

Expanding neutrino oscillation parameter measurements in NOvA using a Bayesian approach

M. A. Acero,² B. Acharya,³² P. Adamson,¹³ N. Anfimov,²⁷ A. Antoshkin,²⁷ E. Arrieta-Diaz,²⁸ L. Asquith,⁴⁰ A. Aurisano,⁷ A. Back,²¹ N. Balashov,²⁷ P. Baldi,²⁶ B. A. Bambah,¹⁸ A. Bat,^{3,12} K. Bays,^{31,20} R. Bernstein,¹³ T. J. C. Bezerra,⁴⁰ V. Bhatnagar,³⁴ D. Bhattarai,³² B. Bhuyan,¹⁶ J. Bian,^{26,31} A. C. Booth,^{36,40} R. Bowles,²¹ B. Brahma,¹⁹ C. Bromberg,²⁹ N. Buchanan,⁹ A. Butkevich,²³ T. J. Carroll,^{42,48} E. Catano-Mur,⁴⁷ J. P. Cesar,⁴² A. Chatla,¹⁸ R. Chirco,²⁰ S. Chaudhary,¹⁶ B. C. Choudhary,¹¹ A. Christensen,⁹ T. E. Coan,³⁹ A. Cooleybeck,⁴⁸ L. Cremonesi,³⁶ G. S. Davies,³² P. F. Derwent,¹³ P. Ding,¹³ Z. Djurcic,¹ M. Dolce,⁴³ D. Doyle,⁹ D. Dueñas Tonguino,⁷ E. C. Dukes,⁴⁵ A. Dye,³² R. Ehrlich,⁴⁵ M. Elkins,²⁵ E. Ewart,²¹ P. Filip,²⁴ J. Franc,¹⁰ M. J. Frank,³⁷ H. R. Gallagher,⁴³ A. Giri,¹⁹ R. A. Gomes,¹⁵ M. C. Goodman,¹ M. Groh,^{9,21} R. Group,⁴⁵ A. Habig,³⁰ F. Hakl,²² J. Hartnell,⁴⁰ R. Hatcher,¹³ M. He,¹⁷ K. Heller,³¹ V. Hewes,⁷ A. Himmel,¹³ B. Jargowsky,²⁶ J. Jarosz,⁹ C. Johnson,⁹ M. Judah,^{9,35} I. Kakorin,²⁷ D. M. Kaplan,²⁰ A. Kalitkina,²⁷ J. Kleykamp,³² O. Klimov,²⁷ L. W. Koerner,¹⁷ L. Kolupaeva,²⁷ R. Kralik,⁴⁰ A. Kumar,³⁴ C. D. Kuruppu,³⁸ V. Kus,¹⁰ T. Lackey,^{13,21} K. Lang,⁴² P. Lasorak,⁴⁰ J. Lesmeister,¹⁷ A. Lister,⁴⁸ J. Liu,²⁶ M. Lokajicek,²⁴ M. MacMahon,⁴⁴ J. A. Lock,⁴⁰ S. Magill,¹ M. Manrique Plata,²¹ W. A. Mann,⁴³ M. T. Manoharan,⁸ M. L. Marshak,³¹ M. Martinez-Casales,²⁵ V. Matveev,²³ B. Mehta,³⁴ M. D. Messier,²¹ H. Meyer,⁴⁶ T. Miao,¹³ V. Mikola,⁴⁴ W. H. Miller,³¹ S. Mishra,⁴ S. R. Mishra,³⁸ R. Mohanta,¹⁸ A. Moren,³⁰ A. Morozova,²⁷ W. Mu,¹³ L. Mualem,⁵ M. Muether,⁴⁶ K. Mulder,⁴⁴ D. Myers,⁴² D. Naples,³⁵ A. Nath,¹⁶ S. Nelleri,⁸ J. K. Nelson,⁴⁷ R. Nichol,⁴⁴ E. Niner,¹³ A. Norman,¹³ A. Norrick,¹³ T. Nosek,⁶ H. Oh,⁷ A. Olshevskiy,²⁷ T. Olson,¹⁷ A. Pal,³³ J. Paley,¹³ L. Panda,³³ R. B. Patterson,⁵ G. Pawloski,³¹ O. Petrova,²⁷ R. Petti,³⁸ A. Rafique,¹ L. R. Prais,³² V. Raj,⁵ M. Rajaoalisoa,⁷ B. Ramson,¹³ B. Rebel,^{13,48} P. Roy,⁴⁶ O. Samoylov,²⁷ M. C. Sanchez,^{14,25} S. Sánchez Falero,²⁵ P. Shanahan,¹³ P. Sharma,³⁴ S. Shukla,⁴ A. Sheshukov,²⁷ D. K. Singha,¹⁸ W. Shorrock,⁴⁰ I. Singh,¹¹ P. Singh,^{36,11} V. Singh,⁴ E. Smith,²¹ J. Smolik,¹⁰ P. Snopok,²⁰ N. Solomey,⁴⁶ A. Sousa,⁷ K. Soustruznik,⁶ M. Strait,³¹ L. Suter,¹³ A. Sutton,^{14,25,45} K. Sutton,⁵ S. Swain,³³ C. Sweeney,⁴⁴ A. Sztuc,⁴⁴ B. Tapia Oregui,⁴² P. Tas,⁶ T. Thakore,⁷ J. Thomas,^{44,48} E. Tiras,^{12,25} J. Trokan-Tenorio,⁴⁷ Y. Torun,²⁰ J. Urheim,²¹ P. Vahle,⁴⁷ Z. Vallari,⁵ K. J. Vockerodt,³⁶ T. Vrba,¹⁰ M. Wetstein,²⁵ D. Whittington,⁴¹ D. A. Wickremasinghe,¹³ T. Wieber,³¹ J. Wolcott,⁴³ M. Wrobel,⁹ S. Wu,³¹ W. Wu,²⁶ Y. Xiao,²⁶ B. Yaeggy,⁷ A. Yankelevich,²⁶ K. Yonehara,¹³ Y. Yu,²⁰ S. Zadorozhnyy,²³ J. Zalesak,²⁴ and R. Zwaska¹³

(The NOvA Collaboration)

¹Argonne National Laboratory, Argonne, Illinois 60439, USA

²Universidad del Atlantico, Carrera 30 No. 8-49, Puerto Colombia, Atlantico, Colombia

³Bandirma Onyedi Eylul University, Faculty of Engineering and Natural

Sciences, Engineering Sciences Department, 10200, Bandirma, Balikesir, Turkey

⁴Department of Physics, Institute of Science, Banaras Hindu University, Varanasi, 221 005, India

⁵California Institute of Technology, Pasadena, California 91125, USA

⁶Charles University, Faculty of Mathematics and Physics, Institute of Particle and Nuclear Physics, Prague, Czech Republic

⁷Department of Physics, University of Cincinnati, Cincinnati, Ohio 45221, USA

⁸Department of Physics, Cochin University of Science and Technology, Kochi 682 022, India

⁹Department of Physics, Colorado State University, Fort Collins, CO 80523-1875, USA

¹⁰Czech Technical University in Prague, Brehova 7, 115 19 Prague 1, Czech Republic

¹¹Department of Physics and Astrophysics, University of Delhi, Delhi 110007, India

¹²Department of Physics, Erciyes University, Kayseri 38030, Turkey

¹³Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA

¹⁴Florida State University, Tallahassee, Florida 32306, USA

¹⁵Instituto de Física, Universidade Federal de Goiás, Goiânia, Goiás, 74690-900, Brazil

¹⁶Department of Physics, IIT Guwahati, Guwahati, 781 039, India

¹⁷Department of Physics, University of Houston, Houston, Texas 77204, USA

¹⁸School of Physics, University of Hyderabad, Hyderabad, 500 046, India

¹⁹Department of Physics, IIT Hyderabad, Hyderabad, 502 205, India

²⁰Illinois Institute of Technology, Chicago IL 60616, USA

²¹Indiana University, Bloomington, Indiana 47405, USA

²²Institute of Computer Science, The Czech Academy of Sciences, 182 07 Prague, Czech Republic

²³Institute for Nuclear Research of Russia, Academy of Sciences 7a, 60th October Anniversary prospect, Moscow 117312, Russia

²⁴Institute of Physics, The Czech Academy of Sciences, 182 21 Prague, Czech Republic

²⁵Department of Physics and Astronomy, Iowa State University, Ames, Iowa 50011, USA

²⁶Department of Physics and Astronomy, University of California at Irvine, Irvine, California 92697, USA

²⁷Joint Institute for Nuclear Research, Dubna, Moscow region 141980, Russia

²⁸Universidad del Magdalena, Carrera 32 No 22-08 Santa Marta, Colombia

²⁹ *Department of Physics and Astronomy, Michigan State University, East Lansing, Michigan 48824, USA*

³⁰ *Department of Physics and Astronomy, University of Minnesota Duluth, Duluth, Minnesota 55812, USA*

³¹ *School of Physics and Astronomy, University of Minnesota Twin Cities, Minneapolis, Minnesota 55455, USA*

³² *University of Mississippi, University, Mississippi 38677, USA*

³³ *National Institute of Science Education and Research, Khurda, 752050, Odisha, India*

³⁴ *Department of Physics, Panjab University, Chandigarh, 160 014, India*

³⁵ *Department of Physics, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, USA*

³⁶ *Particle Physics Research Centre, Department of Physics and Astronomy,*

Queen Mary University of London, London E1 4NS, United Kingdom

³⁷ *Department of Physics, University of South Alabama, Mobile, Alabama 36688, USA*

³⁸ *Department of Physics and Astronomy, University of South Carolina, Columbia, South Carolina 29208, USA*

³⁹ *Department of Physics, Southern Methodist University, Dallas, Texas 75275, USA*

⁴⁰ *Department of Physics and Astronomy, University of Sussex, Falmer, Brighton BN1 9QH, United Kingdom*

⁴¹ *Department of Physics, Syracuse University, Syracuse NY 13210, USA*

⁴² *Department of Physics, University of Texas at Austin, Austin, Texas 78712, USA*

⁴³ *Department of Physics and Astronomy, Tufts University, Medford, Massachusetts 02155, USA*

⁴⁴ *Physics and Astronomy Department, University College London, Gower Street, London WC1E 6BT, United Kingdom*

⁴⁵ *Department of Physics, University of Virginia, Charlottesville, Virginia 22904, USA*

⁴⁶ *Department of Mathematics, Statistics, and Physics, Wichita State University, Wichita, Kansas 67260, USA*

⁴⁷ *Department of Physics, William & Mary, Williamsburg, Virginia 23187, USA*

⁴⁸ *Department of Physics, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA*

(Dated: November 3, 2023)

NOvA is a long-baseline neutrino oscillation experiment that measures oscillations in charged-current $\nu_\mu \rightarrow \nu_\mu$ (disappearance) and $\nu_\mu \rightarrow \nu_e$ (appearance) channels, and their antineutrino counterparts, using neutrinos of energies around 2 GeV over a distance of 810 km. In this work we reanalyze the dataset first examined in our previous paper [Phys. Rev. D **106**, 032004 (2022)] using an alternative statistical approach based on Bayesian Markov Chain Monte Carlo. We measure oscillation parameters consistent with the previous results. We also extend our inferences to include the first NOvA measurements of the reactor mixing angle θ_{13} and the Jarlskog invariant. We use these results to quantify the strength of our inferences about CP violation, as well as to examine the effects of constraints from short-baseline measurements of θ_{13} using antineutrinos from nuclear reactors when making NOvA measurements of θ_{23} . Our long-baseline measurement of θ_{13} is also shown to be consistent with the reactor measurements, supporting the general applicability and robustness of the PMNS framework for neutrino oscillations.

I. INTRODUCTION

In the three-flavor neutrino oscillation paradigm, transitions among the three flavor eigenstates ν_e, ν_μ, ν_τ are governed by the matrix elements between these states and the mass eigenstates ν_1, ν_2, ν_3 .¹ The elements of this Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix have been constrained by numerous experiments using neutrinos with a variety of energy spectra over various baselines [1–9]. Decomposing the PMNS matrix yields a set of rotation-like “mixing angles” $\theta_{12}, \theta_{13}, \theta_{23}$, and a phase δ_{CP} . Contemporary neutrino oscillation experiments seek to make precision measurements of these parameters, as well as the differences between the squared mass eigenvalues ($\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$). These results have fundamental implications for models of neutrino mass and lepton flavor [10–14], as well as models of baryogenesis via charge-parity (CP) symmetry violation [15–19]. Where multiple experiments can access the same parameters using different neutrino flavors or energies,

the overall validity of the three-neutrino framework can also be tested.

Long-baseline (LBL) accelerator neutrino oscillation experiments measure oscillations in $\nu_\mu \rightarrow \nu_\mu$ (disappearance) and $\nu_\mu \rightarrow \nu_e$ (appearance) channels. These channels constrain the mixing angles θ_{13}, θ_{23} , the mass-squared splitting Δm_{32}^2 , and the CP-violating phase δ_{CP} . Current measurements of θ_{23} are consistent with maximal mixing ($\theta_{23} = \pi/4$), which would suggest a μ - τ symmetry in their mixing into the ν_3 mass eigenstate; a non-maximal value such as $\theta_{23} < \pi/4$ (lower octant, LO) or $\theta_{23} > \pi/4$ (upper octant, UO) would indicate a preferential coupling of ν_τ or ν_μ , respectively, with ν_3 . Current experimental uncertainties on θ_{23} are the largest among the mixing angles [20]. LBL oscillation measurements, where neutrinos traverse significant quantities of matter, are also impacted by the coherent forward scattering of ν_e s on electrons in the Earth [21]. This modifies the oscillation probabilities $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$ and $\mathcal{P}(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ with opposite signs. The direction of the resulting change in rate depends on whether ν_3 is the heaviest neutrino state (Normal Ordering, NO) or the lightest (Inverted Ordering, IO). Current observations from LBL (and other) experiments prefer the NO, though the strength of the NO hypothesis depends on which measurements are included

¹ The same is true for their antineutrino counterparts $\bar{\nu}_e, \bar{\nu}_\mu, \bar{\nu}_\tau$ and $\bar{\nu}_1, \bar{\nu}_2, \bar{\nu}_3$. Throughout this paper the symbol ν will refer to both neutrinos and antineutrinos unless otherwise specified.

[22]. The value of δ_{CP} is currently essentially unknown; LBL experiments provide the only constraints, but they are weak.

NOvA is a long-baseline neutrino oscillation experiment that observes the ν_{μ} disappearance and ν_e appearance channels using neutrinos of energies around 2 GeV over a distance of 810 km. Previously [23–29], we have presented NOvA constraints on Δm_{32}^2 , $\sin^2 \theta_{23}$, and δ_{CP} using a classical frequentist approach. However, the Feldman–Cousins technique that is required in order to obtain correct classical confidence regions for these variables [30, 31] poses challenges when confronted with highly degenerate sets of parameters. It also does not allow for *post hoc* transformations of the variables considered in the analysis. In this work we present a new analysis of the dataset from Ref. [29], based on Bayesian Markov Chain Monte Carlo, which enables us to extend our inferences to include θ_{13} and the Jarlskog invariant J , for which $J \neq 0$ unambiguously indicates CP violation. We use these results to examine the implications of assuming short-baseline, nuclear-reactor antineutrino constraints on θ_{13} when making measurements of other oscillation parameters in NOvA.

We also investigate the consistency of the PMNS framework by comparing the constraint from reactor experiments with our long-baseline measurement of θ_{13} .

II. 3-FLAVOR NEUTRINO OSCILLATIONS IN NOvA

In this paper, we reanalyze data collected from an exposure of 13.6×10^{20} 14 kton-equivalent protons on target (POT) in the neutrino-enriched beam mode and 12.5×10^{20} POT in the analogous antineutrino mode. The dataset, simulations, reconstruction, and estimation of systematic uncertainties remain unchanged for this analysis. A brief overview of these components follows, with detailed descriptions available in Ref. [29]. Extensive discussion of the new analysis method, its implementations, and the resulting inferences are presented in Sec. III.

A. The NOvA experiment

NOvA observes oscillations using neutrinos from the NuMI beamline [32] at Fermilab using two functionally identical tracking calorimeter detectors that differ primarily in size. The 0.3 kton near detector (ND), the smaller of the detectors, is located 1 km from the neutrino production target, 100 m underground. The far detector (FD), by contrast, is 14 kton and is located on the surface at Ash River, Minnesota, 810 km from the target. Both detectors are built from PVC cells of $3.9 \times 6.6 \text{ cm}^2$ cross-sectional area and 3.9 m (ND) or 15.5 m (FD) length, which are arranged in alternating horizontal and vertical planes and filled with a mineral-oil-based liquid scintillator. A stack of alternating active planes and steel plates

is placed behind the ND to range out muons while a small rock overburden is placed above the FD to aid in rejecting the cosmic-ray background. The detectors are placed 14.6 mrad from the central axis of the beam to receive a narrow-band neutrino flux predominantly between 1 and 3 GeV.

B. Simulation and selection

We use GEANT4 (v10.4) [33, 34] to simulate the production of hadrons from interactions of the primary proton beam with the target as well as their transport through the beam optics. These simulations are reweighted using the Package to Predict the FluX (PPFX) [35] to include constraints from external hadron production data [36–54]. Simulated interactions of neutrinos that arise from decays of those hadrons are generated using GENIE 3.0.6 [55, 56] and modified by corrections we derive from NOvA ND and external data. NOvA ND data is used to produce a NOvA-specific tune of the IFIC València model [57, 58] that describes charged-current neutrino scattering from correlated pairs of nucleons. The NOvA tune also modifies final-state interactions (simulated with the GENIE hN full intranuclear cascade model) using pion–nucleon scattering data [59–65]. The outgoing final-state particles are propagated through the detector using GEANT4 and a custom NOvA readout simulation [66].

We use groups of spatially and temporally proximate cells with activity above threshold to apply basic data quality and containment selection cuts to events in both data and simulation. Within these groups we assign vertices and reconstruct likely particle trajectories. Ultimately we divide the events into ν_{μ} charged-current (CC), ν_e CC, neutral-current (NC), or cosmogenic background categories using NOvA’s convolutional neural network (CNN)-based classifier [67]. Boosted decision trees (BDTs) are used to further reject cosmic backgrounds in the FD samples. Both sets of tags are utilized together to create ν_{μ} CC and ν_e CC candidate samples. Fully contained ν_e candidates at the FD are further divided into high and low purity samples based on the CNN score in order to enhance the signal-to-background rejection capability of the fit. To improve the statistical strength of the fit, we recover an additional sample of “peripheral” events that fail the containment or cosmic rejection BDT but pass stricter particle ID requirements. We estimate neutrino energy for ν_{μ} CC events using the muon track length and the total deposited calorimetric energy of the hadronic system. The energy for ν_e CC events is estimated using a function of calorimetric energy that takes as input the energy of the event’s reconstructed trajectories divided into electromagnetic and hadronic components, as identified by a separate CNN-based classifier [68].

C. Near-to-far extrapolation and systematics

Predictions for the neutrino event rates at the FD are constrained using the high-statistics neutrino interactions measured in the ND. These measurements imply corrections to the ND prediction, which we propagate to the FD by adjusting for the differing efficiency and flux between the detectors using simulations. We call this process “extrapolation.” We apply oscillations to this data-driven prediction when comparing to FD data during inference (Sec. III B below).

The signal spectra for both ν_μ disappearance and ν_e appearance channels are predicted using the ν_μ spectra at the ND. The near-to-far extrapolation for the ν_μ disappearance samples at the FD is performed in quartiles of hadronic energy fraction ($f_{\text{had}} = E_{\text{had}}/E_\nu$, with E_ν being the reconstructed neutrino energy and E_{had} the reconstructed hadronic energy).

Extrapolating in the f_{had} bins has the effect of grouping events that share similar hadronic system characteristics so that compatible events are constrained together in the FD sample, despite the detectors’ somewhat different acceptances. Performing the oscillation fit in these bins enhances sensitivity to the oscillation dip as a function of E_ν since we achieve a finer energy resolution for bins with smaller f_{had} . We further subdivide the near-to-far extrapolation into three bins of reconstructed transverse momentum of the outgoing charged lepton p_T for both appearance and disappearance channels. Like the f_{had} subdivision, this allows us to better match the constraints from the much smaller ND to the FD, in this case adjusting for the differing containment of events with leptons that emerge at large angles relative to the beam direction (i.e., large p_T). The predictions in p_T bins are summed prior to the oscillation fit. We constrain the small beam backgrounds in the ν_e appearance channels with a similar procedure based on ND ν_e candidates after first decomposing them into NC, ν_μ CC, and intrinsic ν_e categories using data-driven constraints. ($\bar{\nu}_e$ beam backgrounds are all constrained together rather than being decomposed this way.) Cosmic backgrounds for all FD samples are determined from dedicated FD cosmic data samples. The remaining minor backgrounds are estimated from simulation. More detail on these procedures can be found in our previous paper [29].

We evaluate the impact of systematic uncertainties on the analysis by repredicting the sample spectra described above with altered parameters in the simulation. Uncertainties in the neutrino flux and interaction model are treated using event reweighting. Uncertainties in the detector calibration and custom modeling of light in the detectors, on the other hand, must be fully resimulated. The 67 systematically shifted simulations that arise from these techniques are extrapolated to the FD using ND data via the same process as above, which constrains the impact of the uncertainties on the predicted spectra. Parameters corresponding to these uncertainties are marginalized over in the oscillation inference in Sec. III.

Figure 1 shows the reconstructed energy spectra of the data observed at the FD using the ν_e and ν_μ CC selections described in Sec. II B, separated by the neutrino-enriched and antineutrino-enriched beam modes. Overlaid on these plots are bands of FD predictions produced according to the extrapolation procedure above. These illustrate the spectra predicted using the 68.3%, 95.4%, and 99.7% highest probability values of the systematics just described and of the relevant oscillation parameters, determined using the MCMC algorithms that will be detailed in the next section.

A Poisson likelihood [20] computed over the bins between the FD data and predictions such as those shown here forms the likelihood used in Bayes’ theorem below.

III. OSCILLATION INFERENCES USING MARKOV CHAIN MONTE CARLO

We derive posterior probability distributions for relevant oscillation parameters using Bayes’ theorem [69]. Marginalizing away the nuisance parameters of our model, which include the tens of systematic uncertainties described in Sec. II C, is a challenging problem because it requires an integral over many dimensions. We therefore turn to a Monte Carlo method for computing the posterior: Markov Chain Monte Carlo (MCMC). In MCMC, we draw a collection of sequential samples from the posterior with frequency proportional to the posterior density. Histograms that approximate the posterior shape (with accuracy governed by the sample count) may be computed in any variable(s) of interest using these samples. In so doing, any dimensions not explicitly summed are implicitly marginalized.²

Numerous algorithms for obtaining MCMC samples exist; we have implemented two for this analysis. The conclusions obtained from them agree with one another. Though descriptions of both methods are readily found in the literature, there are certain implementation choices that must be made for each, which we discuss in Sec. III A. Sec. III B lays out our resulting inferences on the parameters.

A. NOvA MCMC implementations

1. “ARIA:” MR^2T^2 algorithm

The traditional MCMC algorithm, the Metropolis–Rosenbluth–Rosenbluth–Teller–Teller (MR^2T^2) method³

² For an accessible introduction to MCMC, the reader is referred to Ref. [70]. An exhaustive treatment may be found at Ref. [71].

³ Historically this method was known as the “Metropolis” (or “Metropolis–Hastings”) method, referring to the first author of the seminal papers [72, 73]. We follow more recent convention and hereafter refer to it by all of the authors’ names. See also Ref. [74].

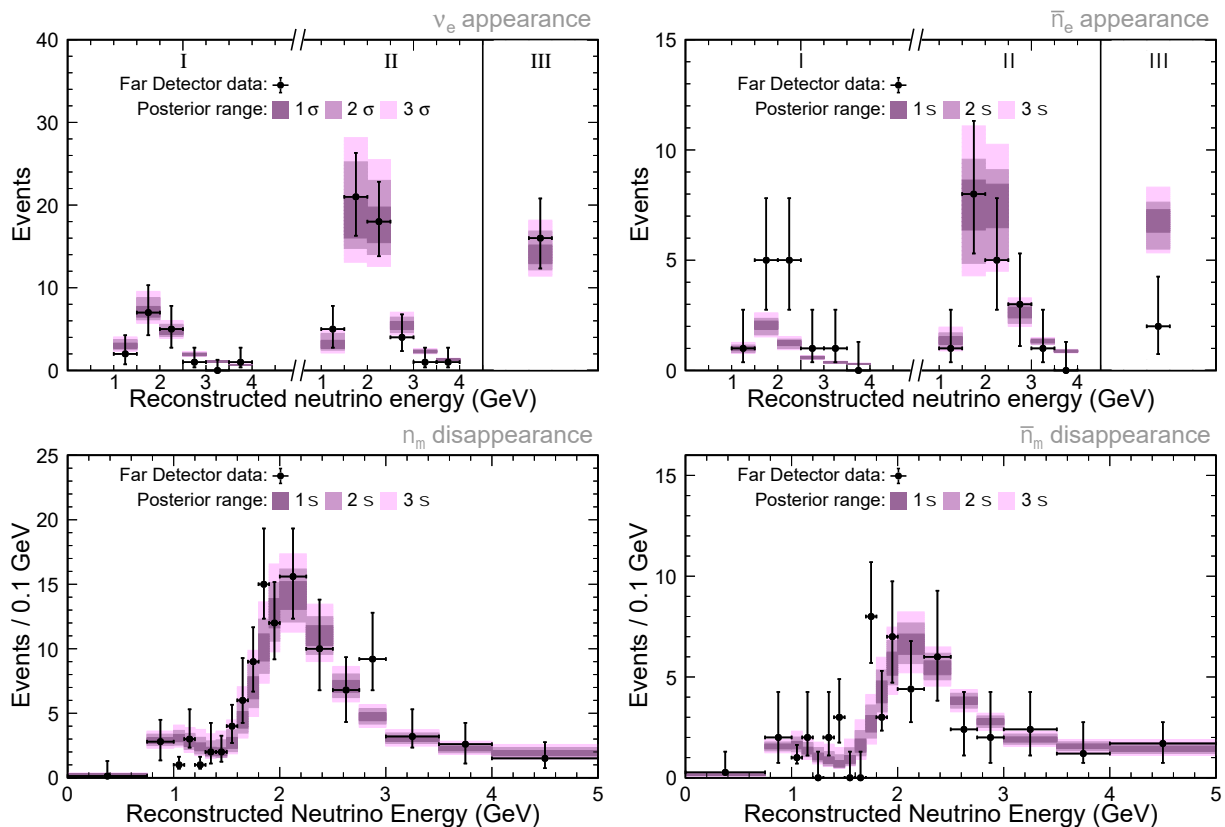


FIG. 1: Reconstructed neutrino energy distribution of selected data events (black crosses) in FD ν_e CC samples (top) and FD ν_μ CC samples (bottom) in neutrino-enriched beam mode (left) and antineutrino-enriched beam mode (right). The colored bands correspond to the range of 68.3% (darkest), 95.4%, and 99.7% (lightest) of the extrapolated FD spectra produced using the combinations of the oscillation and systematic parameters sampled by our MCMC algorithms (illustrating the posterior distributions described further in Sec. III B). The FD ν_e samples are divided into bins of low (I) and high (II) particle ID confidence as well as the peripheral (III) sample discussed in Sec. II B. The four E_{frac} ν_μ subsamples have been combined together in each of the lower two plots.

is straightforward. We call our implementation ARIA, in honor of Arianna Rosenbluth, who first implemented the method in machine code [75]. Proceeding from an initial seed in the parameter space, subsequent samples are selected by proposing a jump to a new set of coordinates, and accepting or rejecting that proposal according to an acceptance rule [20]. This process is repeated until a sufficient number of samples have been collected. There is no explicit stopping criterion. The method also does not specify the distribution to be used in the proposal algorithm. In our implementation, we use the most common choice, which is of a multivariate Gaussian. We determine the characteristic length scale of the Gaussian empirically in order to optimize sampling efficiency (see App. A). We also “thin” the resulting chains to reduce autocorrelations among the samples (see App. B). Our ARIA results below have 5×10^5 effective samples.

2. “Stan:” Hamiltonian MCMC

Though the MR^2T^2 method proposes samples quickly, they are typically highly autocorrelated. Its sampling proposals can also be inefficient if the posterior is sharply concentrated. Other MCMC methods have been developed to address these shortcomings, including one called “Hamiltonian” MCMC inference (HMCMC). We implemented a C++ interface to the Stan modeling platform [76] to obtain HMCMC samples.

The main difference between HMCMC and MR^2T^2 is how proposals are generated. Rather than proposing randomly, HMCMC views the posterior surface as a topographical one that can be explored by a fictitious particle. Samples correspond to trajectories under the influence of a gravitational potential whose gradient corresponds to that of higher posterior density. Endowing the particle with an initial momentum that counterbalances the centripetal force from gravitation produces stable trajectories that traverse the highest density region of posterior space [77]. HMCMC does this by numerically integrat-

ing Hamilton’s equations for the fictitious particle system with its position \vec{q} (which correspond to the parameters of interest) and momentum \vec{p} coordinates, and a Hamiltonian $H = -\log(\text{posterior})$. This approach produces samples that are nearly uncorrelated without thinning at the expense of additional computing cycles to compute the gradient of the posterior. (Compare the 7×10^7 Stan samples we obtained, which require no thinning, to the 5×10^5 effective samples mentioned above for ARIA that remain after thinning.) We find Stan’s default choices of the sampling distribution for the pseudoparticle kinetic energies and the integration stopping condition to be sufficient for our needs (see App. C).

Its topographical nature means HMC/MC is ill-suited to parameters that assume only one of a discrete set of values, which would manifest as discontinuities in the trajectories considered. This presents a difficulty in neutrino oscillation parameter inference, where the absolute value of Δm_{32}^2 is known with relatively good precision, but its sign remains an important unknown. While it is possible to allow HMC/MC to explore the entire range of Δm_{32}^2 , we find in practice that this results in poor exploration, as few trajectories manage to “jump” across the wide disfavored region at low $|\Delta m_{32}^2|$. Instead, at the end of each trajectory determination, we introduce a separate MR²T²-like step which considers the possibility of changing the sign of Δm_{32}^2 according to the prior probability chosen for it (50%, i.e., uniform prior; see Sec. III A 3 below). If the acceptance ratio between both mass orderings satisfies the MR²T² criteria, the proposed sign is retained; if not, it is reverted to its previous value.

3. Choices of prior

In Bayesian inference, the posterior probabilities are influenced by the choice of prior probability densities, which encode assumptions made about the parameters before the data is examined. If the data used for a measurement is sufficient, its constraint on the posterior will overwhelm the prior, and the prior choice is unimportant. However, when data is sparse, the choice of prior may affect the result. The priors we choose differ according to the parameters being considered.

a. Parameters of interest: $|\Delta m_{32}^2|$, $\text{sgn}(\Delta m_{32}^2)$, $\sin^2 \theta_{23}$ We prefer to use “uninformed” priors, which do not favor any particular value, for the physics parameters we intend to measure directly. In practice, this usually amounts to a prior uniform in the variable in question. However, uniformity is not preserved under a change of variable: for instance, a prior uniform in θ_{23} is not uniform in the measured variable $\sin^2 \theta_{23}$. In our results below, we have studied the impact of priors uniform both in a particular variable and relevant functions of it, and we report when the prior choice significantly affects the results.

b. Parameter of interest: δ_{CP} While δ_{CP} is intended to have a uniform prior as well, it receives additional




FIG. 2: Special prior distribution from Eq. 1 used for δ_{CP} .

special treatment. Its cyclical nature as the phase of a complex number results in an infinite set of values having identical consistency with the data. This can cause MCMC samplers never to converge on a single value for δ_{CP} and consequently to sample much more slowly. To combat this problem, we developed a novel special prior over δ_{CP} :

$$\Pi(\delta_{\text{CP}}) = \begin{cases} \frac{1}{2} \sin^2\left(\frac{1}{4}(\delta_{\text{CP}} + \pi)\right), & -1 \leq \delta_{\text{CP}}/\pi \leq 3 \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

This function is illustrated in Fig. 2. This prior forces the samplers to remain near a single phase of δ_{CP} , $0 \leq \delta_{\text{CP}}/\pi \leq 2$, as the prior vanishes outside of $[-1, 3]$. Because it makes the transition to the vanishing regions in a differentiable manner, this prior is suitable for use with HMC/MC. Moreover, the sum of the prior’s value at every point $0 \leq \delta_{\text{CP}}/\pi \leq 2$ with that of all its complex branch points outside that interval is a value that does not depend on δ_{CP} :

$$\sum_{n=-\infty}^{\infty} \Pi(\delta_{\text{CP}} + 2\pi n) = \frac{1}{2}. \quad (2)$$

This implies that the prior behaves identically to a uniform prior when used in conjunction with the oscillation probability. For the rest of this paper, whenever we refer to a “uniform” prior in δ_{CP} , we mean this prior.

When in the subsequent sections we study the effect of choosing a prior uniform in δ_{CP} vs. $\sin(\delta_{\text{CP}})$, we reweight the MCMC samples obtained with the prior above, taking the Jacobian factor $\partial(\sin(\delta_{\text{CP}})) = \cos(\delta_{\text{CP}})$ as the weight.

c. Parameter of interest: $\sin^2 2\theta_{13}$ As discussed further in Sec. III B below, we consider two separate cases for $\sin^2 2\theta_{13}$: one where its prior is treated identically to that of $\sin^2 \theta_{23}$ (see III A 3 a), and one where measurements from short-baseline reactor antineutrino oscillations are applied as a constraint. In the latter case,

we impose a Gaussian prior with standard deviation obtained from the 2019 world average of reactor measurements [78], analogous to the treatment in our most recent frequentist result [29]: $\sin^2 2\theta_{13} = 0.085 \pm 0.003$.

d. Systematic uncertainties For all systematic uncertainties we use a unit Gaussian distribution as their prior.

B. Oscillation inferences: results and discussion

In this section we describe our inferences regarding the PMNS neutrino oscillation parameters, given the data, model, and the Bayesian methodology described above. Two separate sets of results are obtained. The first uses a Gaussian prior on $\sin^2 2\theta_{13}$, imposing a constraint from reactor antineutrino experiments (Sec. III B 2-III B 3). The second uses a prior uniform in $\sin^2 2\theta_{13}$, yielding results constrained only by the NOvA data (Sec. III B 4). Samples obtained from ARIA and Stan produce essentially identical distributions in all the variables considered (App. D).

1. Goodness of fit

To evaluate the goodness of the fit for this Bayesian analysis we use posterior predictive p -values (PPP) [79]. In a PPP test, the coordinates from each MCMC sample are used to form a prediction for the observed spectra. A χ^2 statistic is computed between each prediction and the data, which we denote as χ^2_{data} . A second prediction is made for each sample by applying Poisson fluctuations to the prediction above, and a second $\chi^2_{\text{pseudodata}}$ calculated between this pseudodata and the original prediction. Because the data spectra are unchanged in this process, χ^2_{data} incorporates only variations in the oscillation parameters and systematic uncertainties. Conversely, since the pseudodata distributions have Poisson fluctuations applied but use the same oscillation parameters and systematic uncertainty pulls as the base model for each MCMC sample, $\chi^2_{\text{pseudodata}}$ treats only statistical uncertainties.

The distribution of these $(\chi^2_{\text{data}}, \chi^2_{\text{pseudodata}})$ pairs for the entire ensemble of spectra considered in Fig. 1 is shown as the purple shading in Fig. 3. The PPP then consists of the fraction of points in this ensemble that lie above the $\chi^2_{\text{data}} = \chi^2_{\text{pseudodata}}$ line. In the limit of infinite MCMC samples, a model that perfectly describes the data apart from statistical variations will produce a PPP of 0.5. We observe that the shaded distribution in Fig. 3 is distributed evenly around 1 unit of χ^2 per degree of freedom in both axes, and the PPP we obtain is 0.56. Both of these imply that the model is a good representation of the data.

We also find good p -values for the ν_μ and $\bar{\nu}_\mu$ samples considered independently, whose 68.3% credible regions




FIG. 3: Posterior predictive p -values from real data MCMC samples, with θ_{13} constraint from reactor experiments applied. The purple distribution is a scatterplot of the binned χ^2 computed between the model and real data spectra (x -axis), against a similar χ^2 between the model and pseudodata spectra (y -axis), both divided by the number of degrees of freedom (DOF) in the fit, computed for each MCMC sample. (See the text for more on how the pseudodata spectra are constructed.) The dashed and solid contours show 1σ intervals from the same posterior-predictive distributions calculated only for ν_e (dark blue, left solid), $\bar{\nu}_e$ (red, right solid), ν_μ (light blue, left dashed) and $\bar{\nu}_\mu$ (light red, right dashed) data samples. The posterior predictive p -value is the fraction of the distribution that lies above the diagonal $\chi^2_{\text{data}} = \chi^2_{\text{pseudodata}}$ line (black). The fact that the purple distribution is centered at (1, 1) and is evenly distributed above and below the diagonal, with the associated p -value of 0.56, indicates a good fit.

are indicated by the dashed light blue and light red colored contours in Fig. 3, respectively. By contrast, the effect of fluctuations in the smaller-statistics samples can be seen more readily in the analogous ν_e (solid dark blue) and $\bar{\nu}_e$ (solid red) contours. Both contours are much larger than their ν_μ counterparts, particularly along the y -axis, which corresponds to the dimension where statistical uncertainties are considered. In the ν_e contour, the offset downwards from unity along the x -axis, and the corresponding shift above the diagonal along the y -axis, together suggest a set of fluctuations that are relatively close to the Asimov (unfluctuated) model prediction. Each member of the ensemble of pseudodata spectra thus typically has larger χ^2 relative to the Asimov than that of the data. This relative closeness of the ν_e prediction to the data can also be seen in the top left panel of Fig. 1. The unusual shape of this contour arises

because the PPP distribution for ν_e consists of two modes superimposed upon one another, corresponding to the high-probability regions within the NO that will be discussed below in Fig. 7. On the other hand, the situation is reversed for the $\bar{\nu}_e$ spectra, where the data fluctuations result in larger deviations from the Asimov prediction than the bulk of the pseudodata spectra. The fact that combining all four subsamples together produces a PPP that is closer to 0.5 than any of them are individually is good evidence that the deviations from PPP of 0.5 in the various subsamples are predominantly driven by statistical, rather than systematic, effects.⁴ We conclude therefore that our set of MCMC samples reflect PMNS parameters with a good description of the physics exhibited in our data.

2. PMNS parameter measurements

We produced Bayesian credible regions for the PMNS neutrino oscillation parameters using the data spectra and the MCMC samplers described in sections III A and II C, respectively. These credible regions together with the posterior probability distributions are shown in Figs. 4 and 7. In both cases the reactor θ_{13} constraint used is the 2019 PDG’s combination of extant measurements [78] and applied in the form of a Gaussian prior on $\sin^2 2\theta_{13}$, as explained in section III A 3 c. In the figures in this section, credible intervals that show the normal and inverted mass orderings separately are created by first making one shared credible interval that spans both. The separate panels then display the relevant regions from this shared interval that apply to the specified ordering. By constructing them in this manner, we ensure the NO and IO intervals share a highest-posterior density point and posterior probability distribution, preserving any NOvA preference towards one of the mass orderings in the credible region limits. Similarly, the distributions and intervals labeled “both orderings” are created from MCMC samples over all values for Δm_{32}^2 , summing together the normal and inverted ordering posteriors before extracting the credible intervals.

Table I shows the highest posterior probability points together with the credible intervals enclosing 68.3% of the posterior (henceforth called the 1σ contour in analogy with Gaussian p -values in a frequentist context). These points are given for all the PMNS oscillation parameters of interest, split into both, normal, and inverted mass orderings (using the same methodology regarding the mass ordering as for the figures). For some of the parameters the 1σ region spans disjoint areas; we denote this with a union symbol \cup . These high posterior probability regions are in generally good agreement with the frequentist analysis of the same dataset [29].

TABLE I: The highest posterior density points (HPD) together with the 1σ Bayesian credible interval limits for the PMNS parameters of interest marginalized over all the mass ordering (MO) hypotheses. Marginalization over the mass orderings is explained at the beginning of Sec. III B 2. In these results a Gaussian prior corresponding to reactor constraint on $\sin^2 2\theta_{13}$ (see sec. III A 3 c) is applied.

	MO	HPD	1σ
δ_{CP}	Both	0.91π	$[0.02\pi, 0.31\pi] \cup [0.68\pi, 1.67\pi]$
	Normal	0.89π	$[0.54\pi, 1.07\pi] \cup [1.99\pi, 0.48\pi]$
	Inverted	1.44π	$[1.26\pi, 1.65\pi]$
$\sin^2 \theta_{23}$	Both	0.56	$[0.45, 0.49] \cup [0.52, 0.59]$
	Normal	0.56	$[0.44, 0.59]$
	Inverted	0.56	$[0.55, 0.57]$
Δm_{32}^2 ($\times 10^{-3}$)	Normal	2.39	$[2.32, 2.46]$
	Inverted	-2.44	$[-2.47, -2.41]$

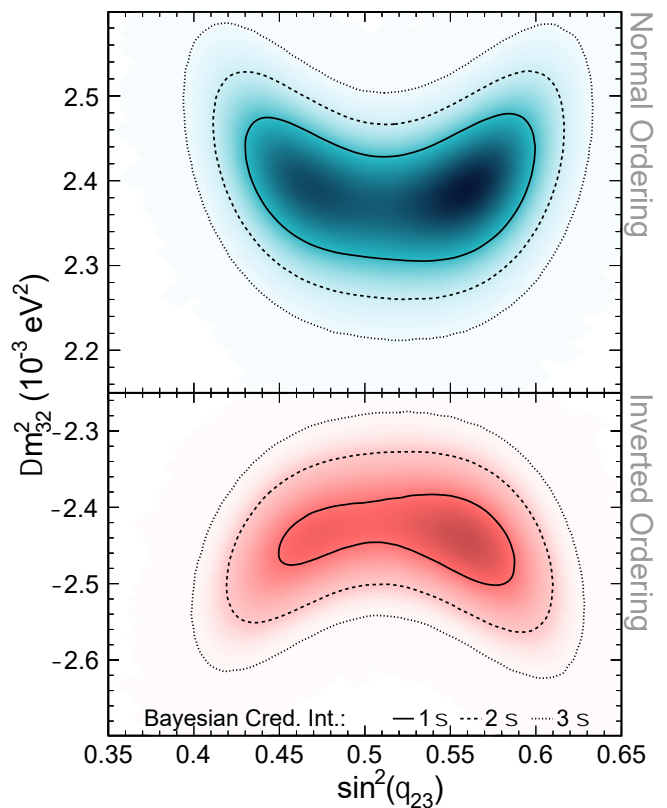


FIG. 4: Binned posterior probability densities (shaded) for $\sin^2 \theta_{23} - \Delta m_{32}^2$, marginalized over both mass orderings and plotted separately for the normal (top) and the inverted (bottom) mass orderings (marginalization over the mass orderings is explained at the beginning of Sec. III B 2). Contours indicate regions enclosing 68.3%, 95.5%, and 99.7% of the posterior probability, and for convenience are labeled in terms of Gaussian standard deviations σ using the corresponding traditional z -scores (i.e., the number of standard deviations away from the central value).

⁴ The reader may find further information on interpretations of posterior predictive p -values in Refs. [80, 81].

Figure 4 shows the $\sin^2 \theta_{23} - \Delta m_{32}^2$ plane, where the denser MCMC samples (darker color) and larger credible regions in the upper panel relative to the lower indicate a mild preference for the normal ordering. This conclusion holds in the presence of the entire systematic uncertainty model discussed in Sec. II C. However, the preferred regions in each mass ordering depend in a nontrivial way on the systematic uncertainties, as illustrated in Fig. 5. The most significant effect is in Δm_{32}^2 , where the systematic effects not only broaden the preferred region, but also shift the central value to larger absolute magnitudes. The most important of the uncertainties contributing to this movement is in the absolute calibration of the calorimetric energy scale. Because this directly affects reconstructed neutrino energies, it shifts the expected number of events in the trough of the ν_μ and $\bar{\nu}_\mu$ disappearance spectra and is thus anticorrelated with Δm_{32}^2 , as can be seen in Fig. 6. Shifting Δm_{32}^2 in this way also moves $\sin^2 \theta_{23}$ closer to maximal disappearance ($\sin^2 \theta_{23}^{\text{MD}} \approx 0.51$; the value depends slightly on $\sin^2 \theta_{13}$). Because the ν_μ disappearance spectra are essentially identical for values reflected across the maximal disappearance line, the credible regions are nearly symmetric around it (the degeneracy is broken only by the ν_e appearance spectra, which have less statistical power). Thus, the credible regions for smaller $|\sin^2 \theta_{23} - \sin^2 \theta_{23}^{\text{MD}}|$ appear narrower, even though the sensitivity is unchanged. This effect is responsible for the apparently slightly tighter constraint on $\sin^2 \theta_{23}$ observed in Fig. 5 under the effect of systematic uncertainties.

The situation is different in the $\delta_{\text{CP}} - \sin^2 \theta_{23}$ plane, which is shown in Fig. 7. Here, we observe preferences for CP-nonconserving values of δ_{CP} (i.e., nonintegral values of δ_{CP}/π) in both normal and inverted orderings, and for the upper octant ($\sin^2 \theta_{23} > 0.5$) of θ_{23} (reflected also in Fig. 4), though the CP-conserving points are only weakly disfavored in NO. However, in contrast to the conclusions for $\sin^2 \theta_{23} - \Delta m_{32}^2$, these inferences are minimally affected by the presence of systematic uncertainties, as Fig. 8 makes clear—even if the resolution does degrade slightly and the credible regions grow. Here the minimal impact of systematic uncertainties owes primarily to the smaller statistics of the ν_e and especially $\bar{\nu}_e$ appearance samples that drive the sensitivity to these parameters, which results in statistical uncertainties dominating the uncertainty budget.

To more quantitatively assess the mass ordering and octant preferences, we give the posterior probabilities inferred for each combination of hypotheses in Table II. We also express them in a less prior-dependent way using Bayes factors. In both cases the evidence for either option is weak.⁵ The analogous Gaussian p -values also correspond to significances of less than 1σ . The interpre-

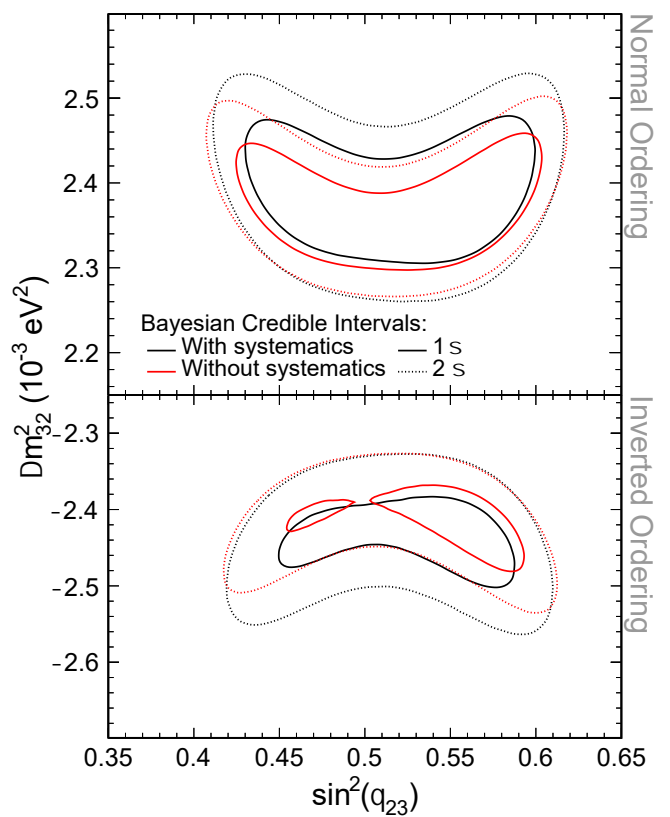


FIG. 5: Credible interval comparisons for $\sin^2 \theta_{23} - \Delta m_{32}^2$ between statistical-only fits and fits including all the NOvA systematic parameters. The external constraint on θ_{13} from the reactor experiments was applied in these fits.




FIG. 6: Binned posterior probability density (shaded) with 1, 2, and 3 σ credible intervals in $|\Delta m_{32}^2|$ and the NOvA absolute calibration systematic uncertainty, where the horizontal axis is measured in units of standard deviations from the nominal value. The external constraint on θ_{13} from reactor experiments was applied in this fit.

⁵ The reader unfamiliar with the interpretation of Bayes factors is referred to the standard treatments of Jeffreys [82] or Kass & Raftery [83].




FIG. 7: Binned posterior probability density (shaded) with 1, 2, and 3 σ credible intervals for $\delta_{\text{CP}} - \sin^2 \theta_{23}$, marginalized over both mass orderings for the normal mass ordering (top, blue) and the inverted mass ordering (bottom, red). In this fit a Gaussian θ_{13} prior from reactor experiments was applied.

tations of the octant and the mass ordering hypothesis preferences are in good agreement with the 2020 frequentist analysis of the same dataset, which used profiling with Feldman–Cousins corrections instead of marginalization [29, 31]. This general agreement is also true for the parameters’ intervals, with small differences expected given the differing statistical methods used. To examine the CP-conservation situation more comprehensively, taking the whole PMNS matrix into consideration, we will use the Jarlskog invariant measure, explored in the next section.

3. CP violation – Jarlskog invariant

The Jarlskog invariant is a measure of the strength of charge-parity violation that is independent of how the mixing matrix is parameterized. The neutrino-mixing variant [84] parallels its development in the quark sector [85]. Under the three-neutrino-flavor assumption, its

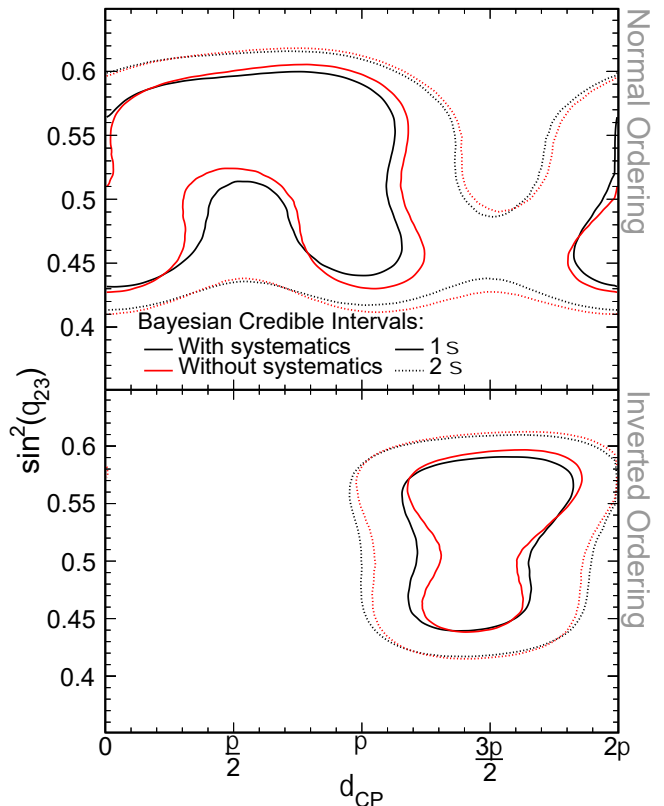


FIG. 8: Credible interval comparisons for $\delta_{\text{CP}} - \sin^2 \theta_{23}$ between statistical-only fits and fits including all the NOvA systematic parameters, with external constraint on θ_{13} from reactor experiments.

TABLE II: Bayes factors (posterior probabilities) in all the θ_{23} octant and mass ordering hypotheses combinations. A weak preference towards the normal mass ordering and upper octant is observed. Numbers extracted from a fit with the external θ_{13} constraint. Probabilities summed across rows or columns may differ slightly from the totals due to rounding. In these results the reactor constraint on $\sin^2 2\theta_{13}$ is applied.

	Normal Ordering	Inverted Ordering	Total
Upper Octant	0.71 (0.42)	0.26 (0.21)	1.67 (0.63)
Lower Octant	0.35 (0.26)	0.13 (0.12)	0.60 (0.38)
Total	2.08 (0.68)	0.48 (0.33)	(1.0)

definition is:

$$J \equiv \cos(\theta_{12}) \cos^2(\theta_{13}) \cos(\theta_{23}) \sin(\theta_{12}) \times \sin(\theta_{13}) \sin(\theta_{23}) \sin(\delta_{\text{CP}}), \quad (3)$$

where if any of the factors is zero, the invariant vanishes and the neutrino mixing matrix is CP-conserving. Nonzero values of J , on the other hand, indicate CP violation. We produce this measurement by taking the

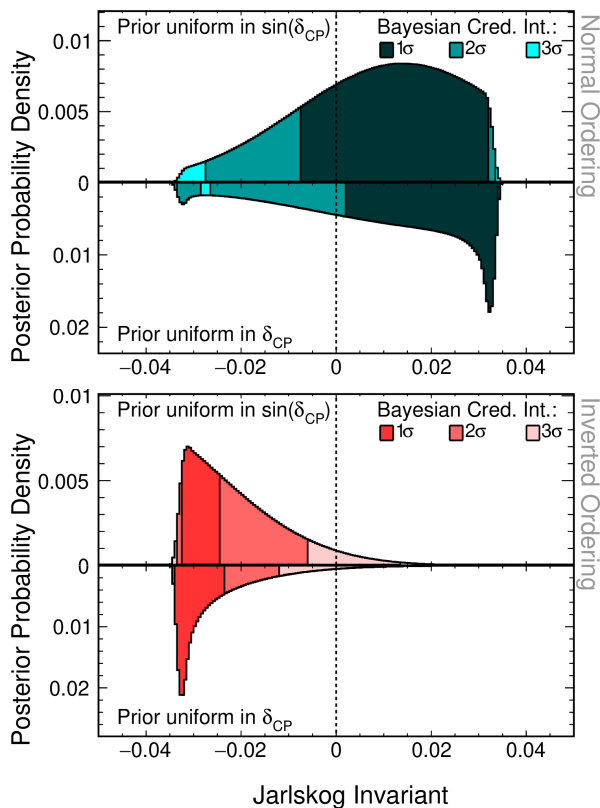


FIG. 9: Posterior probability density for the Jarlskog invariant in NO (top plot) and IO (bottom plot). The top half of each panel shows the posterior with a prior uniform in $\sin \delta_{\text{CP}}$, while the bottom half uses a prior uniform in δ_{CP} . A CP-conserving line is drawn at $J = 0$. The external constraint in θ_{13} from the reactor experiments is used.

MCMC chain with all the oscillation parameters' values and calculating the Jarlskog invariant at each MCMC step. As NOvA is not sensitive to the value of θ_{12} , for J we use a Gaussian prior analogous to the one used for the reactor constraint (sec. III A 3 c) but whose range consists of the PDG's 2019 average of solar and LBL reactor neutrino measurements: $\sin^2 \theta_{12} = 0.307 \pm 0.013$ [78], the same value used in our previous results.

As described in section III A 3, we use uninformed priors for the oscillation parameters where possible, meaning that the priors are uniform in the variable that is being shown. Since the Jarlskog invariant is written in terms of $\sin \delta_{\text{CP}}$, the natural formulation is in terms of a prior uniform in $\sin \delta_{\text{CP}}$. At the same time, there are theoretical considerations that suggest a prior uniform in δ_{CP} may be more appropriate [86]. We therefore consider both a prior uniform in $\sin \delta_{\text{CP}}$ and one uniform in δ_{CP} . Priors on the other oscillation parameters were investigated; these showed very weak effects on the J posterior due to the far higher sensitivity of NOvA data to the other PMNS parameters constrained above besides δ_{CP} .

Figure 9 shows the inferred values of the Jarlskog in-

variant extracted from a fit to the NOvA data with the external constraint on θ_{13} from the reactor experiments, marginalized over the normal (top) and inverted (bottom) mass orderings. The regions most favored by the data tend towards CP violation in the NO, although the 1σ intervals lie very close to $J = 0$. The IO's preferred regions are more markedly distant from CP conservation, particularly for the prior uniform in δ_{CP} . As expected, these trends mirror those observed when considering CP-conserving and CP-violating values of δ_{CP} in Sec. III B 2.

A more quantitative way of measuring the NOvA preference for CP conservation with the Jarlskog invariant is through Bayes factors. These can be calculated with the Savage–Dickey density ratio method, which computes Bayes factors for point hypotheses [87, 88]. Using this method we can calculate the Bayes factor for the CP-conserving value $J = 0$, nested under the unconstrained hypothesis where J can take any value. Bayes factors always depend formally on the choice of prior, but in many circumstances (such as those discussed above), if two hypotheses being compared use the same prior, it will cancel. This is not the case in the Savage–Dickey method. Therefore, we compute the Bayes factor for $J = 0$ under both priors for δ_{CP} considered above. We then invert it to obtain the associated Bayes factors for CP violation over CP conservation, $J \neq 0$.

TABLE III: Bayes factors for preference of CP violation over CP conservation, extracted using the Savage–Dickey method at $J = 0$. Priors uniform in δ_{CP} and $\sin \delta_{\text{CP}}$ are both shown. The preferences are given for the normal (NO), inverted (IO), and both (BO) mass orderings.

Prior	NO	IO	BO
Uniform $\sin \delta_{\text{CP}}$	1.3	3.5	1.5
Uniform δ_{CP}	1.1	4.4	1.5

Table III shows the Bayes factors for CP violation against CP conservation for normal, inverted, and both mass orderings, calculated for priors uniform in $\sin \delta_{\text{CP}}$ or δ_{CP} . All of these probabilities point towards a preference for CP nonconservation, although the preferences are minimal regardless of the δ_{CP} prior or the mass ordering (and an analogous frequentist p -value would indicate significances less than 1σ , apart from the inverted ordering, where they range $1.1 - 1.2\sigma^6$).

⁶ We emphasize that these significances cannot be read off of Fig. 9. The probability density shown there is marginalized independently at each value of J . However, the binary hypothesis test $J = 0$ vs. $J \neq 0$ requires a simultaneous marginalization across the whole $J \neq 0$ space. For the latter, a point-hypothesis treatment, such as the the Savage–Dickey formalism we use here, is necessary.

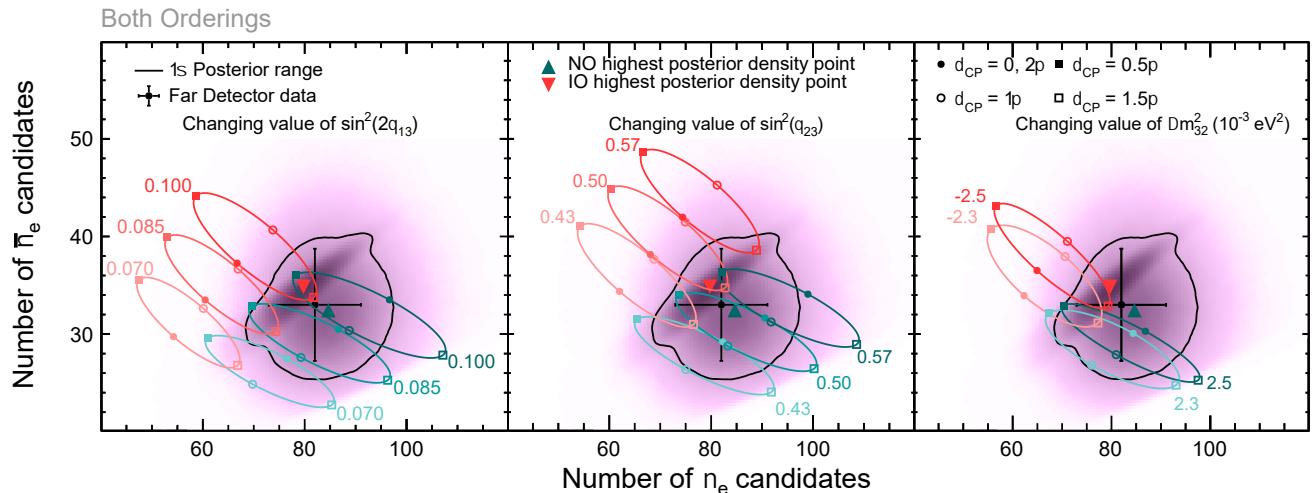


FIG. 10: Posterior probability density (purple shading)—that is, the probability that a given $(\nu_e, \bar{\nu}_e)$ pair of counts is the true underlying value in nature (see text)—compared to the measurement and associated expected statistical variance (black cross) of the total number of $\bar{\nu}_e$ candidates vs. ν_e candidates. The black solid contours indicate the regions enclosing 68.3% of the posterior. Each panel overlays a set of ellipses corresponding to the predicted event counts over the range of possible values of δ_{CP} , with the parameter notated nearby held at that value (blue for NO, red for IO, with varying shading depending on the parameter value) and other parameter values as given in Table IV. The triangular markers show the highest posterior density in the ν_e - $\bar{\nu}_e$ space for when restricting to NO (blue) and IO (red) hypotheses. The posteriors shown here employ a uniform prior over $\sin^2 2\theta_{13}$.

4. Using only NOvA constraints on θ_{13}

NOvA’s oscillation measurements simultaneously constrain a combination of mixing angles (θ_{23} , θ_{13}), the mass-squared splitting Δm_{32}^2 , the CP phase δ_{CP} , and the neutrino mass ordering. Therefore, applying a strong external 1D constraint on θ_{13} from reactor antineutrino experiments—and thereby reducing the available solution space—increases the sensitivity to other parameters. This is especially true of the octant of θ_{23} , as we will show below. However, using an uninformed (uniform) prior in $\sin^2 2\theta_{13}$ enables us to directly compare a NOvA-only measurement against that of the reactor experiments. In so doing we can examine the robustness of the PMNS description across short-baseline reactor antineutrino experiments and long-baseline accelerator-based oscillation experiments. We also may study whether our preferences change in the presence of an external constraint on the data.

Without an external constraint on $\sin^2 2\theta_{13}$, the relevant degrees of freedom present a complex space with many intercorrelations among the parameters and numerous possible solutions. We can use the type of posterior distributions shown in Fig. 1 to explore these possibilities: for example, by computing the total number of predicted ν_e and $\bar{\nu}_e$ candidates for the parameters of each MCMC sample and comparing the distribution of these predictions in $(\nu_e, \bar{\nu}_e)$ candidate space to what we observe in the data. This is shown in Fig. 10. To guide our intuition, we overlay ellipses corresponding to the ranges

of predicted number of events for the possible values of δ_{CP} at fixed values of the other parameters, subject to constraint from our data in that the highest posterior density (HPD) value in each ordering is used for parameters not explicitly varied. We note several important features. First, the highest-density (highest-probability) region lies essentially equally along the overlap region between NO and IO ellipses in all panels, meaning we will find good solutions for either, given an appropriate value of δ_{CP} . However, the NO ellipses subtend more of the posterior region, which will result in a slight overall preference for NO. Second, it is clear from the left panel that the value of $\sin^2 2\theta_{13}$ preferred by the reactor average, $\sin^2 2\theta_{13} = 0.085$, is very close to the highest-probability region for NOvA. Third, comparing the left and center panels, there is obvious degeneracy between $\sin^2 2\theta_{13}$ and $\sin^2 \theta_{23}$, since independently varying them produces similar effects on the predictions. And finally, the right panel demonstrates that the value of Δm_{32}^2 plays a minor (though not negligible) role in the appearance channel; its primary constraints arise from the ν_μ disappearance measurement. These features will be examined more quantitatively in the distributions that follow. Table IV shows the highest posterior probability points together with the 1σ credible intervals for all the PMNS oscillation parameters of interest, extracted from the fit with a prior uniform in $\sin^2 2\theta_{13}$. Similarly to Table I, the values are split into both, normal, and inverted mass orderings with any disjoint 1σ regions denoted with a union symbol \cup . As compared to the results of the fit

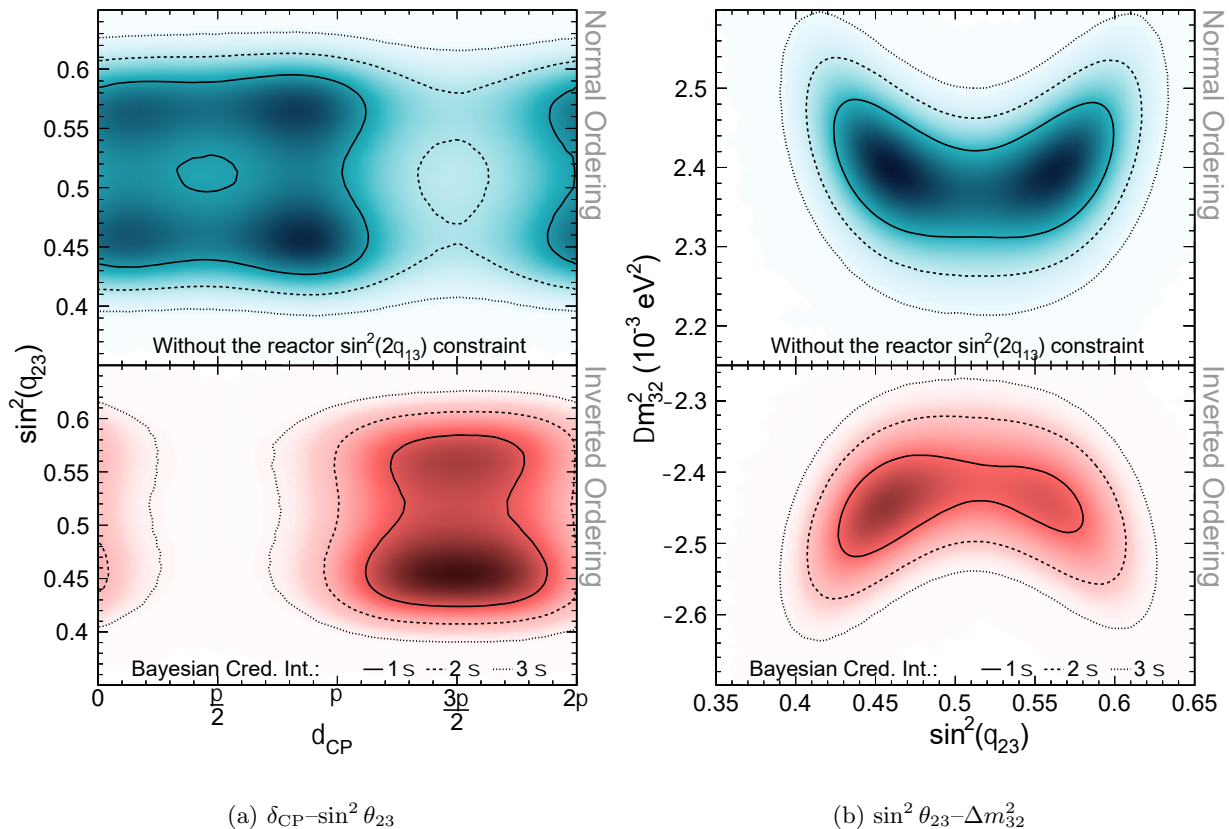


FIG. 11: Posterior probability density (shaded) with 1, 2, and 3 σ credible intervals for $\delta_{\text{CP}} - \sin^2 \theta_{23}$ (left) and $\sin^2 \theta_{23} - \Delta m_{32}^2$ (right), marginalized over both mass orderings, plotted separately for NO (top, blue) and IO (bottom, red). Contours extracted from a fit with prior uniform in $\sin^2 2\theta_{13}$.

with the external constraint from reactor experiments, the central values for δ_{CP} and $\sin^2 \theta_{23}$ shown here are the most different. This is not unexpected since there is also a degeneracy between these two parameters, and both δ_{CP} and $\sin^2 \theta_{23}$ have multiple areas of high probability.

Fig. 11 shows our preferred regions in $\delta_{\text{CP}} - \sin^2 \theta_{23}$ and $\sin^2 \theta_{23} - \Delta m_{32}^2$ spaces without the external θ_{13} constraint. Comparing with Figs. 4 and 7, it is evident that removing the constraint diminishes the NOvA sensitivity especially to the octant of θ_{23} (there is now near symmetry across $\sin^2 \theta_{23} = 0.5$). Although the octant preference is weakened, the central values of the PMNS parameters $\sin^2 \theta_{23}$, Δm_{32}^2 , and especially δ_{CP} do not otherwise see any significant change. This is as expected from Fig. 10. The results without the reactor constraint are therefore fully compatible with the standard NOvA results with the external constraint on $\sin^2 2\theta_{13}$ applied. We also observe that although the sensitivity is reduced, and although that for a given value of δ_{CP} there is always a combination of parameters within either mass ordering compatible with the data, certain $\delta_{\text{CP}} - \sin^2 \theta_{23}$ -ordering combinations are still excluded with reasonable confidence, such as ($\delta_{\text{CP}} = \frac{\pi}{2}$, IO) and ($\delta_{\text{CP}} = \frac{3\pi}{2}$, $\sin^2 \theta_{23} = 0.5$, NO).

Combinations such as these produce strong asymmetries in the $(\nu_e, \bar{\nu}_e)$ counts, and as Fig. 10 makes clear, such parameter values lie in regions outside the posterior point cloud, and are thus disfavored. Table V examines the posterior probabilities for each of the octant and mass-ordering hypotheses more quantitatively. Compared to Table II, we note a similar weakened preference for the upper octant of θ_{23} for the fit without the reactor constraint.

In Fig. 12 we study the impact of applying the reactor constraint on our mass-ordering inference. As expected from Fig. 10, the normal ordering contains more posterior probability, as evidenced by the larger NO contour in the lower panel. Moreover, as we see in the marginal posterior distribution shown in the top panel, the posterior restricted to IO prefers generally larger values of $\sin^2 2\theta_{13}$. The associated correlation with $\sin^2 \theta_{23}$ also pushes θ_{23} into the lower octant, as seen in the lower panel; this point will be developed further momentarily. Because the reactor value (the yellow bar indicates its 1σ range) lies at lower $\sin^2 2\theta_{13}$, this results in the slightly stronger preference for the NO in Table II as compared to V. However, the difference is small, which indicates that the (mild) NO preference observed in the data is largely

TABLE IV: The highest posterior density points (HPD) together with the 1σ Bayesian credible interval limits for the PMNS parameters of interest marginalized over all the mass ordering (MO) hypotheses. Marginalization over the mass orderings is explained at the beginning of Sec. III B 2. Extracted from a fit with prior uniform in $\sin^2 2\theta_{13}$.

	MO	HPD	1σ
δ_{CP}	Both	1.53π	$[0.67\pi, 1.88\pi] \cup [0.06\pi, 0.12\pi]$
	Normal	0.85π	$[0.48\pi, 1.04\pi] \cup [1.99\pi, 0.41\pi]$
	Inverted	1.52π	$[1.25\pi, 1.75\pi]$
$\sin^2 2\theta_{13}$	Both	0.087	$[0.071, 0.107]$
	Normal	0.084	$[0.065, 0.108]$
	Inverted	0.094	$[0.083, 0.106]$
$\sin^2 \theta_{23}$	Both	0.46	$[0.43, 0.50] \cup [0.53, 0.58]$
	Normal	0.46	$[0.43, 0.59]$
	Inverted	0.46	$[0.44, 0.48]$
Δm_{32}^2 ($\times 10^{-3}$)	Normal	2.39	$[2.33, 2.46]$
	Inverted	-2.44	$[-2.48, -2.40]$

TABLE V: Bayes factors (posterior probabilities) for all the θ_{23} octant and mass ordering hypotheses, with a marginal preference towards the normal mass ordering and upper octant. Numbers extracted from a fit with a uniform prior in $\sin^2 2\theta_{13}$. Probabilities summed across rows or columns may differ slightly from the totals due to rounding.

	Normal Ordering	Inverted Ordering	Total
Upper Octant	0.53 (0.35)	0.20 (0.17)	1.05 (0.51)
Lower Octant	0.39 (0.28)	0.42 (0.20)	0.95 (0.49)
Total	1.70 (0.63)	0.59 (0.37)	(1.0)

independent of the reactor constraint.

The degeneracy in the measurement between $\sin^2 \theta_{23}$ and $\sin^2 2\theta_{13}$ as noted in Fig. 10 can be studied directly by examining their joint posterior probability distribution, marginalized over all other parameters and the mass ordering. We show this in Fig. 13. The central panel exhibits a clear anticorrelation between the octant of θ_{23} and the value of $\sin^2 2\theta_{13}$, which is expected since both parameters enter at leading order in the $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation probabilities. Here the overlap of the reactor measurements (again indicated by the yellow hatched bar) with our marginal posterior for $\sin^2 2\theta_{13}$ (right panel) favors the upper octant over the lower octant when we constrain the results to specifically the upper or lower octant of θ_{23} . Thus, we see that the preference for the upper octant of θ_{23} in Table II is an emergent behavior that arises from the application of the reactor constraint. We also note that though the marginal posterior for $\sin^2 \theta_{23}$, in the top panel, shows a higher posterior density in the lower octant, the total posterior probability integrated across the upper octant is slightly larger

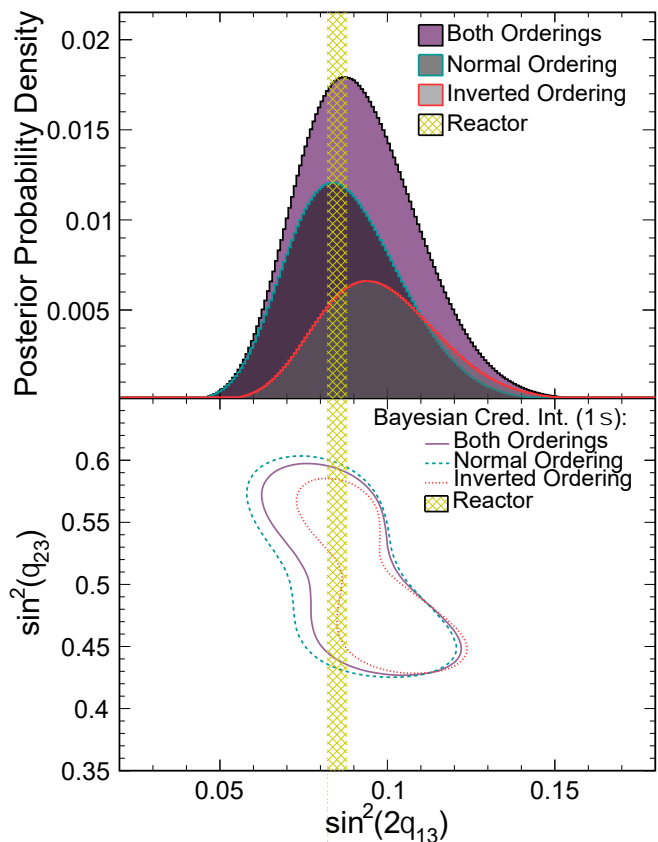


FIG. 12: Comparison of the 2D $\sin^2 2\theta_{13}$ – $\sin^2 \theta_{23}$ credible intervals (bottom) and 1D $\sin^2 2\theta_{13}$ posterior probability densities (top) between both (purple), normal (light-blue) and inverted (red) mass orderings with a prior uniform in $\sin^2 2\theta_{13}$ and $\sin^2 \theta_{23}$. The reactor experiments' 1σ interval in $\sin^2 2\theta_{13}$ from the PDG 2019 [78] is shown in yellow hatched bar.

than the corresponding lower octant probability; but as Table V makes clear, this preference is entirely insignificant.

We emphasize that reactor neutrino experiments and accelerator neutrino experiments measure the PMNS $\sin^2 2\theta_{13}$ by examining different sectors of neutrino oscillations, over a wide range of baselines. Reactor neutrino experiments measure the $\bar{\nu}_e$ survival probability $\mathcal{P}(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ with low-energy (few-MeV) $\bar{\nu}_e$ s over a short (few-km) baseline. Conversely, accelerator neutrino experiments simultaneously measure ν_e appearance in a few-GeV ν_μ beam, and $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam, both over a long (hundreds-of-km) baseline. In long-baseline measurements $\mathcal{P}(\nu_\mu \rightarrow \nu_\mu)$, $\mathcal{P}(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu)$, $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$, and $\mathcal{P}(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ are all exploited to constrain the PMNS oscillation parameters, including $\sin^2 2\theta_{13}$. Thus, the consistency observed between long- and short-baseline measurements lends strong support to the PMNS interpretation of neutrino oscillations.

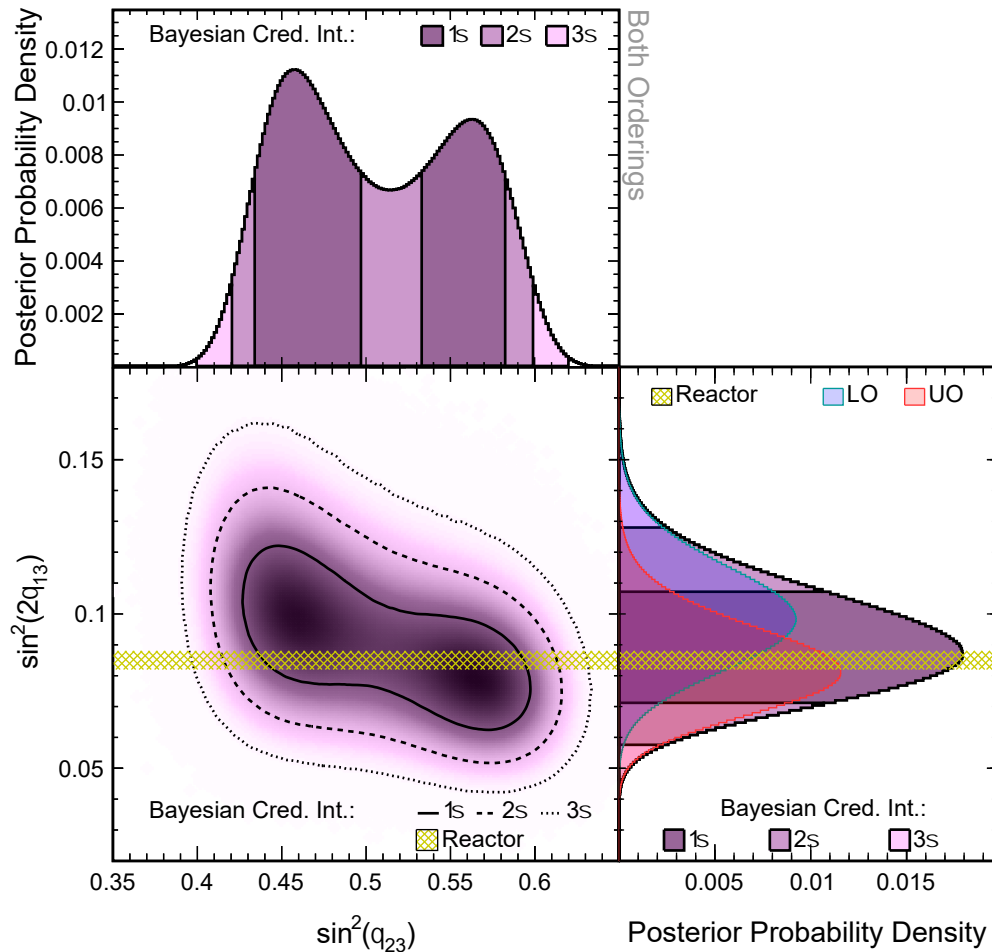


FIG. 13: $\sin^2 \theta_{23}$ – $\sin^2 2\theta_{13}$ posterior probability densities with 1, 2, and 3 σ credible intervals, marginalized over both mass orderings. The posterior density was extracted from a fit with a prior uniform in $\sin^2 2\theta_{13}$. The reactor experiments' 1σ interval in $\sin^2 2\theta_{13}$ from the PDG 2019 [78] is shown in the yellow hatched bar. The right panel shows the posterior probability for $\sin^2 2\theta_{13}$, with its contributions from the upper octant (UO, transparent, red outline) and lower octant (LO, transparent, blue outline) of θ_{23} .

IV. CONCLUSIONS

We have refit the NOvA dataset of 13.6×10^{20} POT in neutrino beam mode and 12.5×10^{20} POT in antineutrino mode using a Bayesian statistical approach. NOvA data continues to be consistent with maximal mixing for $\sin^2 \theta_{23}$ and the regions of the $\sin^2 \theta_{23}$ – Δm_{32}^2 and δ_{CP} – $\sin^2 \theta_{23}$ spaces preferred by this analysis are consistent with those of the previous analysis done using a frequentist fit.

With the introduction of the Bayesian analysis, we also expand the neutrino oscillation parameters measured by NOvA. We report for the first time NOvA measurements that do not require constraints on $\sin^2 2\theta_{13}$ from reactor antineutrino oscillations. Moreover, we are also able to include new results for $\sin^2 2\theta_{13}$ and the Jarlskog invariant J ; these were impractical to produce under the preceding frequentist method due to the necessity of

Feldman–Cousins corrections. The inferences on J provide a model-independent measurement of CP violation and indicate that NOvA data has a weak preference for CP violation, which becomes more pronounced when assuming the inverted mass ordering.

As NOvA measures a convolution of mixing angles, δ_{CP} , and mass ordering in the oscillation probabilities, removing the external constraint on θ_{13} reduces our constraining power to determine the octant of θ_{23} , mass ordering, and δ_{CP} . While our sensitivity is reduced without the external constraint (particularly for the octant), we note that the conclusions arising from the analysis remain unchanged.

This measurement of $\sin^2 2\theta_{13}$ using electron neutrinos and antineutrinos with energies in the GeV range and propagating for hundreds of kilometers is fully consistent with measurements performed using few-MeV electron antineutrinos from nuclear reactors propagating for a few

kilometers. The consistency of results using the PMNS framework across a broad regime of conditions bolsters its applicability. More stringent tests of CP violation and the consistency of our $\sin^2 2\theta_{13}$ measurement with those of reactors will be possible with increased statistics in upcoming NOvA measurements.

V. ACKNOWLEDGMENTS

This document was prepared by the NOvA collaboration using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359. This work was supported by the U.S. Department of Energy; the U.S. National Science Foundation; the Department of Science and Technology, India; the European Research Council; the MSMT CR, GA UK, Czech Republic; the RAS, MSHE, and RFBR, Russia; CNPq and FAPEG, Brazil; UKRI, STFC and the Royal Society, United Kingdom; and the state and University of Minnesota. We are grateful for the contributions of the staffs of the University of Minnesota at the Ash River Laboratory, and of Fermilab. For the purpose of open access, the author

has applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising.

Appendix A: Determining step sizes for ARIA

As noted in Sec. III A 1, an MR²T² chain is derived sample-by-sample using a repeated two-step procedure:

1. **Proposal:** The coordinates of a potential new sample are selected from a probability distribution centered on the current sample (or initial seed).
2. **Acceptance:** The proposal selected above is either accepted or rejected according to the rule of detailed balance, i.e., that every step in the chain be exactly reversible.

If accepted, the proposed coordinates become the next sample. If rejected, the previous sample is repeated to become the next sample.

The MR²T² algorithm does not specify the distribution to be used in step 1 above, however. In our implementation, we use the most common choice, a multivariate Gaussian:

$$g(\vec{x}'|\vec{x}) = (2\pi)^{-\frac{N}{2}} (\det \Sigma)^{-\frac{1}{2}} \exp\left(-\frac{1}{2} (\vec{x}' - \vec{x})^T \Sigma^{-1} (\vec{x}' - \vec{x})\right), \quad (\text{A1})$$

where \vec{x} represents the current sample coordinates, \vec{x}' the proposed next coordinates, and N the dimensionality of the coordinate space. The matrix Σ imposes a length scale on the “distance” between successive samples, and (especially when it is diagonal) its elements are usually called the “step sizes” of the sampling for each degree of freedom. The ideal asymptotic fraction of samples accepted in step 2, α , is 23.4% under a wide range of circumstances [89, 90]. Though this figure is strictly true only for $N \rightarrow \infty$, it has been shown to hold approximately even for parameter counts as low as $N = 5$ [91]. Because the outcome of step 2 is related to the proposals generated in step 1, we tuned the values of Σ to arrive at $\alpha = 23.4\%$.

Our overall heuristic in the tuning procedure is to maintain step sizes that yield similar autocorrelations (defined rigorously below) across all the parameters. This results in the most efficient exploration of the parameter space [71]. We first optimized the step sizes for the parameters of interest, θ_{13} , θ_{23} , $|\Delta m_{32}^2|$, and δ_{CP} . We constructed a chain that sampled only those parameters using a unit matrix for Σ . We computed α for this chain and scaled the relevant elements of Σ in order to arrive at a tolerable preliminary acceptance rate of about 20%. We

then computed the k -lag autocorrelation for each parameter θ , which measures the average correlation between MCMC sample n and sample $n + k$ across all n [92]:

$$r_k = \frac{\sum_{n=1}^{N-k} (\theta_n - \bar{\theta})(\theta_{n+k} - \bar{\theta})}{\sum_{n=1}^N (\theta_n - \bar{\theta})^2}, \quad (\text{A2})$$

where θ_n refers to the value of parameter θ at step n , and $\bar{\theta}$ is its mean value. These autocorrelations are shown in Fig. 14. Using these r_k , we further adjusted the elements of Σ so the oscillation parameters would have similar autocorrelations.

To optimize the step sizes for the nuisance parameters (systematic uncertainties), we constructed a chain sampling only those parameters, again beginning with Σ entries of unity for them. As with the oscillation parameters, we adjusted Σ to ensure none of the nuisance parameters had significantly different autocorrelations from the others. We then constructed a new, much longer chain, and subsequently computed a covariance matrix over the nuisance parameters from it. The Cholesky decomposition of this matrix, L , was used in the next step.

A final chain, this time sampling both oscillation and nuisance parameters, was constructed using the adjusted Σ entries for both. While sampling, the proposed val-




FIG. 14: Lag autocorrelation (see Eq. A2) computed using a MR^2T^2 chain sampling only oscillation parameters.

ues for the nuisance parameters were multiplied by the decomposed covariance scaled by a tunable factor, βL . Using this chain, we recomputed the autocorrelations for all the parameters. The elements of Σ were readjusted to obtain similar autocorrelations from the oscillation parameters, now in the presence of the nuisance parameters. We also adjusted β to yield similar autocorrelations to those of the oscillation parameters. A final global scale was applied to Σ and β to finally arrive at $\alpha = 23.4\%$. The autocorrelations for the oscillation parameters at the end of the tuning procedure are shown in Fig. 15.




FIG. 15: Lag autocorrelation (see Eq. A2) computed using a MR^2T^2 chain after step size tuning with all parameters included.

Appendix B: Determining warmup and thinning fractions for ARIA

The property of sample proportionality to posterior density in the MR^2T^2 method is only guaranteed as

asymptotic behavior. Therefore, it is usually necessary to discard some number of samples N_0 at the beginning of the Markov chain while the chain “burns in” or “warms up.” Moreover, it is usually impossible to find a choice of step sizes Σ in Eq. A1 that proposes samples that are both fully uncorrelated with the previous one and whose acceptance probabilities are high enough to not impose severe computing requirements. We “thin” our chains by discarding all but the k th sample to reduce autocorrelations. The resulting fraction $N_{\text{eff}} = \frac{1}{k}(N - N_0)$, the “number of effective samples,” corresponds to the statistical power of the chain.

Because autocorrelations in our analysis are relatively long (see Fig. 15) and ARIA runs fairly quickly, we produced very long chains of 5×10^6 samples each. Thus, when we compared our posterior densities using $N_0 = 0$, $N_0 = \{1, 3, 5\} \times 10^3$, $N_0 = \{1, 3, 5\} \times 10^4$, and $N_0 = \{1, 3\} \times 10^5$, we found them all to be indistinguishable.

Proper thinning results in minimal autocorrelations. Though there is no universally agreed-upon threshold below which autocorrelations are considered acceptable, we found the plateau in Fig. 15 of around or less than 1% to be sufficient for our needs. Therefore we chose to thin at $k = 10^4$.

Appendix C: Algorithm choices for HMCMC

As noted in Sec. III, HMCMC generates proposals by numerically integrating a Hamiltonian for a fictitious particle, whose potential arises from treating the log-posterior in analogy to gravity:

$$\begin{aligned} \frac{d\vec{q}}{dt} &= \frac{\partial H}{\partial \vec{p}} = \frac{\partial T}{\partial \vec{p}} \\ \frac{d\vec{p}}{dt} &= -\frac{\partial H}{\partial \vec{q}} = -\frac{\partial T}{\partial \vec{q}} - \frac{\partial V}{\partial \vec{q}} \end{aligned} \quad (\text{C1})$$

where T and V are the kinetic and potential energies of the system, respectively.

There are two ingredients of HMCMC left unspecified by the method. In both cases Stan’s default choices were found to be suitable for our needs. The first is the distribution of kinetic energies from which T in Eq. C1 is chosen. Stan’s default is the Euclidean-Gaussian kinetic energy distribution:

$$T(\vec{q}, \vec{p}) = \frac{1}{2} \vec{p}^T M^{-1} \vec{p} + \log |M| + \text{const}. \quad (\text{C2})$$

Here the mass matrix M (analogous to the effect of mass in gravitation) is a parameter that is automatically inferred by Stan during its warm-up sampling by iterative adjustments based on a running covariance over the samples. The second implementation choice is how long the integrator is allowed to run for each particular trajectory. Stan uses an algorithm called No-U-Turns (NUTS) [93], which is a heuristic method that halts integration when two trajectories, extending in each direction from the




FIG. 16: Posterior probability density in Δm_{32}^2 for normal ordering (left) and inverted ordering (right). Samples from Stan are shown as the shaded distributions (1σ , 2σ , 3σ regions indicated by darkest to lightest fill). Samples from ARIA are shown as the red (normal ordering) or blue (inverted ordering) overlaid lines (solid, dashed, dotted lines represent the boundaries of the 1σ , 2σ , 3σ regions, respectively).




FIG. 17: Posterior probability density in $\sin^2 \theta_{23}$ (left) and δ_{CP} (right), marginalized over both mass orderings. Line styles as in Fig. 16.

starting point along the initial momentum, begin to converge towards one another. An upper limit of integrator steps is also supplied as a parameter to Stan; in this analysis, we find that all our trajectories end within 2^{11} steps.

Appendix D: Equivalence of ARIA and Stan results

We extracted the posterior distributions for all of the results shown in this paper using both the ARIA and

Stan samplers described in Sec. III. The posteriors were in all cases nearly indistinguishable, with tiny differences that occasionally caused the boundaries of credible intervals to shift by single bins. This is illustrated in Figs. 16–17.

-
- [1] Y. Fukuda *et al.* (Super-Kamiokande), Phys. Rev. Lett. **81**, 1562 (1998), arXiv:hep-ex/9807003 [hep-ex].
 - [2] S. Fukuda *et al.* (Super-Kamiokande), Phys. Lett. B **539**, 179 (2002), arXiv:hep-ex/0205075 [hep-ex].
 - [3] Q. R. Ahmad *et al.* (SNO), Phys. Rev. Lett. **89**, 011301 (2002), arXiv:nucl-ex/0204008 [nucl-ex].
 - [4] K. Eguchi *et al.* (KamLAND), Phys. Rev. Lett. **90**, 021802 (2003), arXiv:hep-ex/0212021 [hep-ex].
 - [5] D. G. Michael *et al.* (MINOS), Phys. Rev. Lett. **97**, 191801 (2006), arXiv:hep-ex/0607088 [hep-ex].
 - [6] K. Abe *et al.* (T2K), Phys. Rev. Lett. **107**, 041801 (2011), arXiv:1106.2822 [hep-ex].
 - [7] Y. Abe *et al.* (Double Chooz), Phys. Rev. Lett. **108**, 131801 (2012), arXiv:1112.6353 [hep-ex].
 - [8] F. P. An *et al.* (Daya Bay), Phys. Rev. Lett. **108**, 171803 (2012), arXiv:1203.1669 [hep-ex].

- [9] J. K. Ahn *et al.* (RENO), Phys. Rev. Lett. **108**, 191802 (2012), arXiv:1204.0626 [hep-ex].
- [10] R. N. Mohapatra and A. Y. Smirnov, *Elementary particle physics. Proceedings, Corfu Summer Institute, CORFU2005, Corfu, Greece, September 4-26, 2005*, Ann. Rev. Nucl. Part. Sci. **56**, 569 (2006), arXiv:hep-ph/0603118 [hep-ph].
- [11] H. Nunokawa, S. J. Parke, and J. W. F. Valle, Prog. Part. Nucl. Phys. **60**, 338 (2008), arXiv:0710.0554 [hep-ph].
- [12] G. Altarelli and F. Feruglio, Rev. Mod. Phys. **82**, 2701 (2010), arXiv:1002.0211 [hep-ph].
- [13] S. F. King, J. Phys. G **42**, 123001 (2015), arXiv:1510.02091 [hep-ph].
- [14] S. T. Petcov, Eur. Phys. J. C **78**, 709 (2018), arXiv:1711.10806 [hep-ph].
- [15] M. Fukugita and T. Yanagida, Phys. Lett. B **174**, 45 (1986).
- [16] W. Buchmuller and M. Plumacher, Phys. Lett. B **389**, 73 (1996), arXiv:hep-ph/9608308.
- [17] W. Buchmuller, P. Di Bari, and M. Plumacher, Annals Phys. **315**, 305 (2005), arXiv:hep-ph/0401240.
- [18] W. Buchmuller, R. D. Peccei, and T. Yanagida, Ann. Rev. Nucl. Part. Sci. **55**, 311 (2005), arXiv:hep-ph/0502169.
- [19] A. Pilaftsis, Phys. Rev. D **56**, 5431 (1997), arXiv:hep-ph/9707235.
- [20] R. L. Workman *et al.* (Particle Data Group), PTEP **2022**, 083C01 (2022).
- [21] L. Wolfenstein, Phys. Rev. D **17**, 2369 (1978).
- [22] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, JHEP **09**, 178, arXiv:2007.14792 [hep-ph].
- [23] P. Adamson *et al.* (NOvA), Phys. Rev. D **93**, 051104 (2016), arXiv:1601.05037 [hep-ex].
- [24] P. Adamson *et al.* (NOvA), Phys. Rev. Lett. **116**, 151806 (2016), arXiv:1601.05022 [hep-ex].
- [25] P. Adamson *et al.* (NOvA), Phys. Rev. Lett. **118**, 151802 (2017), arXiv:1701.05891 [hep-ex].
- [26] P. Adamson *et al.* (NOvA), Phys. Rev. Lett. **118**, 231801 (2017), arXiv:1703.03328 [hep-ex].
- [27] M. A. Acero *et al.* (NOvA), Phys. Rev. D **98**, 032012 (2018), arXiv:1806.00096 [hep-ex].
- [28] M. A. Acero *et al.* (NOvA), Phys. Rev. Lett. **123**, 151803 (2019), arXiv:1906.04907 [hep-ex].
- [29] M. A. Acero *et al.* (NOvA), Phys. Rev. D **106**, 032004 (2022), arXiv:2108.08219 [hep-ex].
- [30] G. J. Feldman and R. D. Cousins, Phys. Rev. D **57**, 3873 (1998), arXiv:physics/9711021.
- [31] M. A. Acero *et al.* (NOvA), Submitted to *Phys. Rev. D* (2022), arXiv:2207.14353 [hep-ex].
- [32] P. Adamson *et al.*, Nucl. Instrum. Methods Phys Res., Sect. A **806**, 279 (2016), arXiv:1507.06690 [hep-ex].
- [33] S. Agostinelli *et al.* (GEANT4), Nucl. Instrum. Meth. A **506**, 250 (2003).
- [34] Geant4 Collaboration, <https://geant4-data.web.cern.ch/https://geant4-data.web.cern.ch/ReleaseNotes/ReleaseNotes4.10.4.html> (2017).
- [35] L. Aliaga *et al.* (MINERvA), Phys. Rev. D **94**, 092005 (2016), [Addendum: Phys. Rev. D **95**, no.3, 039903 (2017)], arXiv:1607.00704 [hep-ex].
- [36] J. M. Paley *et al.* (MIPP), Phys. Rev. **D90**, 032001 (2014), arXiv:1404.5882 [hep-ex].
- [37] C. Alt *et al.* (NA49), Eur. Phys. J. **C49**, 897 (2007), arXiv:hep-ex/0606028 [hep-ex].
- [38] N. Abgrall *et al.* (NA61/SHINE), Phys. Rev. **C84**, 034604 (2011), arXiv:1102.0983 [hep-ex].
- [39] D. S. Barton *et al.*, Phys. Rev. **D27**, 2580 (1983).
- [40] S. M. Seun, *Measurement of $\pi - K$ ratios from the NuMI target*, Ph.D. thesis, Harvard U. (2007).
- [41] G. M. Tinti, *Sterile neutrino oscillations in MINOS and hadron production in pC collisions*, Ph.D. thesis, Oxford U. (2010).
- [42] A. V. Lebedev, *Ratio of pion kaon production in proton carbon interactions*, Ph.D. thesis, Harvard U. (2007).
- [43] B. Baatar *et al.* (NA49), Eur. Phys. J. **C73**, 2364 (2013), arXiv:1207.6520 [hep-ex].
- [44] P. Skubic *et al.*, Phys. Rev. **D18**, 3115 (1978).
- [45] S. P. Denisov, S. V. Donskov, Yu. P. Gorin, R. N. Krasnokutsky, A. I. Petrukhin, Yu. D. Prokoshkin, and D. A. Stoyanova, Nucl. Phys. **B61**, 62 (1973).
- [46] A. S. Carroll *et al.*, Phys. Lett. **80B**, 319 (1979).
- [47] K. Abe *et al.* (T2K), Phys. Rev. **D87**, 012001 (2013), [Addendum: Phys. Rev. **D87**, no.1, 019902 (2013)], arXiv:1211.0469 [hep-ex].
- [48] T. K. Gaisser, G. B. Yodh, V. D. Barger, and F. Halzen, in *14th International Cosmic Ray Conference (ICRC 1975) Munich, Germany, August 15-29, 1975* (1975) pp. 2161–2166.
- [49] J. W. Cronin, R. Cool, and A. Abashian, Phys. Rev. **107**, 1121 (1957).
- [50] J. V. Allaby *et al.* (IHEP-CERN), Phys. Lett. **30B**, 500 (1969).
- [51] M. J. Longo and B. J. Moyer, Phys. Rev. **125**, 701 (1962).
- [52] B. M. Bobchenko *et al.*, Sov. J. Nucl. Phys. **30**, 805 (1979), [Yad. Fiz. **30**, 1553 (1979)].
- [53] V. B. Fedorov, Yu. G. Grishuk, M. V. Kosov, G. A. Leksin, N. A. Pivnyuk, S. V. Shevchenko, V. L. Stolin, A. V. Vlasov, and L. S. Vorobev, Sov. J. Nucl. Phys. **27**, 222 (1978), [Yad. Fiz. **27**, 413 (1978)].
- [54] R. J. Abrams, R. L. Cool, G. Giacomelli, T. F. Kycia, B. A. Leontic, K. K. Li, and D. N. Michael, Phys. Rev. **D1**, 1917 (1970).
- [55] C. Andreopoulos *et al.*, Nucl. Instrum. Meth. A **614**, 87 (2010), arXiv:0905.2517 [hep-ph].
- [56] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, *The GENIE Neutrino Monte Carlo Generator: Physics and User Manual* (2015), arXiv:1510.05494 [hep-ph].
- [57] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas, Phys. Rev. C **83**, 045501 (2011), arXiv:1102.2777 [hep-ph].
- [58] R. Gran, J. Nieves, F. Sanchez, and M. J. Vicente Vacas, Phys. Rev. D **88**, 113007 (2013), arXiv:1307.8105 [hep-ph].
- [59] B. W. Allardyce *et al.*, Nucl. Phys. A **209**, 1 (1973).
- [60] A. Saunders, S. Hoeibraten, J. J. Kraushaar, B. J. Kriss, R. J. Peterson, R. A. Ristinen, J. T. Brack, G. Hofman, E. F. Gibson, and C. L. Morris, Phys. Rev. C **53**, 1745 (1996).
- [61] O. Meirav, E. Friedman, R. R. Johnson, R. Olszewski, and P. Weber, Phys. Rev. C **40**, 843 (1989).
- [62] S. M. Levenson *et al.*, Phys. Rev. C **28**, 326 (1983).
- [63] D. Ashery, I. Navon, G. Azuelos, H. K. Walter, H. J. Pfeiffer, and F. W. Schlepütz, Phys. Rev. C **23**, 2173 (1981).
- [64] D. Ashery *et al.*, Phys. Rev. C **30**, 946 (1984).
- [65] E. S. Pinzon Guerra *et al.* (DUET), Phys. Rev. C **95**, 045203 (2017), arXiv:1611.05612 [hep-ex].

- [66] A. Aurisano, C. Backhouse, R. Hatcher, N. Mayer, J. Musser, R. Patterson, R. Schroeter, and A. Sousa (NOvA), *Proceedings, 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13-17, 2015*, J. Phys. Conf. Ser. **664**, 072002 (2015).
- [67] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, JINST **11** (09), P09001, arXiv:1604.01444 [hep-ex].
- [68] F. Psihas, E. Niner, M. Groh, R. Murphy, A. Aurisano, A. Himmel, K. Lang, M. D. Messier, A. Radovic, and A. Sousa, Phys. Rev. D **100**, 073005 (2019), arXiv:1906.00713 [physics.ins-det].
- [69] Mr. Bayes and Mr. Price, Philosophical Transactions (1683-1775) **53**, 370 (1763).
- [70] J. S. Speagle, A conceptual introduction to markov chain monte carlo methods (2020), arXiv:1909.12313 [stat.OT].
- [71] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, *Handbook of Markov Chain Monte Carlo* (2011) pp. 1-592.
- [72] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
- [73] W. K. Hastings, Biometrika **57**, 97 (1970).
- [74] J. E. Gubernatis, Phys. Plasmas **12**, 057303 (2005).
- [75] M. Rosenbluth, AIP Conference Proceedings **690**, 22.
- [76] Stan Development Team, *Stan Modeling Language Users Guide and Reference Manual, version 2.26.1*, Stan Development Team (2021).
- [77] M. Betancourt, arXiv:1701.02434 [stat.ME].
- [78] M. Tanabashi *et al.* (Particle Data Group), Phys. Rev. D **98**, 030001 (2018 and 2019 update).
- [79] A. Gelman, X.-L. Meng, and H. Stern, Statistica Sinica **6**, 733 (1996).
- [80] A. Gelman, Y. Goegebeur, F. Tuerlinckx, and I. v. Van Mechelen, Journal of the Royal Statistical Society Series C: Applied Statistics **49**, 247 (2002).
- [81] A. Gelman, Electronic Journal of Statistics **7**, 2595 (2013).
- [82] H. Jeffreys, *Theory of Probability*, 3rd ed. (Oxford, Oxford, England, 1961).
- [83] R. E. Kass and A. E. Raftery, Journal of the American Statistical Association **90**, 773 (1995).
- [84] P. B. Denton and S. J. Parke, Phys. Rev. D **100**, 053004 (2019).
- [85] C. Jarlskog, Phys. Rev. Lett. **55**, 1039 (1985).
- [86] P. B. Denton and R. Pestes, JHEP **05**, 139, arXiv:2006.09384 [hep-ph].
- [87] J. M. Dickey, The Annals of Mathematical Statistics **42**, 204 (1971).
- [88] J. Mulder, E. J. Wagenmakers, and M. Marsman 10.48550/arxiv.2004.09899 (2020).
- [89] A. Gelman, W. R. Gilks, and G. O. Roberts, The Annals of Applied Probability **7**, 110 (1997).
- [90] G. O. Roberts, Stochastics and Stochastic Reports **62**, 275 (1998).
- [91] G. O. Roberts and J. S. Rosenthal, Statistical Science **16**, 351 (2001).
- [92] C. Croarkin and P. Tobias, eds., *NIST/SEMATECH e-Handbook of Statistical Methods* (2013).
- [93] M. D. Hoffman and A. Gelman, Journal of Machine Learning Research **15**, 1593 (2014).