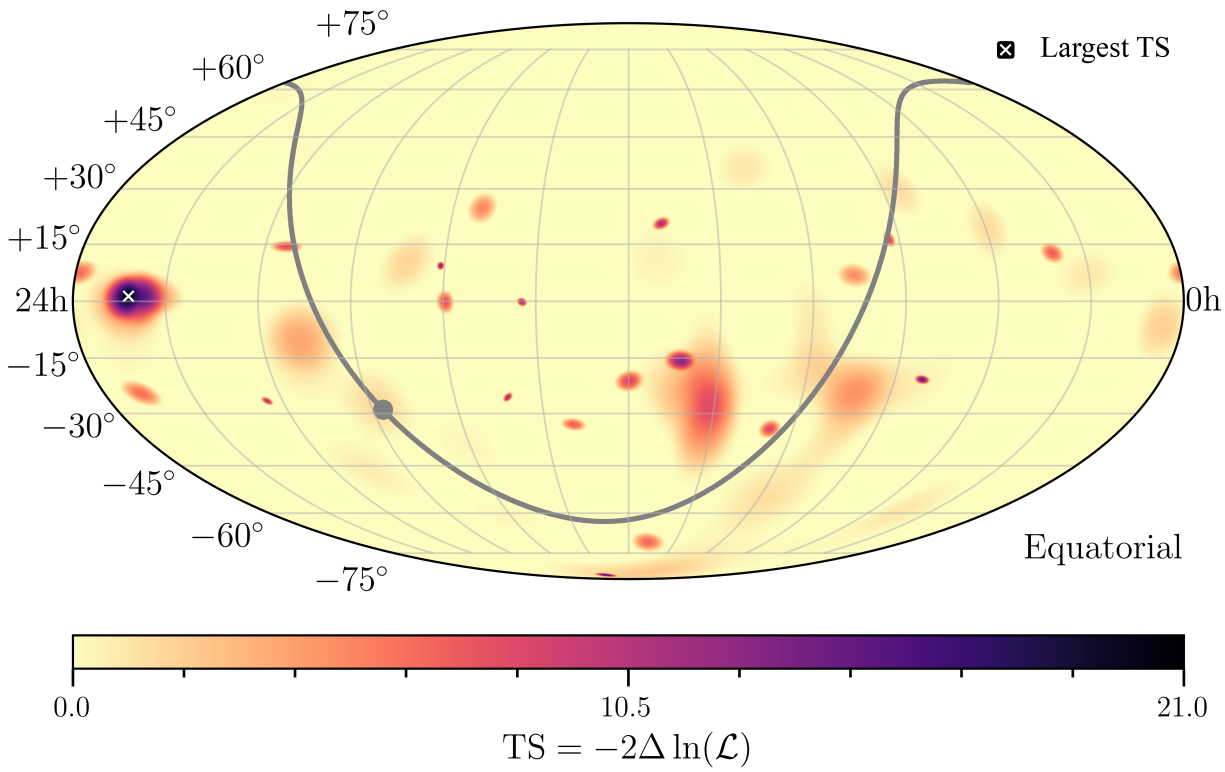

Figure A.1: ***Point source*** TS ***map.*** The TS at each point on the sky is indicated by the blue color scale. The Galactic center and Galactic plane are indicated by the gray dot and gray curve, respectively. Unlike in previous analyses [59, 60, 61, 62], the most significant position is well separated from the Galactic plane, at $(\alpha, \delta) = (342.1°, 1.3°)$, indicated by the white $\times$.

To conclude, the new reconstruction techniques used for this analysis are more conservative, and generally increase the reported angular error, a detailed accounting of these changes is found in Section A.2. With these searches, we find no evidence of clustering in the sample or correlation with tested sources. Other, more sensitive, searches for neutrino sources have been performed with IceCube data and provide better constraints on the scenarios they test [198, 201].

A.2 Event comparison

Event ID	HPD Area (7.5yr) [sr]	HPD Area (prev.) [sr]	HPD Overlap Area [sr]
1	50 %: 0.564	50 %: 0.253	50 %: 0.116
	90 %: 1.71	90 %: 0.826	90 %: 0.584

2	50 %: 0.940	50 %: 0.605	50 %: 0.226
	90 %: 3.02	90 %: 2.75	90 %: 1.17
3	50 %: 1.26×10^{-3}	50 %: 1.81×10^{-3}	50 %: 0.00
	90 %: 4.19×10^{-3}	90 %: 6.06×10^{-3}	90 %: 7.31×10^{-4}
4	50 %: 0.0421	50 %: 0.0481	50 %: 0.00
	90 %: 0.141	90 %: 0.160	90 %: 0.0413
5	-	50 %: 1.33×10^{-3}	-
		90 %: 4.43×10^{-3}	
6	-	50 %: 0.0913	-
		90 %: 0.302	
7	50 %: 1.10	50 %: 0.545	50 %: 0.00
	90 %: 3.18	90 %: 1.76	90 %: 0.195
8	50 %: 1.87×10^{-3}	50 %: 1.65×10^{-3}	50 %: 3.64×10^{-4}
	90 %: 6.22×10^{-3}	90 %: 5.51×10^{-3}	90 %: 3.09×10^{-3}
9	50 %: 1.93	50 %: 0.259	50 %: 0.259
	90 %: 5.78	90 %: 0.848	90 %: 0.848
10	50 %: 0.0308	50 %: 0.0624	50 %: 0.00
	90 %: 0.102	90 %: 0.206	90 %: 4.31×10^{-3}
11	50 %: 0.345	50 %: 0.264	50 %: 0.0179
	90 %: 1.12	90 %: 0.863	90 %: 0.357
12	50 %: 0.0881	50 %: 0.0920	50 %: 0.0429
	90 %: 0.296	90 %: 0.304	90 %: 0.219
13	50 %: 0.0269	50 %: 1.36×10^{-3}	50 %: 0.00
	90 %: 0.0893	90 %: 4.51×10^{-3}	90 %: 0.00
14	50 %: 0.0445	50 %: 0.167	50 %: 0.0445
	90 %: 0.148	90 %: 0.547	90 %: 0.148

15	50 %: 0.219	50 %: 0.367	50 %: 0.0242
	90 %: 0.536	90 %: 1.19	90 %: 0.290
16	50 %: 0.787	50 %: 0.358	50 %: 0.109
	90 %: 2.20	90 %: 1.18	90 %: 0.489
17	50 %: 0.0309	50 %: 0.128	50 %: 0.0291
	90 %: 0.103	90 %: 0.421	90 %: 0.103
18	50 %: 3.58×10^{-3}	50 %: 1.65×10^{-3}	50 %: 1.17×10^{-3}
	90 %: 0.0119	90 %: 5.49×10^{-3}	90 %: 5.39×10^{-3}
19	50 %: 0.265	50 %: 0.0897	50 %: 0.0346
	90 %: 0.887	90 %: 0.297	90 %: 0.262
20	50 %: 0.0440	50 %: 0.109	50 %: 0.0169
	90 %: 0.148	90 %: 0.360	90 %: 0.125
21	50 %: 0.990	50 %: 0.414	50 %: 0.249
	90 %: 2.72	90 %: 1.37	90 %: 0.785
22	50 %: 0.0367	50 %: 0.140	50 %: 7.09×10^{-3}
	90 %: 0.123	90 %: 0.460	90 %: 0.108
23	50 %: 1.69	50 %: 3.58×10^{-3}	50 %: 3.58×10^{-3}
	90 %: 5.65	90 %: 0.0135	90 %: 0.0135
24	50 %: 0.202	50 %: 0.229	50 %: 0.0759
	90 %: 0.660	90 %: 0.751	90 %: 0.422
25	50 %: 1.66	50 %: 1.94	50 %: 0.362
	90 %: 4.59	90 %: 5.67	90 %: 3.24
26	50 %: 0.279	50 %: 0.134	50 %: 0.0886
	90 %: 0.915	90 %: 0.442	90 %: 0.408
27	50 %: 0.0785	50 %: 0.0419	50 %: 0.0148
	90 %: 0.259	90 %: 0.139	90 %: 0.0985

28	50 %: 2.32×10^{-3}	50 %: 1.61×10^{-3}	50 %: 7.19×10^{-4}
	90 %: 7.72×10^{-3}	90 %: 5.36×10^{-3}	90 %: 4.20×10^{-3}
29	50 %: 0.0710	50 %: 0.0522	50 %: 0.0458
	90 %: 0.235	90 %: 0.173	90 %: 0.164
30	50 %: 0.0739	50 %: 0.0603	50 %: 0.0418
	90 %: 0.247	90 %: 0.200	90 %: 0.181
31	50 %: 1.61	50 %: 0.635	50 %: 0.307
	90 %: 4.09	90 %: 2.71	90 %: 2.39
32	-	-	-
33	50 %: 0.0126	50 %: 0.173	50 %: 0.00
	90 %: 0.0419	90 %: 0.568	90 %: 0.00
34	50 %: 2.24	50 %: 1.66	50 %: 0.733
	90 %: 5.55	90 %: 5.13	90 %: 3.89
35	50 %: 0.0565	50 %: 0.239	50 %: 0.0201
	90 %: 0.190	90 %: 0.781	90 %: 0.166
36	50 %: 0.427	50 %: 0.130	50 %: 3.20×10^{-5}
	90 %: 1.17	90 %: 0.428	90 %: 0.131
37	50 %: 5.43×10^{-4}	50 %: 1.43×10^{-3}	50 %: 2.88×10^{-4}
	90 %: 1.84×10^{-3}	90 %: 4.75×10^{-3}	90 %: 1.61×10^{-3}
38	50 %: 5.59×10^{-4}	50 %: 1.33×10^{-3}	50 %: 0.00
	90 %: 1.85×10^{-3}	90 %: 4.41×10^{-3}	90 %: 6.99×10^{-4}
39	50 %: 0.218	50 %: 0.192	50 %: 0.0151
	90 %: 0.655	90 %: 0.639	90 %: 0.261
40	50 %: 0.0852	50 %: 0.130	50 %: 0.0152
	90 %: 0.292	90 %: 0.429	90 %: 0.174
41	50 %: 0.103	50 %: 0.118	50 %: 0.0527
	90 %: 0.342	90 %: 0.389	90 %: 0.257

42	-	50 %: 0.405 90 %: 1.39	-
43	50 %: 0.842 90 %: 2.38	50 %: 1.55×10^{-3} 90 %: 5.10×10^{-3}	50 %: 1.55×10^{-3} 90 %: 5.10×10^{-3}
44	50 %: 9.75×10^{-3} 90 %: 0.0324	50 %: 1.41×10^{-3} 90 %: 4.67×10^{-3}	50 %: 1.40×10^{-3} 90 %: 4.67×10^{-3}
45	50 %: 1.44×10^{-4} 90 %: 4.79×10^{-4}	50 %: 1.33×10^{-3} 90 %: 4.42×10^{-3}	50 %: 1.44×10^{-4} 90 %: 4.79×10^{-4}
46	50 %: 0.0461 90 %: 0.152	50 %: 0.0551 90 %: 0.183	50 %: 0.0261 90 %: 0.113
47	50 %: 0.171 90 %: 0.555	50 %: 1.41×10^{-3} 90 %: 4.67×10^{-3}	50 %: 0.00 90 %: 1.97×10^{-3}
48	50 %: 0.0923 90 %: 0.306	50 %: 0.0629 90 %: 0.208	50 %: 0.00 90 %: 0.0156
49	50 %: 1.18 90 %: 3.24	50 %: 0.450 90 %: 1.54	50 %: 0.301 90 %: 0.992
50	50 %: 0.806 90 %: 3.22	50 %: 0.0633 90 %: 0.210	50 %: 0.0633 90 %: 0.210
51	50 %: 0.448 90 %: 1.73	50 %: 0.0401 90 %: 0.133	50 %: 0.0401 90 %: 0.133
52	50 %: 0.113 90 %: 0.333	50 %: 0.0578 90 %: 0.192	50 %: 0.0190 90 %: 0.124
53	-	50 %: 1.37×10^{-3} 90 %: 4.56×10^{-3}	-
54	50 %: 0.0916 90 %: 0.311	50 %: 0.128 90 %: 0.422	50 %: 0.0678 90 %: 0.276
55	-	-	-

56	50 %: 0.0797	50 %: 0.0408	50 %: 0.0408
	90 %: 0.270	90 %: 0.135	90 %: 0.135
57	50 %: 0.574	50 %: 0.197	50 %: 0.146
	90 %: 1.93	90 %: 0.651	90 %: 0.648
58	50 %: 0.0500	50 %: 1.72×10^{-3}	50 %: 1.72×10^{-3}
	90 %: 0.168	90 %: 5.73×10^{-3}	90 %: 5.73×10^{-3}
59	50 %: 0.0273	50 %: 0.0732	50 %: 0.00
	90 %: 0.0908	90 %: 0.242	90 %: 0.0131
60	50 %: 0.549	50 %: 0.168	50 %: 0.138
	90 %: 1.76	90 %: 0.551	90 %: 0.551
61	-	50 %: 1.34×10^{-3}	-
		90 %: 4.45×10^{-3}	
62	-	50 %: 1.56×10^{-3}	-
		90 %: 5.16×10^{-3}	
63	-	50 %: 1.34×10^{-3}	-
		90 %: 4.46×10^{-3}	
64	50 %: 0.150	50 %: 0.107	50 %: 0.0417
	90 %: 0.515	90 %: 0.354	90 %: 0.246
65	50 %: 2.09	50 %: 0.290	50 %: 0.117
	90 %: 6.55	90 %: 0.951	90 %: 0.808
66	50 %: 0.446	50 %: 0.317	50 %: 5.85×10^{-3}
	90 %: 1.25	90 %: 1.05	90 %: 0.338
67	50 %: 0.0431	50 %: 0.0473	50 %: 0.0168
	90 %: 0.145	90 %: 0.157	90 %: 0.0918
68	50 %: 0.0761	50 %: 0.131	50 %: 0.0742
	90 %: 0.256	90 %: 0.433	90 %: 0.254

69	-	50 %: 0.234 90 %: 0.762	-
70	50 %: 0.0992 90 %: 0.332	50 %: 0.144 90 %: 0.475	50 %: 0.0315 90 %: 0.218
71	50 %: 2.72×10^{-4} 90 %: 9.11×10^{-4}	50 %: 1.35×10^{-3} 90 %: 4.48×10^{-3}	50 %: 2.72×10^{-4} 90 %: 9.11×10^{-4}
72	50 %: 0.442 90 %: 1.81	50 %: 0.360 90 %: 1.10	50 %: 0.229 90 %: 0.612
73	-	50 %: 0.0456 90 %: 0.151	-
74	50 %: 4.75×10^{-3} 90 %: 0.0158	50 %: 0.153 90 %: 0.504	50 %: 4.75×10^{-3} 90 %: 0.0158
75	50 %: 0.514 90 %: 1.80	50 %: 0.163 90 %: 0.515	50 %: 0.163 90 %: 0.515
76	50 %: 1.41×10^{-3} 90 %: 4.70×10^{-3}	50 %: 1.33×10^{-3} 90 %: 4.43×10^{-3}	50 %: 6.90×10^{-4} 90 %: 3.30×10^{-3}
77	50 %: 0.193 90 %: 0.659	50 %: 0.0490 90 %: 0.163	50 %: 0.00 90 %: 0.0634
78	50 %: 3.28×10^{-3} 90 %: 0.0109	50 %: 1.48×10^{-3} 90 %: 4.92×10^{-3}	50 %: 1.11×10^{-3} 90 %: 4.90×10^{-3}
79	50 %: 0.257 90 %: 0.793	50 %: 0.203 90 %: 0.666	50 %: 0.0899 90 %: 0.499
80	50 %: 0.342 90 %: 0.988	50 %: 0.246 90 %: 0.816	50 %: 0.0580 90 %: 0.462
81	50 %: 3.98×10^{-3} 90 %: 0.0132	50 %: 0.175 90 %: 0.574	50 %: 0.00 90 %: 0.00

82	50 %: 1.44×10^{-4}	50 %: 1.36×10^{-3}	50 %: 1.44×10^{-4}
	90 %: 4.79×10^{-4}	90 %: 4.50×10^{-3}	90 %: 4.79×10^{-4}

Table A.1: ***Comparison of event reconstructions.*** As the reconstruction used for the neutrino source searches has changed in this analysis as opposed to previous analyses, this table provides a comparison of the reconstructed event properties. The angular size of the 50 % and 90 % directional highest posterior density regions, and the overlap of these regions are given for each reconstruction. Only events with previously released properties are shown in the table; the reconstructed properties of all events, including full directional posterior distributions can be found in the data release. The highest posterior density regions of the directions can be directly computed for the reconstruction used in this analysis, and so are used directly in this comparison. For the previously used reconstructions, the reconstruction likelihood as a function of direction is converted to a directional probability density with a smoothing representative of the average angular resolution of each morphology. This probability density is then re-calibrated to match the angular resolution obtained from the re-simulations of the events, and the highest posterior density regions computed in this way are used for the comparison. Events with no properties listed were determined to be coincident muons. Events without properties listed from this analysis were removed from the sample for reasons discussed in Section 4.1.

A.3 Source catalog

Source	Type	RA [°]	DEC [°]	TS	n_s
TYCHO	SNR	6.36	64.18	0.00	0.00
Cas A	SNR	350.85	58.81	0.200	0.52
IC443	SNR	94.18	22.53	0.00	0.00
LSI 303	XB/mqso	40.13	61.23	0.00	0.00
Cyg X-3	XB/mqso	308.11	40.96	0.0136	0.13
Cyg X-1	XB/mqso	299.59	35.20	0.0209	0.18

HESS J0632+057	XB/mqso	98.25	5.80	0.00	0.00
SS433	XB/mqso	287.96	4.98	0.00	0.00
Cyg OB2	SFR	308.08	41.51	0.0169	0.15
MGRO J2019+37	PWN	305.22	36.83	3.32×10^{-3}	0.07
Crab Nebula	PWN	83.63	22.01	9.72×10^{-3}	0.14
Geminga	PWN	98.48	17.77	0.00	0.00
Sgr A*	GC	266.42	−29.01	1.56	1.66
MGRO J1908+06	NI	286.98	6.27	0.00	0.00
W28	SNR	270.43	−23.34	1.17	1.86
RX J1713.7-3946	SNR	258.25	−39.75	0.00	0.00
RX J0852.0-4622	SNR	133.00	−46.37	0.0800	0.54
RCW 86	SNR	220.68	−62.48	0.00	0.00
LS 5039	XB/mqso	276.56	−14.83	0.517	1.97
GX 339-4	XB/mqso	255.70	−48.79	0.00	0.00
Cir X-1	XB/mqso	230.17	−57.17	1.20×10^{-3}	0.05
Vela X	PWN	128.75	−45.60	0.148	0.68
HESS J1632-478	PWN	248.04	−47.82	0.00	0.00
HESS J1616-508	PWN	243.78	−51.40	0.00	0.00
HESS J1023-575	PWN	155.83	−57.76	0.00	0.00
MSH 15-52	PWN	228.53	−59.16	0.0285	0.21
HESS J1303-631	PWN	195.74	−63.20	0.00	0.00

PSR B1259-63	PWN	197.55	−63.52	0.00	0.00
HESS J1356-645	PWN	209.00	−64.50	0.00	0.00
HESS J1614-518	cluster	63.58	−51.82	0.00	0.00
HESS J1837-069	NI	98.69	−8.76	0.00	0.00
HESS J1834-087	NI	278.69	−8.76	1.73	2.77
HESS J1741-302	NI	265.25	−30.20	1.37	1.57
HESS J1503-582	NI	226.46	−58.74	4.09×10^{-4}	0.03
HESS J1507-622	NI	226.72	−62.34	0.0473	0.24
S5 0716+71	BL Lac	110.47	71.34	0.00	0.00
1ES 1959+650	BL Lac	300.00	65.15	0.00	0.00
1ES 2344+514	BL Lac	356.77	51.71	0.00	0.00
3C66A	BL Lac	35.67	43.04	0.00	0.00
H 1426+428	BL Lac	217.14	42.67	0.00	0.00
BL Lac	BL Lac	330.68	42.28	0.00	0.00
Mrk 501	BL Lac	253.47	39.76	0.00	0.00
Mrk 421	BL Lac	166.11	38.21	0.00	0.00
W Comae	BL Lac	185.38	28.23	0.00	0.00
1ES 0229+200	BL Lac	38.20	20.29	0.00	0.00
PKS 0235+164	BL Lac	39.66	16.62	0.00	0.00
PKS 2155-304	BL Lac	329.72	−30.23	0.00	0.00
PKS 0537-441	BL Lac	84.71	−44.09	0.00	0.00

4C 38.41	FSRQ	248.82	38.13	0.00	0.00
3C 454.3	FSRQ	343.49	16.15	0.00	0.00
PKS 0528+134	FSRQ	82.74	13.53	0.00	0.00
PKS 1502+106	FSRQ	226.10	10.49	0.00	0.00
3C 273	FSRQ	187.28	2.05	0.00	0.00
3C279	FSRQ	194.05	−5.79	0.00	0.00
QSO 2022-077	FSRQ	306.42	−7.64	0.00	0.00
PKS 1406-076	FSRQ	212.24	−7.87	0.00	0.00
QSO 1730-130	FSRQ	263.26	−13.08	0.00	0.00
PKS 1622-297	FSRQ	246.52	−29.86	3.26×10^{-3}	0.11
PKS 1454-354	FSRQ	224.36	−35.65	0.00	0.00
M82	SRG	148.97	69.68	0.00	0.00
NGC 1275	SRG	49.95	41.51	0.00	0.00
Cyg A	SRG	299.87	40.73	0.114	0.36
3C 123.0	SRG	69.27	29.67	0.00	0.00
M87	SRG	187.71	12.39	0.00	0.00
Cen A	SRG	201.37	−43.02	0.00	0.00
1ES 0347-121	BL Lac	57.35	−11.99	0.00	0.00
1ES 1101-232	BL Lac	165.91	−23.49	2.64	3.67
H 2356-309	BL Lac	359.78	−30.63	0.00	0.00
PKS 0548-322	BL Lac	87.67	−32.27	0.0558	0.55

PKS 0426-380	BL Lac	67.17	−37.93	0.00	0.00
PKS 2005-489	BL Lac	302.37	−48.82	0.00	0.00
PKS 0727-11	FSRQ	112.58	−11.70	0.510	1.71
PKS 0454-234	FSRQ	74.27	−23.43	0.126	0.86
ESO 139-G12	Seyfert	264.41	−59.94	0.00	0.00

Table A.2: ***Sources in catalogue search.*** The sources used for the search described in Section A.1 are listed above. The name of each source, the source category, and source direction are provided. Additionally, the best-fit number of source neutrinos and the TS from the search are provided as well.

Appendix B

HESE diffuse flux measurements

B.1 Data release for additional characterization of the astrophysical neutrino flux

The different astrophysical and atmospheric models explored in this paper represent only a small portion of the theoretical model space. To facilitate better model tests and the combination of results from different experiments and data sets, we provide the data and simulated event information used in the analyses described in Section 8.1. The data release information is provided in [202].

The data used for the analyses described in Section 8.1 is provided in a `json` formatted file, which contains the 102 data events that pass the selection described in Section 4.1. For each data event we provide the following variables:

- `recoDepositedEnergy` - The reconstructed deposited energy of the event, given in GeV.
- `recoMorphology` - The inferred morphology of the event, where `0, 1, 2` correspond to cascades, tracks, and double cascades, respectively.
- `recoZenith` - The reconstructed zenith direction of the event, given in radians.
- `recoLength` - The reconstructed length of the event, given in meters.

A second `json` formatted file contains the MC events used to compute the expected data event rates. For each event we provide the following variables:

- `primaryType` - The simulated initial particle flavor, given in the Monte Carlo numbering scheme outlined by the Particle Data Group [72].
- `primaryEnergy` - The simulated true energy of the initial particle (neutrino or muon), given in GeV.

- `primaryZenith` - The simulated true zenith direction of the initial particle (neutrino or muon), given in radians.
- `trueLength` - The simulated true length between cascades in a double cascade the event, given in meters.
- `interactionType` - The simulated neutrino interaction of the initial particle. Values of `1, 2, 3` correspond to CC, NC, and GR interactions, respectively. For simulated atmospheric muons a value of `0` is given.
- `weightOverFluxOverLivetime` - The MC weight of the neutrino event, divided by the simulated flux and detector livetime, given in units of GeV sr cm^2. This is set to zero for atmospheric muon events.
- `muonWeightOverLivetime` - The MC weight for each muon event, divided by the detector livetime. This is set to zero for neutrino events.
- `pionFlux` - The nominal conventional atmospheric neutrino flux from pion decay, as described in Chapter 6, given in units of $GeV^{-1}s^{-1}sr^{-1}cm^{-2}$.
- `kaonFlux` - The nominal conventional atmospheric neutrino flux from kaon decay, as described in Chapter 6, given in units of $GeV^{-1}s^{-1}sr^{-1}cm^{-2}$.
- `promptFlux` - The nominal prompt atmospheric flux neutrino flux, as described in Chapter 6, given in units of $GeV^{-1}s^{-1}sr^{-1}cm^{-2}$.
- `conventionalSelfVetoCorrection` - The veto passing fraction for conventional neutrinos, computed as described in Chapter 6.
- `promptSelfVetoCorrection` - The veto passing fraction for prompt neutrinos, computed as described in Chapter 6.
- `recoDepositedEnergy` - The reconstructed deposited energy of the event, given in GeV.
- `recoMorphology` - The inferred morphology of the event. `0, 1, 2` correspond to cascades, tracks, and double cascades, respectively.
- `recoLength` - The reconstructed length of the event, given in meters.
- `recoZenith` - The reconstructed zenith angle of the event, given in radians.

We provide a `python3` example code of using the MC information, where the energy dependent neutrino effective area of the selection is computed and the corresponding distributions are plotted. The output of this code is shown in Fig. 4.3.

The interpolating b-splines used to perform the systematic corrections, described in Sections 3.2 and 7.2, are provided as `fits` files which can be read with the `PHOTOSPLINE` software package [103].

The final component of this data release is a sample code, also written in `python3`, which reproduces the fit of the data to a single power-law astrophysical flux in the same manner as described in Section 8.1.1.1. The primary goal of these scripts are to provide a working example utilizing the information provided in the data files, and we encourage readers to use these files as a jumping point into their own analyses. A more detailed description of the files in this example are provided in the accompanying `README` file in [202].

To also support analyses of the source search variety, we provide the directional posterior distributions used in the analyses described in Section A.1. These posterior distributions are parameterized with the Fisher-Bingham eight-parameter distribution (FB8). The directional probability density of the FB8 distribution can be written as

$$f_8(\vec{x}) = c_8(\kappa, \beta, \eta, \vec{\nu})^{-1} e^{\kappa \vec{\nu} \cdot \Gamma^T \vec{x} + \beta \left[(\vec{\gamma_2} \cdot \vec{x})^2 + \eta (\vec{\gamma_3} \cdot \vec{x})^2 \right]}. \tag{B.1}$$

In this parameterization Γ is a rotation matrix which also specifies $\vec{\gamma_2}$ and $\vec{\gamma_3}$, which are the second and third columns of Γ respectively. The unit vector $\vec{\nu}$ specifies a direction. The parameters κ, β, and η, are scalars that satisfy the conditions $\kappa \geq 0$, $\beta \geq 0$, and $|\eta| \leq 1$. The term $c_8(\kappa, \beta, \eta)$ normalizes the distribution in direction and in this work is computed using the method described in [246]. In total these can be specified by eight parameters; the matrix by three parameters, the vector by two parameters, and the scalars by three parameters. For each event in the sample we provide the parameters of the corresponding FB8 distribution, as well as a `HEALPix` compatible `fits` file that contains an estimate of the integrated probability within each of the 786432 considered pixels.