

This is the accepted manuscript made available via CHORUS. The article has been published as:

Inference of bipolar neutrino flavor oscillations near a core-collapse supernova based on multiple measurements at Earth

Eve Armstrong, Amol V. Patwardhan, A. A. Ahmetaj, M. Margarete Sanchez, Sophia Miskiewicz, Marcus Ibrahim, and Ishaan Singh

Phys. Rev. D **105**, 103003 — Published 3 May 2022

DOI: [10.1103/PhysRevD.105.103003](https://doi.org/10.1103/PhysRevD.105.103003)

Inference of bipolar neutrino flavor oscillations near a core-collapse supernova, based on multiple measurements at Earth

Eve Armstrong,^{1,2,*} Amol V. Patwardhan,^{3,†} A.A. Ahmetaj,¹ M. Margarette Sanchez,¹ Sophia Miskiewicz,⁴ Marcus Ibrahim,¹ and Ishaan Singh¹

¹*Department of Physics, New York Institute of Technology, New York, NY 10023, USA*

²*Department of Astrophysics, American Museum of Natural History, New York, NY 10024, USA*

³*SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA*

⁴*Department of Physics, Stevens Institute of Technology, Hoboken, NJ 07030, USA*

(Dated: March 23, 2022)

Neutrinos in compact-object environments, such as core-collapse supernovae, can experience various kinds of collective effects in flavor space, engendered by neutrino-neutrino interactions. These include “bipolar” collective oscillations, which are exhibited by neutrino ensembles where different flavors dominate at different energies. Considering the importance of neutrinos in the dynamics and nucleosynthesis in these environments, it is desirable to ascertain whether an Earth-based detection could contain signatures of bipolar oscillations that occurred within a supernova envelope. To that end, we continue examining a cost-function formulation of statistical data assimilation (SDA) to infer solutions to a small-scale model of neutrino flavor transformation. SDA is an inference paradigm designed to optimize a model with sparse data. Our model consists of two mono-energetic neutrino beams with different energies emanating from a source and coherently interacting with each other and with a matter background, with radially-varying interaction strengths. We attempt to infer flavor transformation histories of these beams using simulated measurements of the flavor content at locations “in vacuum” (that is, far from the source), which could in principle correspond to earth-based detectors. Within the scope of this small-scale model, we found that: (i) based on such measurements, the SDA procedure is able to infer *whether* bipolar oscillations had occurred within the protoneutron star envelope, and (ii) if the measurements sample the full amplitude of the neutrino oscillations in vacuum, then the amplitude of the prior bipolar oscillations is well predicted. This result intimates that the inference paradigm can well complement numerical integration codes, via its ability to infer flavor evolution at physically inaccessible locations.

I. Introduction

The physics of neutrino flavor evolution can significantly influence the dynamics and nucleosynthesis in core-collapse supernovae (CCSN) and neutron star binary mergers [1–11], and the era of multi-messenger astrophysics offers us an unprecedented vantage point on these events. Understanding flavor evolution is critical for leveraging gravitational wave and electromagnetic observations so as to deepen our understanding of energy, entropy, and lepton number transport at these sites. Owing to the large fluxes of neutrinos in these environments, their flavor evolution can be significantly impacted by neutrino coherent forward scattering off of other neutrinos, resulting in a variety of interesting collective phenomena in flavor space (see, e.g., the reviews in [12–16] and references therein).

One important question pertaining to neutrino flavor evolution in a CCSN environment is to ascertain whether “bipolar” oscillations [17–23] occur within the supernova envelope. These are driven by neutrino-neutrino coherent forward scattering, and arise in systems where the

initial state of the interacting neutrino ensemble exhibits a dominance of different flavors at different energies. Bipolar oscillations involve neutrinos at different energies rapidly and repeatedly swapping flavors as they propagate. In numerical solutions of flavor evolution in CCSN environments, neutrinos typically experience these types of oscillations at earlier radii; that is, prior to undergoing the Mikheyev-Smirnov-Wolfenstein (MSW) resonance [24–26], a phenomenon that arises due to coherent forward scattering with matter. These oscillations dramatically change the flavor evolution histories of neutrinos compared to the otherwise-simple MSW-only scenario. As a result of occurring deeper within the supernova envelope (where the neutrino fluxes are higher) compared to the MSW resonance, the effects of bipolar oscillations are potentially more significant with regard to energy transport and nucleosynthesis.

There exist powerful numerical integration codes for obtaining solutions to the flavor evolution problem in compact object environments [27–29]. Utilizing these codes, however, requires making definite choices regarding the relative flavor content of the neutrinos at the point of emission from the proto-neutron star, based on reasonable assumptions about the physics of dense nuclear matter and neutrino decoupling. Typically, the decoupling of neutrinos from chemical and thermal equi-

* evearmstrong.physics@gmail.com

† apatward@slac.stanford.edu

librium is approximated to be instantaneous at the surface of the proto-neutron star, represented by a single, sharp “neutrino-sphere”. As a result, the initial states of neutrinos at the neutrino-sphere radius are taken to be definite flavor states. It has been shown in recent years that relaxing these assumptions regarding uniform, instantaneous neutrino decoupling can result in an emission-angle dependence in the initial flavor content of neutrinos, resulting in the phenomenon of “fast” flavor oscillations (Ref. [16] and references therein).

Moreover, including the effects of direction-changing scattering of neutrinos can result in a small, non-outward-propagating component of the neutrino flux, which can nevertheless significantly contribute to the forward-scattering potential experienced by the outgoing neutrinos, as a result of the large intersection angles between their trajectories [30–33]. This “halo effect” potentially changes how this problem must be approached — not as an initial-value problem with flavor content fully specified at the source surface, but rather as a boundary-value problem, with flavor information propagating both outward and inward. As a result of such discoveries, it has become pertinent to ask how much can be learned about neutrino oscillations near a supernova from a future earth-based neutrino detection, without any *a priori* assumptions about the initial conditions [34].

In this paper, we avoid assumed knowledge of flavor evolution at inaccessible locations within the supernova envelope, by adopting an inverse approach. Using a small-scale model with simulated data, we ask: what information can we *infer* regarding the realm of bipolar oscillations, using measurements made only in the vacuum oscillations regime?

Specifically, we seek to ascertain whether multiple measurements of flavor made in vacuum contain a signature of the flavor evolution history within the supernova envelope, where the neutrino-matter and neutrino-neutrino potentials are dominant. By “multiple” measurements, we mean: measurements spaced out in location but clustered within the vacuum oscillations regime — a proxy for multiple Earth-based detectors. Importantly, the critical differences between this formulation and the forward integration approach are that we do not assume knowledge of (i) unmeasurable model state variables, or (ii) any (measurable or unmeasurable) state variables at physically inaccessible locations. We ask whether the accessible information is sufficient to infer the complete flavor transformation histories of neutrinos back to the emission surface.

To adopt this formalism we employ an inference procedure. Inference is a means to optimize a model given measurements, where the measurements are assumed to arise from model dynamics. Importantly, an inference procedure need not be formulated as an initial-value problem. Rather, we formulate the procedure using

partial information at one bound (near Earth) and zero information at the other (at emission).

The specific inference technique used in this paper is statistical data assimilation (SDA). SDA was invented for numerical weather prediction [35–40] for the case of sparse data. It has since gained traction in neurobiology [41–47], for estimating cellular and synaptic properties given sparse neuronal electrical signals. Within astrophysics, the known applications of SDA include exoplanet modelling [48] and solar cycle prediction [49, 50]. In recent years, the utility of SDA has been explored in the context of inferring solutions to small-scale flavor evolution models [34, 51–53].

In this paper we find, for a small-scale steady-state coherent forward-scattering model of flavor evolution, that multiple measurements of neutrino flavor in the vacuum-oscillations regime could contain a signature of the frequency and amplitude — and to some degree the complexity of the waveform — of bipolar oscillations that had occurred near the point of emission. We quantify the robustness of this result, and discuss implications regarding a real detection.

II. Model

A. Formulation

Our model has been fully described in Refs. [34, 51, 52], and we refer the reader there for details. Here we briefly describe the model’s equations of motion, and note one important feature of the collective neutrino oscillations problems: nonlinearity.

We consider a single-angle, two-flavor scenario wherein two mono-energetic neutrino beams with different energies interact with each other and with a background consisting of particles carrying weak charge, such as nuclei, free nucleons, and electrons. The densities of the background particles and of the neutrino beams dilute as some functions of a position coordinate r , which we interpret as the distance from the neutrino-sphere in a supernova. That is: on their journey through the supernova envelope, the neutrinos interact coherently with each other and with the dense ejecta surrounding the star immediately after core collapse. Importantly, the model is a forward-scattering-only scenario, rendering it solvable via traditional forward-integration techniques — a consistency check for SDA solutions.

We write the equations of motion for flavor evolution of each neutrino in terms of “polarization vectors” \vec{P}_i , after decomposing the density matrices and Hamiltoni-

ans, respectively, into bases of Pauli spin matrices¹ (for details see Ref. [54, 55]):

$$\frac{d\vec{P}_i}{dr} = \left(\Delta_i \vec{B} + V(r)\hat{z} + \mu(r) \sum_{j \neq i} \vec{P}_j \right) \times \vec{P}_i \quad (1)$$

In Equation 1, $\Delta_i = \delta m^2 / (2E_i)$ are the vacuum oscillation frequencies of the two neutrinos with energies E_1 and E_2 , with δm^2 being the mass-squared difference in vacuum. The unit vector $\vec{B} = \sin(2\theta)\hat{x} - \cos(2\theta)\hat{z}$ represents flavor mixing in vacuum, with mixing angle θ . The functions $V(r)$ and $\mu(r)$ are potentials for neutrino-matter and neutrino-neutrino coupling, respectively. They take the forms $V(r) = C_m / (r + r_m)^3$ and $\mu(r) = Q / (r + r_\nu)^4$, respectively; C_m and Q are constant numbers, and r_m and r_ν are offsets which determine the reference values of V and μ at $r = 0$. This form for the neutrino-neutrino coupling reflects the manner in which coupling strength varies in the neutrino bulb model calculations that employ the single-angle approximation.

All the model parameters here are expressed in dimensionless units, but the parameter value choices (shown in Table I) are motivated by certain physical considerations. For instance, the strengths of Q and C_m relative to Δ_i were chosen such that the domain of the solution, chosen here to be $r = [0 : 2]$, would encompass a transition from a matter/neutrino dominated regime to a vacuum oscillation dominated regime. Moreover, the offsets r_m and r_ν were chosen such that $r = 0$ would be near the characteristic radius where the “bipolar” oscillation mode (see Sec. II B) becomes active (for instance, see the definition of R_{coll}^- in Sec. 6 of Ref. [13]). All model parameters are taken to be constant and known to the SDA procedure throughout all of the experimental setups described in this work, with the exception of Sec. V D. We again emphasize that the equations of motion are fiercely nonlinear — and that SDA was designed to perform state-and-parameter estimation for nonlinear models.

The \hat{z} component of the neutrino polarization vector denotes the net flavor content of electron flavor minus “ x ” flavor, the latter being a superposition of muon and tau flavors. In this scenario, we assume that flavor evolution is driven entirely by coherent forward-scattering. At certain distances from the emission surface for each neutrino, the forward scattering potential arising from neutrino-matter and neutrino-neutrino interactions leads

to an in-medium effective neutrino mass level crossing, referred to as the “MSW resonance.” [24–26]. The MSW resonance is associated with an enhanced $e \leftrightarrow x$ flavor conversion probability.

TABLE I: **Model parameters taken to be known and fixed during the estimation procedure.** The Δ_i are the vacuum oscillation frequencies of the neutrinos, and C and Q are the multiplicative factors governing the neutrino-matter, and neutrino-neutrino coupling potentials $V(r)$ and $\mu(r)$, respectively, and r_m and r_ν are the radial offsets. Parameter θ is the mixing angle in vacuum. Each of these parameters are taken to be known in all of the SDA experiments, with the exception of the matter coefficient C being left unknown in one experiment (Section V D), where the SDA procedure was tasked with inferring its value.

Parameter	Value	Parameter	Value
Δ_1	1000	Δ_2	2500
Q	100.0	C_m	3308
r_ν	0.51	r_m	0.50
θ	0.1		

B. BIPOLAR OSCILLATIONS

Collective neutrino oscillations in spherically symmetric models are known to generically exhibit two types of flavor oscillation phenomena: “synchronized” and “bipolar” [13, 19–22]. The synchronized mode is exhibited, for instance, in a dense neutrino gas where all the neutrinos are initially of the same flavor; that is: all the individual polarization vectors P_ω are aligned with one another. This can cause the system to collectively oscillate with a unified angular frequency Ω_{sync} .

The bipolar mode, on the other hand, may be exhibited by systems consisting of Polarization vectors pointing in opposite directions (e.g., ν_e and ν_x), or ν_e and $\bar{\nu}_e$. Bipolar oscillations can be understood analytically by considering a simple toy system of two neutrinos with equal and oppositely aligned polarization vectors. This simple two-neutrino system permits an analogy with an inverted pendulum (the initial state for bipolar oscillations represents an unstable equilibrium configuration). Through this analogy, it can be shown that the characteristic frequency of these bipolar oscillations is $\sim \sqrt{\Delta\mu}$ [20], where Δ and μ are the vacuum oscillation frequency and the neutrino-neutrino interaction strength, as defined in Eq. (1).

¹ The polarization vectors, or Bloch vectors, are defined in terms of the neutrino density matrices: $\rho_i = \frac{1}{2}(\mathbb{1} + \vec{\sigma} \cdot \vec{P}_i)$. The Hamiltonian can be decomposed in the same manner as $H_i = \frac{1}{2}(\text{Tr}(H_i) + \vec{\sigma} \cdot \vec{V}_i)$. Here, \vec{V}_i contains contributions from vacuum oscillations, neutrino-matter interactions, and neutrino-neutrino interactions, as shown in Eq. 1.

III. Method

A. GENERAL FORMULATION

Statistical data assimilation is an inference procedure wherein any measured quantities are assumed to arise from the dynamics of a physical model, which may be nonlinear in nature, and where only a subset of the state variables can be experimentally accessed. This model F can be written as a set of D ordinary differential equations that evolve in some parameterization r as:

$$\frac{dx_a(r)}{dr} = F_a(\mathbf{x}(r), \mathbf{p}(r)); \quad a = 1, 2, \dots, D, \quad (2)$$

where the components x_a of the vector \mathbf{x} are the model state variables. Unknown parameters to be estimated are contained in \mathbf{p} , and may be variable *in general, although for the scope of this work, we take all the parameters to be constant.*

A subset L of the D state variables is associated with measured quantities. One seeks to estimate the evolution of all state variables that is consistent with the measurements provided, to predict model evolution at parameterized locations where measurements are not present.

A prerequisite for estimation using real experimental data is the design of simulated experiments, where the true model evolution is known. Simulated experiments offer the opportunity to ascertain which and how few

experimental measurements and constraints, in principle, are sufficient to complete a model. This is a critical question for cases wherein available measurements are extremely sparse — as will be the case for an Earth-based neutrino detection from a future core-collapse supernova. Finally, in this paper, we use forward integration to generated simulated data, as a consistency check for SDA solutions.

B. OPTIMIZATION FORMULATION

We formulate the SDA procedure as an optimization wherein a cost function is extremized, and we write the cost function in two terms. One term represents “measurement error” the difference between state prediction and any measurements made. The second term represents “model error” the difference between state prediction and adherence to the model dynamics². It will be shown below in this Section that treating the model error as finite offers a systematic method to identify the lowest minimum, in a specific region of state-and-parameter space, of a non-convex cost function. We search the surface of the cost function via the variational method. The procedure in its entirety — that is: a variational approach to minimization coupled with an annealing method to identify a lowest minimum of the cost function — is referred to as variational annealing (VA).

The cost function A_0 used in this paper is written as:

$$\begin{aligned} A_0 &= R_f A_{\text{model}} + R_m A_{\text{meas}} \\ A_{\text{model}} &= \frac{1}{ND} \sum_{n \in \{\text{odd}\}}^{N-2} \sum_{a=1}^D \left[\left\{ x_a(r_{n+2}) - x_a(r_n) - \frac{\delta r}{6} [F_a(\mathbf{x}(r_n), \mathbf{p}) + 4F_a(\mathbf{x}(r_{n+1}), \mathbf{p}) + F_a(\mathbf{x}(r_{n+2}), \mathbf{p})] \right\}^2 \right. \\ &\quad \left. + \left\{ x_a(r_{n+1}) - \frac{1}{2} (x_a(r_n) + x_a(r_{n+2})) - \frac{\delta r}{8} [F_a(\mathbf{x}(r_n), \mathbf{p}) - F_a(\mathbf{x}(r_{n+2}), \mathbf{p})] \right\}^2 \right] \\ A_{\text{meas}} &= \frac{1}{N_{\text{meas}}} \sum_j \sum_{l=1}^L (y_l(r_j) - x_l(r_j))^2. \end{aligned} \quad (3)$$

One seeks the path $\mathbf{X}^0 = \{\mathbf{x}(r_0), \dots, \mathbf{x}(r_N), \mathbf{p}\}$ in state space on which A_0 attains a minimum value. One can

derive this cost function by considering the classical physical Action on a path in a state space, where the path of lowest Action corresponds to the correct solution [56]. Hereafter we shall refer to the cost function of Eq. (3) as the Action. In a previous publication [52], it was shown that the action formulation offers a litmus test for identifying correct solutions: they are solutions that correspond to the path of least action.

A_{model} in Eq. (3) incorporates the model evolution of all D state variables x_a . Here, the outer sum on n is

² In previous works (Refs [51, 52]), the cost function also included an equality constraint to impose unitarity. Eliminating that term has two advantages. One is the easing of the computational burden. The other is that this makes the approach more amenable to flavor evolution studies including the collision terms [53] — a scenario in which unitarity is not necessarily conserved.

taken over all odd-numbered discretized radial locations of the model equations of motion. The sum on a is taken over all D state variables³. In our model, the state variables are all three polarization components for each neutrino beam, or: $D = 6$. A_{meas} governs the transfer of information from measurements y_l to model states x_l . Here, the summation on j runs over all discretized radial locations J at which measurements are made, which may be some subset of all integrated locations of the model. The summation on l is taken over all L measured quantities⁴. In our model, these measured quantities are the P_z component of the polarization vector for each neutrino beam, or: $L = 2$.

The procedure searches a $(D(N + 1) + p)$ -dimensional state space, where N is the number of discretized steps, and p is the number of unknown parameters in the model.

C. Annealing to identify a lowest minimum of the cost function

Our model is nonlinear, and thus the Action surface will be non-convex. The complete VA procedure anneals in terms of the ratio of model and measurement error, R_f and R_m , respectively⁵, to gradually freeze out a lowest-minimum of the Action [57]. This iteration works as follows.

We define the coefficient of measurement error R_m to be 1.0, and write the coefficient of model error R_f as: $R_f = R_{f,0}\alpha^\beta$, where $R_{f,0} = 10^{-1}$, $\alpha = 1.5$, and β is initialized at zero. Parameter β is the annealing parameter. When $\beta = 0$, relatively free from model constraints the Action surface is smooth and convex, and therefore there are no additional local minima. Then we increase the weight of the model term slightly, via an integer increment in β , and recalculate the Action so that the procedure can again be tasked with finding the minimum. We do this recursively toward the deterministic limit of $R_f \gg R_m$. The aim is to remain sufficiently near to the lowest minimum so as not to become trapped in a local minimum as the surface acquires the structure imposed by the model dynamics.

IV. Experiments

A. Specific physics of interest: presence of bipolar oscillations?

Using forward-integration simulations, we permit two neutrino beams of different energies to be emitted from the source (here the “neutrino sphere” of a proto-neutron star) in two different sets of flavor-state initial conditions. In the first set, the two beams are emitted as pure electron-flavor eigenstates. They evolve synchronously and smoothly through the MSW resonance. In the second set, the beams are emitted in nearly opposite polarization states: one pure electron-flavor ($P_z = 1.0$) and the other nearly pure x -flavor ($P_z = -0.8$). This second set gives rise to bipolar oscillations, as described in Section II B.

We seek to examine these two scenarios for the following reason. In a CCSN environment, typically the neutrino flux during the early shock breakout, or “neutronization burst” phase, is dominated by electron neutrinos over all other flavors of neutrinos and anti-neutrinos [58]. Such initial conditions typically give rise to synchronous oscillations. Conversely, at later times during the supernova explosion, such as the neutrino-driven wind phase, neutrinos are emitted in a rough equipartition among flavors, but with different average energies. As a result, there is a dominance of different flavors at different energies in the initial distribution of neutrinos — leading to bipolar oscillations. As the initial conditions will impact the subsequent flavor evolution and nucleosynthesis throughout the envelope, we seek to eliminate *a priori* assumptions and instead ask what information regarding the early flavor evolution is contained in measurable quantities at an Earth-based detector.

The challenge for the SDA procedure is to infer - based on measurements made in vacuum near Earth - which scenario had occurred at earlier radii: the synchronous behavior or the bipolar oscillations. Translating to a larger-scale model, the question will become: *In principle, can multiple measurements of flavor near Earth yield information about the flavor states at earlier radii within the matter-dominated region?*

B. DETAILS OF THE PROCEDURE

In this paper, we give the SDA procedure full knowledge of the model parameters, and measurements were provided at three locations in the vacuum-dominated region. Figure 1 offers a schematic. Within the context of this simple model, by “measurement” we mean the value of $P_{1,z}$ and $P_{2,z}$ – the z -component of the polarization vector for Neutrino Beams 1 and 2, respectively

³ This term can be derived via consideration of Markov-chain transition probabilities [56]. For details, please also refer to Ref [51].

⁴ The measurement error term derives from the mutual information of probability theory [56].

⁵ More generally, R_m and R_f are inverse covariance matrices for the measurement and model errors, respectively. In this paper the measurements are taken to be mutually independent, rendering these matrices diagonal.

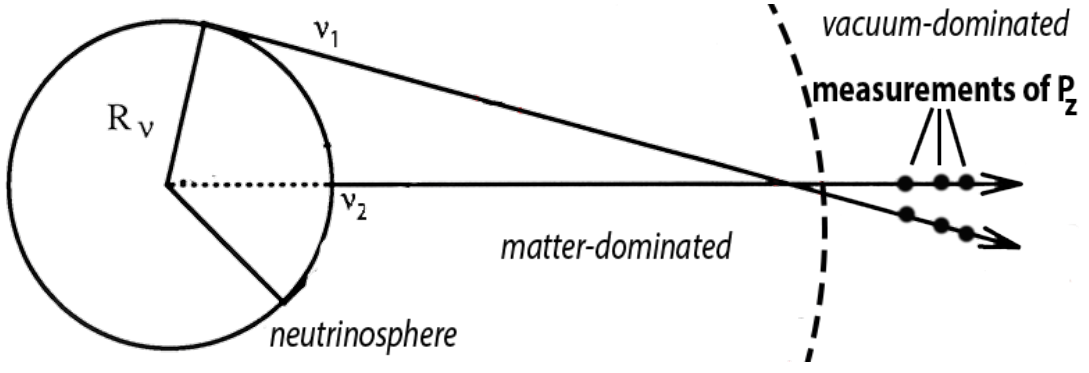


FIG. 1. **Schematic of the small-scale simulation.** Two neutrino beams, ν_1 and ν_2 , emanate from an infinitely sharp "neutrino-sphere," which lies at radius R_v from the center of the proto-neutron star. The beams interact coherently through the matter-dominated envelope, and arrive at Earth. Three detectors, clustered in the vacuum regime, sample the P_z components (i.e. electron-flavor content) of each polarization vector. In a realistic scenario, these three measurement locations represent three satellite- or Earth-based detectors.

⁶. The procedure is provided no information regarding flavor outside of the three locations in vacuum. The task is to take those sparse measurements, together with the model dynamics, to predict the complete flavor evolution history; that is: the values of P_x , P_y , and P_z for each beam at each radial location between emission from the neutrino-sphere at $r = 0$ and detection at $r \sim R$. To obtain the prediction, the SDA procedure is permitted to search the full dynamical range for each variable – of $[-1.0:1.0]$ – at each location⁷. The confidence check on the SDA prediction is simulated "data" generated by forward-integration.

Our specific question is whether vacuum oscillations sampled near Earth contain a signature of whether bipolar oscillations occurred prior to the MSW resonance. To this end, we perform two variations on the experimental design described above.

In the first variation, the simulated data takes as initial conditions (at $r = 0$) $P_{1,z} = 1.0$ and $P_{2,z} = 1.0$ (P_x and P_y are initialized at 0 for both neutrino beams.) Initially aligned in the pure ν_e state, no bipolar oscillation occur, and the trajectory through the MSW resonance is smooth and synchronous. The two beams emerge in nearly-pure ν_x flavor and no appreciable vacuum oscillations occur near the detector (at $r = R$).

In the second variation, the forward integration is instead initialized (at $r = 0$) with $P_{1,z} = +1.0$ and $P_{2,z} = -0.8$ (P_x and P_y are again initialized at 0 for both beams.) These represent two nearly opposite polarization states,

which – as described in Section II B – give rise to bipolar oscillations prior to MSW. The beams then emerge from MSW in mixed states and display high-frequency and high-amplitude vacuum oscillations near the detector (at $r = R$).

In each variation, the SDA procedure is challenged to predict *which behavior had occurred* within the supernova envelope: smooth evolution or bipolar oscillations.

To discretize the neutrino path, we record the output of the simulated forward-integration model at 50,001 discretized steps and a step size δr of 0.00004. The optimization procedure uses the same grid. The units for distance ($r = [0 : 2]$) are arbitrary, in keeping with previous publications [51, 52]. These numbers ensure that bipolar oscillation frequency would be well resolved. Measurements of $P_{1,z}$ and $P_{2,z}$ are taken at the final location ($n = 50,001$) and at two locations within 1,500 steps of that final location. These locations lie sufficiently far beyond the MSW resonance that the vacuum term dominated the Hamiltonian there. The two additional locations are varied, to determine the solution's robustness to the specific choices of locations. In total, we conduct 67 independent experiments, corresponding to 67 distinct choices for the locations of the second and third measurement locations. For each these 67 experiments, four paths are searched, beginning at randomly generated initial conditions for state variables.

The forward integration is performed by Python's odeINT package, which discretizes via an adaptive step size. The optimization is performed by the open-source Interior-point Optimizer (Ipopt) [59], which employs a Hermite-Simpson method of discretization and a constant step size. The discretization of state space, calculations of the model Jacobean and Hessian matrices, and the annealing procedure are performed via an interface with Ipopt that was written in C and Python [60]. Simulations are run on a computing cluster equipped with

⁶ Of course, a real detector will measure a spectrum convolved with contamination.

⁷ In previous publications (Refs [51, 52]), we used search ranges for P_x and P_y that were roughly three times stricter. Broadening those ranges to encompass the full possible dynamics – while it increased the computational expense – has rendered the procedure more robust.

201 GB of RAM and 24 GenuineIntel CPUs (64 bits), each with 12 cores.

V. Result

Key results are as follows:

- Sampling the P_z components of the neutrino beams at multiple vacuum-regime locations reliably predicted *whether* synchronous evolution or bipolar oscillations occurred at earlier radii.
- For the case of bipolar oscillations, and given three measurement locations, the degree to which the measurements sampled the vacuum oscillation amplitude correlated strongly with the strength of the predicted amplitude of earlier bipolar oscillations. Further, if the vacuum oscillation amplitude was well sampled, the Fourier transforms of the evolution of the P_z components at early radii captured some degree of the complexity of the true waveform.
- Using two, rather than three, radial locations, the procedure correctly inferred that bipolar oscillations had occurred, but poorly predicted their amplitude.
- A preliminary examination suggests that performing parameter estimation in addition to state estimation, using multiple measurements in vacuum, will be significantly more challenging than performing state estimation alone.

A. PREDICTION OF SYNCHRONOUS EVOLUTION

The left panel of Figure 2 shows the true (dotted blue) versus predicted (solid) state variable evolution for the case in which initial conditions on $P_{1,z}$ and $P_{2,z}$ were +1.0 and +1.0: pure electron-flavor states. The measured and unmeasured states are in solid red and black, respectively. These initial conditions yield smooth, synchronous evolution through the MSW resonance, as described in Section II. In this case, measurements of $P_{1,z}$ and $P_{2,z}$ were taken at three locations in vacuum.

Here we remind the reader that the measurements used to obtain this prediction were the measurable state variables $P_{1,z}$ and $P_{2,z}$ at three locations near the detector, outside the matter-dominated region; that is: at three out of the 50,001 discretized locations on the path. Given this sparse information – which captured no vacuum oscillations, the procedure correctly inferred that the beams had been emitted in aligned pure states and that no bipolar oscillations had occurred. This result was

robust to ten percent noise added to the measurements of $P_{1,z}$ and $P_{2,z}$ (not shown).

B. PREDICTION OF BIPOLAR OSCILLATIONS

The right panel of Figure 2 shows the true (dotted blue) versus predicted (solid) state variable evolution for one of the 67 experiments with initial conditions on $P_{1,z}$ and $P_{2,z}$ set to +1.0 and -0.8, respectively. In the simulation obtained by forward integration, nearly oppositely-aligned, the beams' interactions yield bipolar oscillations. Given three measurements, which collectively were able to sample the amplitude of vacuum oscillations near Earth, the procedure predicted that bipolar oscillations had occurred at earlier radii. The prediction of the frequency of these bipolar oscillations was robust to ten percent noise added to the measurements of $P_{1,z}$ and $P_{2,z}$ (not shown).

We sought to quantify in more detail the degree to which the structure of bipolar oscillations at early radii was predicted via observations at later radii in vacuum, over all 67 experiments (as noted in Section IV B, the 67 experiments represent 67 distinct choices of two out of three measurement locations: those two lying within 1500 discretized steps of the final location at $n = 50,001$).

First we offer one "good" and one "bad" representative example, out of the 67 total. Figure 3 shows the "good." The right panels in Figure 3 show the flavor evolution in vacuum near the detector ($n = [48500 : 50001]$) for $P_{1,z}$ (top) and $P_{2,z}$ (bottom). True versus predicted are blue and red, respectively; green circles denote the observation locations. The left panels show the corresponding predicted earlier bipolar oscillations ($n = [0 : 1000]$).

Note that together, the three observations (Figure 3 right) capture well the amplitude of vacuum oscillations — and that the corresponding prediction of earlier bipolar oscillation amplitude is strong (Figure 3 left)⁸.

To further analyze the structure of the predicted bipolar oscillations – that is, for $n = [0 : 1000]$, we examined the Fourier decomposition of the evolution of the two P_z components in that region. A Fourier decomposition was called for because those bipolar oscillations may evolve in radius within that region, and may not be strictly sinusoidal.

Figure 4 shows the resulting Fourier power spectrum, for $P_{1,z}$ (top) and $P_{2,z}$ (bottom), where blue and red are true versus predicted, respectively. The predicted amplitude of the strongest harmonic is well matched to

⁸ The difference between true versus predicted initial conditions at $r = 0$ is likely due to different discretization methods employed by the optimization-versus-integration procedures.

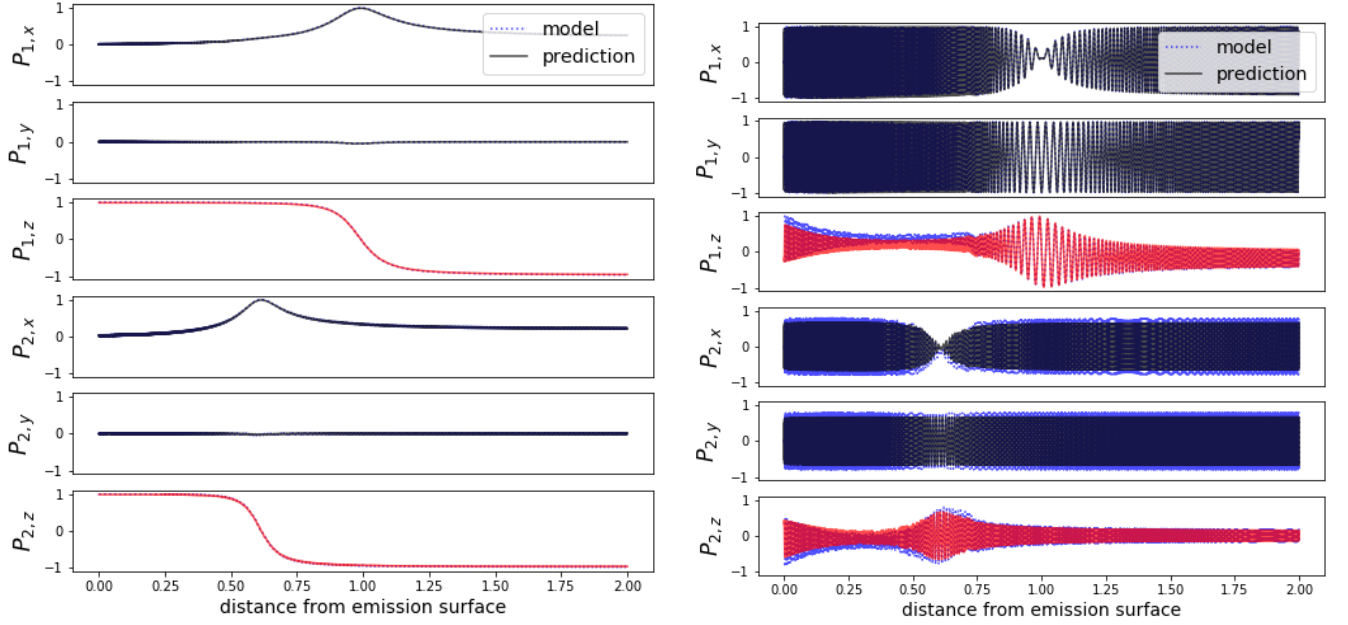


FIG. 2. **True and predicted flavor evolution histories given three measurements of $P_{1,z}$ and $P_{2,z}$ near the detector.** From Top, the columns are: P_x, P_y, P_z for Beam 1, and P_x, P_y, P_z for Beam 2. Black and red lines are predictions for unmeasured and measured state variables, respectively; true model evolution is dotted blue. *Left:* The initial conditions – unknown to the SDA procedure – for $P_{1,z}$ and $P_{2,z}$ were, respectively: +1.0, +1.0, yielding smooth synchronous evolution through the MSW resonance. *Right:* The initial conditions were instead +1.0 and -0.8, yielding bipolar oscillations. The units for distance ($r = [0 : 2]$) are arbitrary, in keeping with previous publications (Refs. [51, 52]).

the true value, as was indicated by the wave-forms at left in Figure 3⁹.

Further, and more interestingly, the predicted Fourier transform captures to some degree the complexity of structure present in the correct solution. First, the peak frequency is not a delta function, but rather has a finite width, indicating that its value is evolving within the range of $n = [0 : 1000]$. Second, the second harmonic is also predicted. (The features at higher frequencies are likely due to differences in the discretization methods used by the Python forward integration versus the optimizer.)

Figure 5 shows a representative "bad" solution out of the 67 total; it is formatted identically to Figure 3. Note that the three measurements (right) poorly sample the vacuum amplitude, and that that poor estimate is reflected in a poor prediction of the bipolar oscillation amplitude (Figure 5 left, top and middle). The Fourier transforms of Figure 6, akin to Figure 4 for the "good"

example, also reflect a poorer match to the power in the first harmonic of the bipolar oscillations waveform, and the complexity of the Fourier decomposition of the true solution is not captured strongly.

These two examples alone suggest a pattern: in the experiments where the multiple measurements in the detector region are able to sample the full extent of the vacuum oscillation amplitude, the prediction of the bipolar oscillation waveform near the source becomes significantly better. To quantify the "goodness" of the bipolar oscillation prediction, the metric that we used was the absolute value of the difference between the peak amplitudes of the strongest frequency in the Fourier transforms (FFTs), of the true and predicted P_z waveforms near the source (the first 1000 grid points), summed over both the neutrino energies. We will call this quantity $\Delta\text{FFT}^{\text{peak, bipolar}}$ and formally write:

$$\Delta\text{FFT}^{\text{peak, bipolar}} = \sum_j \left| \max \left\{ \text{FFT} \left[P_{j,z}^{(\text{true})}(n \in [0 : 1000]) \right] \right\} - \max \left\{ \text{FFT} \left[P_{j,z}^{(\text{pred})}(n \in [0 : 1000]) \right] \right\} \right|, \quad (4)$$

⁹ The difference in the precise value of the predicted versus true peak frequency corresponds to the minimum difference set by the sampling rate. That nonzero difference is likely due to different discretization methods employed by the optimization-versus-integration procedures.

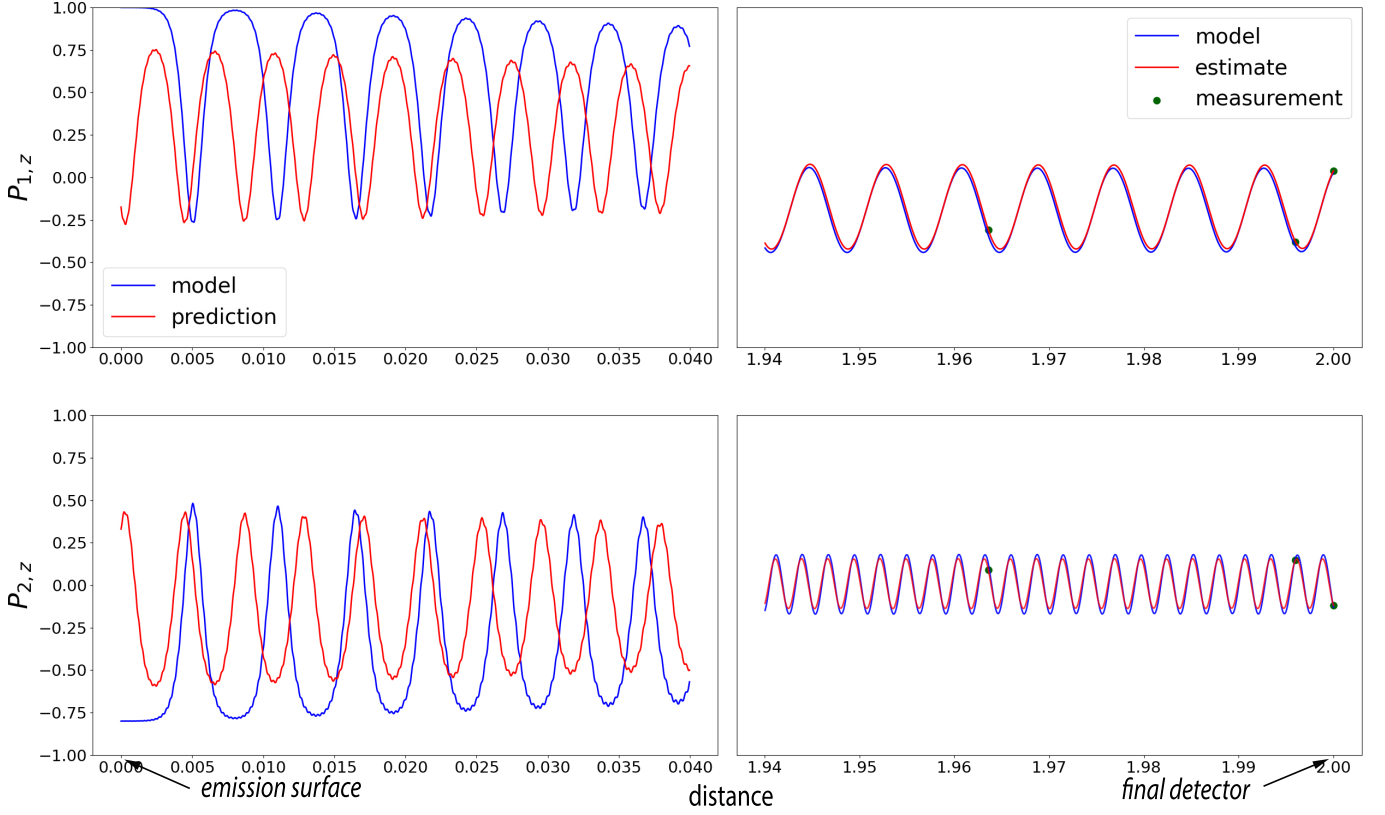


FIG. 3. **Estimate of flavor oscillations in vacuum, juxtaposed with prediction of bipolar oscillations at earlier radii: a "good" example.** Top and bottom: $P_{1,z}$ and $P_{2,z}$, respectively. Right: estimate during observations window in vacuum, with final detector location at far right. Observation locations denoted by green dots; blue and red are true evolution versus estimate, respectively. Left: prediction of bipolar oscillations at earlier radii. Note that adequate sampling of the vacuum oscillations amplitude (right) yields a strong prediction of the bipolar oscillations amplitude at earlier radii (left).

where “max” refers to the maximum strength of the FFT across all Fourier modes (except the zero-frequency mode). We correlated that quantity with the degree to

which the observations captured the vacuum oscillation amplitudes of both neutrinos. Formally, we defined the difference in true versus estimated amplitude near the detector as $\Delta P_z^{ampl, vacuum}$:

$$\Delta P_z^{ampl, vacuum} = \sum_j \left| \text{ampl} \left\{ P_{j,z}^{(true)}(n \in [48500 : 50000]) \right\} - \text{ampl} \left\{ P_{j,z}^{(pred)}(n \in [48500 : 50000]) \right\} \right|, \quad (5)$$

where “ampl” refers to the amplitude of each P_z waveform over the specified domain (in this case, the discrete grid locations numbered 48500 to 50000).

Indeed, across all 67 experiments, we discovered that these two metrics were well correlated. This is shown in Fig. 7, where $\Delta \text{FFT}^{peak, bipolar}$ (y -axis) is plotted against $\Delta P_z^{ampl, vacuum}$ (x -axis).

C. VARYING THE NUMBER OF MEASUREMENTS

Prior to conducting the experiments described above, which employed three measurement locations, we had attempted to find solutions using just one measurement location for the P_z components, and then using two locations. Results were as follows.

Measuring the P_z components only at the final location (at $r = R$) yielded zero model dynamics (not shown). Adding a second measurement location yielded

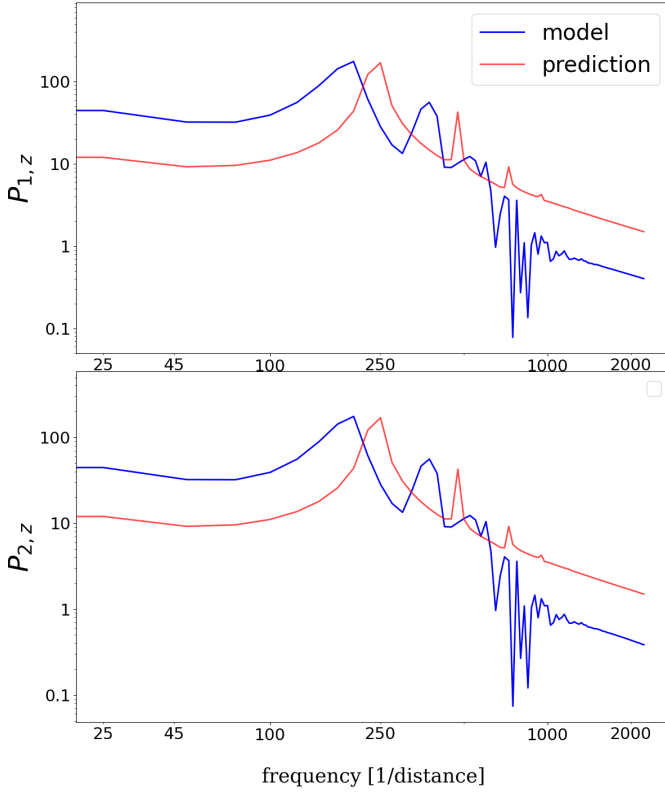


FIG. 4. **The Fourier decomposition of the region of bipolar oscillations shown at left in Figure 3, for $P_{1,z}$ (top) and $P_{2,z}$ (bottom).** Blue and red are true and predicted, respectively. The predicted amplitude of the first harmonic is well matched to the true value, and some complex structure is captured.

a correct inference that bipolar oscillations had occurred, although the amplitude of those oscillations was predicted poorly (not shown). Adding a third measurement significantly improved the prediction of the amplitude of bipolar oscillations, as shown in Figure 2 Right Panel and Figure 3.

To interpret this result, one must recognize that measuring P_z at more than one location yields information about the *derivative* of P_z – which depends on the unmeasurable variables P_x and P_y . As one increases the number of locations at which P_z is measured, one is effectively reconstructing the dynamics of the imaginary P_x and P_y . See Section VI.

D. PREDICTION WITH AN UNKNOWN PARAMETER IN THE MODEL’S MATTER POTENTIAL

With a nonlinear model, rendering a single parameter to be an unknown quantity significantly increases the mathematical challenge for the SDA procedure, compared to state estimation alone. We sought to ascertain how well the SDA design described in this paper would

navigate such an increase in complexity. To this end, we repeated the experiments, this time setting one model parameter to be an unknown quantity to be estimated along with the state variables. We chose as this unknown parameter the constant coefficient C_m in the matter potential (as described in *Model*), because the matter potential is of keen theoretical interest and may impart a signature upon a detection.

For both the cases with and without bipolar oscillations, we initialized ten independent paths, using as measurements the $P_{1,z}$ and $P_{2,z}$ values at three locations near $r = R$, as before. The true value of C_m was 3308.0, and the permitted search range was: [0:10,000.]

For both cases with and without bipolar oscillations, the estimates of C_m were scattered within the permitted search range, with half of the paths estimating the upper bound of 10,000. For each estimate, we confirmed via forward integration that the corresponding state variable evolution was as it should be, were the true value of C_m indeed the estimated value. Importantly, for both cases, the SDA procedure still correctly inferred *whether* bipolar oscillations occurred, and captured the amplitude and frequency of those oscillations as faithfully as it had in the original experiment that had taken C_m to be known. We can conclude that the three measurements of P_z near $r = R$ contain significantly more information about the state variable evolution than they do about the specific strength of the matter potential.

The failure to estimate C_m using measurements in vacuum stems from the degeneracy inherent in these measurements, given the simplicity of the matter potential chosen for the model. Adding additional measurements within the vacuum regime will not reduce this degeneracy. For an additional measurement to help, it would need to be placed within the matter envelope [52] – which would belie the aim of this paper. Alternatively, adding more realistic complexity to the matter potential, for example, including shocks, may improve the ability of in-vacuum measurements to home in on its precise form. Moreover, a detailed study of the procedure’s ability to handle parameter estimation will require an examination of the model’s sensitivity to specific parameter values, and is beyond the scope of this paper.

VI. Discussion

We have made significant progress beyond previous work, having eliminated assumptions regarding flavor content throughout the matter-dominated regime, and instead inferring flavor evolution histories via measurements at accessible locations in vacuum. We have learned that obtaining a measurement of polarization vector component P_z at multiple locations in vacuum yields information about whether bipolar oscillations

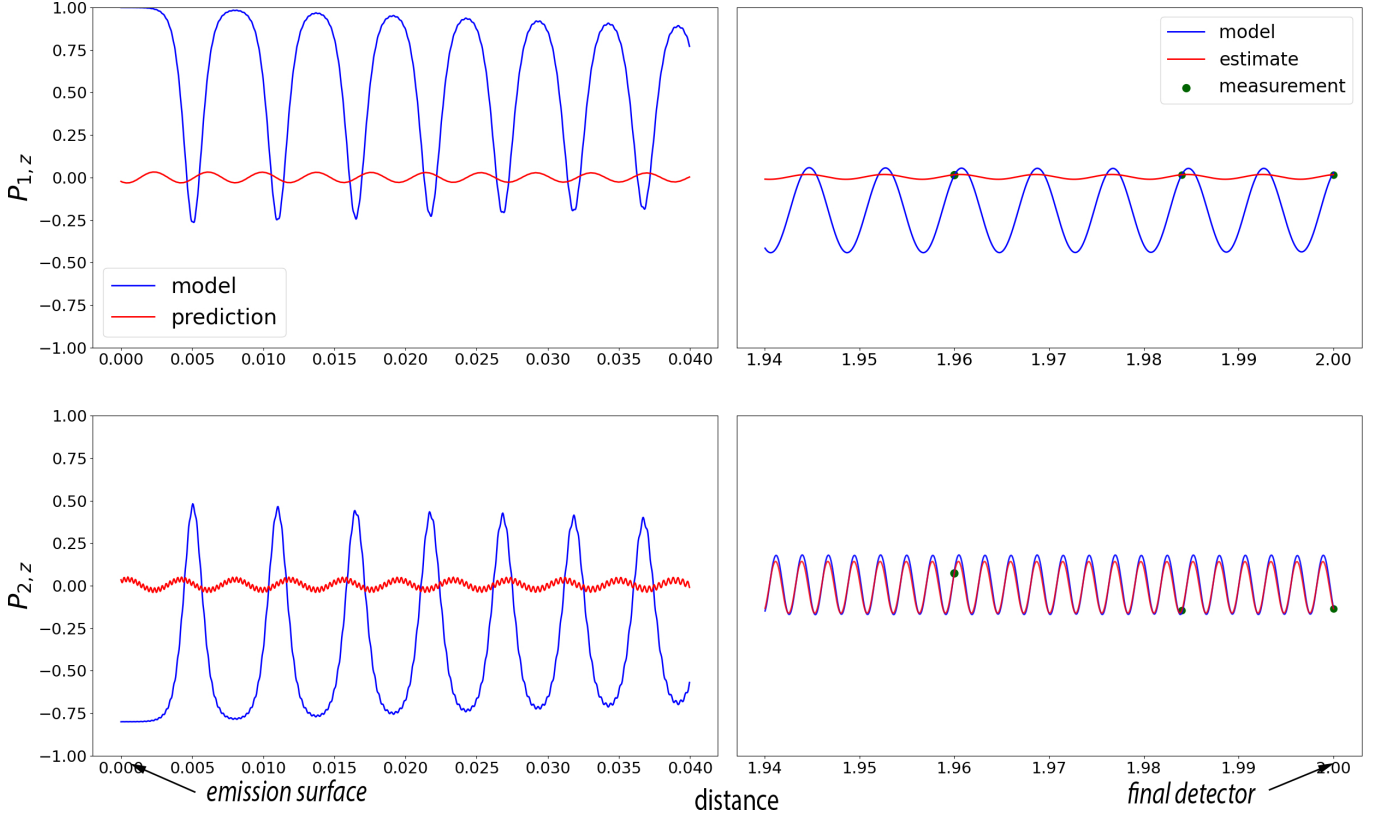


FIG. 5. **Estimate of flavor oscillations in vacuum, juxtaposed with prediction of bipolar oscillations at earlier radii: a "bad" example.** *Top and bottom:* $P_{1,z}$ and $P_{2,z}$, respectively. *Right:* estimate during observations window in vacuum, with final detector location at far right. Observation locations denoted by green dots; blue and red are true evolution versus estimate, respectively. *Left:* prediction of bipolar oscillations at earlier radii. The poor estimate of vacuum oscillation amplitude (right) yields a poor prediction of bipolar oscillation amplitude earlier (left).

occurred at earlier radii, prior to the MSW transition.

A. HOW DO MULTIPLE MEASUREMENTS OF P_z IN VACUUM PREDICT BIPOLAR OSCILLATIONS AT EMISSION?

As noted in Section V B, one measurement in vacuum of the P_z components of the two neutrino beams yielded failed inference of the flavor evolution history. Two measurement locations correctly showed whether bipolar oscillations had occurred at emission, but with poor matches to the oscillation amplitude. Three measurement locations significantly enhanced that amplitude prediction. What is the significance of "at least three measurements"?

This question brings to mind the time-delay embedding theorem from dynamical systems. At the core of that prescription is the notion that one can represent a state space in n variables, or equivalently in one variable at n distinct temporal locations [61]. The concept is intimately related to the information contained in the *derivatives* of a time series [62]. The relevant scenario for our purposes is that, taken together, multiple measure-

ments of P_z represent the derivative of P_z . According to Equation 1, that derivative is dictated in part by the instantaneous values of P_x and P_y . That is: the derivative of P_z contains information about P_x and P_y – and hence *phase information*.

A single measurement location of $P_{1,z}$ and $P_{2,z}$ contains no information regarding the relative phases of the respective polarization vectors. Two measurement locations, however, yield some crude approximation of the derivatives, and hence can reconstruct to some degree the instantaneous values of P_x and P_y . Thus, a pair of measurements of P_z contains some information about the relative phases of the two beams, and hence *whether* bipolar oscillations could have occurred at prior radii.

Adding yet a third measurement further improves the accuracy of the predicted derivative of the P_z components. Having a third measurement greatly increases the likelihood of sampling the amplitude of the vacuum oscillation waveforms of the individual neutrino modes in the detector region. As shown in Fig. 7, sampling the full amplitude of the vacuum oscillations is well-correlated with soundly predicting the amplitude of bipolar oscillations in the source region. Based on this correlation,

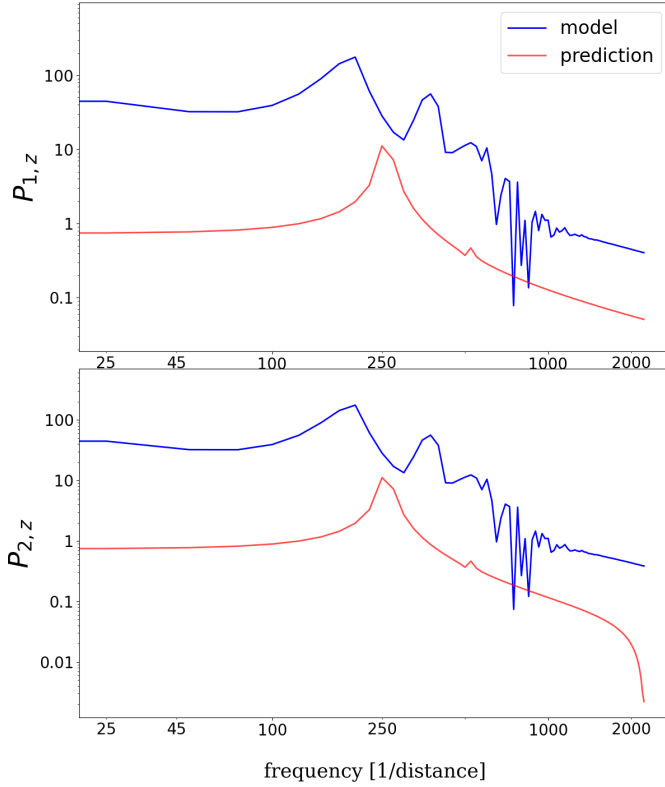


FIG. 6. The Fourier decomposition of the region of bipolar oscillations shown at left in Figure 5. The predicted amplitude of the first harmonic is poor, and complex structure is not captured well.

we might expect that, increasing the number of measurements beyond three would further improve the bipolar oscillation predictions near the source. To this end, we conducted two preliminary tests that employed four and five measurement locations, respectively, using just one set of measurement locations for each test. Both yielded excellent predictions, comparable to the best result obtained over all 67 experiments that had employed three measurements — i.e. Figure 2, Right Panel. This finding is unsurprising: the more independent locations of P_z sampled in vacuum, the more precisely its derivative can be estimated.

Finally, the reader might have noted that, in measuring the P_z vector components at three distinct locations, we

replaced the six boundary conditions used by forward integration by six different boundary conditions. In the case of forward integration, the six are: all measurable (P_z) and unmeasurable state variables (P_x and P_y) for both beams at a single physically inaccessible location ($r = 0$). By contrast, within the SDA formulation, the six boundary conditions were the *measurable* P_z for both beams at three physically accessible regions (in vacuum near $r = R$). In the future, it might prove instructive to formalize a translation between these two formalisms.

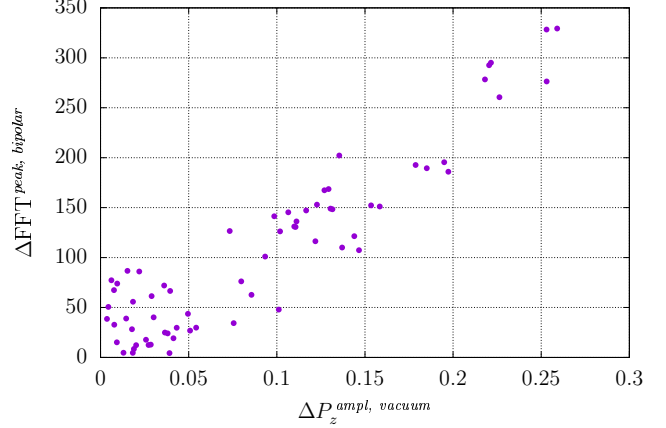


FIG. 7. Correlation between the goodness of the bipolar oscillation prediction, as quantified by the difference between the peak FFT amplitudes of the true and predicted P_z waveforms in the source region ($\Delta\text{FFT}^{\text{peak, bipolar}}$, as defined in Eq. (4)), and the difference between the true and predicted P_z oscillation amplitudes in the detector region ($\Delta P_z^{\text{ampl, vacuum}}$, as defined in Eq. (5)). Each dot on the plot represents one experiment.

VII. ACKNOWLEDGEMENTS

E. A., A. A., M. S., and M. I. acknowledge an Institutional Support for Research and Creativity grant from New York Institute of Technology. E. A. acknowledges NSF grant 2139004. S. M. acknowledges the NSF summer Research Experience for Undergraduates program. The work of A. V. P. was supported by the U.S. Department of Energy under contract number DE-AC02-76SF00515.. As always, eternal thanks to the good people of Doylestown, Ohio.

[1] G. M. Fuller, R. Mayle, B. S. Meyer, and J. R. Wilson, *Astrophys. J.* **389**, 517 (1992).
 [2] Y.-Z. Qian, G. M. Fuller, G. J. Mathews, R. Mayle, J. R. Wilson, and S. E. Woosley, *Phys. Rev. Lett.* **71**, 1965 (1993).
 [3] G. M. Fuller, *Phys. Rept.* **227**, 149 (1993).
 [4] G. M. Fuller and B. S. Meyer, *Astrophys. J.* **453**, 792 (1995).

[5] H. Duan, A. Friedland, G. McLaughlin, and R. Surman, *J. Phys. G* **38**, 035201 (2011), arXiv:1012.0532 [astro-ph.SR].
 [6] M.-R. Wu, Y.-Z. Qian, G. Martinez-Pinedo, T. Fischer, and L. Huther, *Phys. Rev. D* **91**, 065016 (2015), arXiv:1412.8587 [astro-ph.HE].
 [7] M.-R. Wu, G. Martinez-Pinedo, and Y.-Z. Qian, *EPJ Web Conf.* **109**, 06005 (2016), arXiv:1512.03630 [astro-ph.HE].

- [8] H. Sasaki, T. Kajino, T. Takiwaki, T. Hayakawa, A. B. Balantekin, and Y. Pehlivan, Phys. Rev. D **96**, 043013 (2017), arXiv:1707.09111 [astro-ph.HE].
- [9] A. B. Balantekin, AIP Conf. Proc. **1947**, 020012 (2018), arXiv:1710.04108 [nucl-th].
- [10] Z. Xiong, M.-R. Wu, and Y.-Z. Qian, (2019), 10.3847/1538-4357/ab2870, arXiv:1904.09371 [astro-ph.HE].
- [11] Z. Xiong, A. Sieverding, M. Sen, and Y.-Z. Qian, Astrophys. J. **900**, 144 (2020), arXiv:2006.11414 [astro-ph.HE].
- [12] H. Duan and J. P. Kneller, J. Phys. G **36**, 113201 (2009), arXiv:0904.0974 [astro-ph.HE].
- [13] H. Duan, G. M. Fuller, and Y.-Z. Qian, Ann. Rev. Nucl. Part. Sci. **60**, 569 (2010), arXiv:1001.2799 [hep-ph].
- [14] A. Mirizzi, I. Tamborra, H.-T. Janka, N. Saviano, K. Scholberg, R. Bollig, L. Hudepohl, and S. Chakraborty, Riv. Nuovo Cim. **39**, 1 (2016), arXiv:1508.00785 [astro-ph.HE].
- [15] S. Chakraborty, R. Hansen, I. Izaguirre, and G. Raffelt, Nucl. Phys. B **908**, 366 (2016), arXiv:1602.02766 [hep-ph].
- [16] I. Tamborra and S. Shalgar, Ann. Rev. Nucl. Part. Sci. **71**, 165 (2021), arXiv:2011.01948 [astro-ph.HE].
- [17] V. A. Kostelecky and S. Samuel, Phys. Rev. D **52**, 621 (1995), arXiv:hep-ph/9506262.
- [18] S. Samuel, Phys. Rev. D **53**, 5382 (1996), arXiv:hep-ph/9604341.
- [19] H. Duan, G. M. Fuller, and Y.-Z. Qian, Phys. Rev. D **74**, 123004 (2006), arXiv:astro-ph/0511275.
- [20] S. Hannestad, G. G. Raffelt, G. Sigl, and Y. Y. Wong, Phys. Rev. D **74**, 105010 (2006), [Erratum: Phys.Rev.D 76, 029901 (2007)], arXiv:astro-ph/0608695.
- [21] H. Duan, G. M. Fuller, J. Carlson, and Y.-Z. Qian, Phys. Rev. D **74**, 105014 (2006), arXiv:astro-ph/0606616.
- [22] H. Duan, G. M. Fuller, J. Carlson, and Y.-Z. Qian, Phys. Rev. Lett. **97**, 241101 (2006), arXiv:astro-ph/0608050.
- [23] G. G. Raffelt and G. Sigl, Phys. Rev. D **75**, 083002 (2007), arXiv:hep-ph/0701182.
- [24] S. Mikheev and A. Y. Smirnov, arXiv preprint arXiv:0706.0454 (2007).
- [25] S. Mikheev and A. Y. Smirnov, Soviet Journal of Nuclear Physics **42**, 913 (1985).
- [26] L. Wolfenstein, Physical Review D **17**, 2369 (1978).
- [27] H. Duan, G. M. Fuller, J. Carlson, and Y.-Z. Qian, Physical Review D **74**, 105014 (2006).
- [28] H. Duan, G. M. Fuller, and J. Carlson, Computational Science & Discovery **1**, 015007 (2008).
- [29] S. A. Richers, G. C. McLaughlin, J. P. Kneller, and A. Vlasenko, Physical Review D **99**, 123014 (2019).
- [30] J. F. Cherry, J. Carlson, A. Friedland, G. M. Fuller, and A. Vlasenko, Physical review letters **108**, 261104 (2012).
- [31] J. F. Cherry, J. Carlson, A. Friedland, G. M. Fuller, and A. Vlasenko, Phys. Rev. D **87**, 085037 (2013), arXiv:1302.1159 [astro-ph.HE].
- [32] M. Zaizen, J. F. Cherry, T. Takiwaki, S. Horiuchi, K. Kotake, H. Umeda, and T. Yoshida, JCAP **06**, 011 (2020), arXiv:1908.10594 [astro-ph.HE].
- [33] J. F. Cherry, G. M. Fuller, S. Horiuchi, K. Kotake, T. Takiwaki, and T. Fischer, Phys. Rev. D **102**, 023022 (2020), arXiv:1912.11489 [astro-ph.HE].
- [34] E. Rrapaj, A. V. Patwardhan, E. Armstrong, and G. M. Fuller, Physical Review D **103**, 043006 (2021).
- [35] R. Kimura, Journal of Wind Engineering and Industrial Aerodynamics **90**, 1403 (2002).
- [36] E. Kalnay, *Atmospheric modeling, data assimilation and predictability* (Cambridge university press, 2003).
- [37] G. Evensen, *Data assimilation: the ensemble Kalman filter* (Springer Science & Business Media, 2009).
- [38] J. T. Betts, *Practical methods for optimal control and estimation using nonlinear programming*, Vol. 19 (Siam, 2010).
- [39] W. G. Whartenby, J. C. Quinn, and H. D. Abarbanel, Monthly Weather Review **141**, 2502 (2013).
- [40] Z. An, D. Rey, J. Ye, and H. D. Abarbanel, Nonlinear Processes in Geophysics (Online) **24** (2017).
- [41] S. J. Schiff, in *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE, 2009)* pp. 3318–3321.
- [42] B. A. Toth, M. Kostuk, C. D. Meliza, D. Margoliash, and H. D. Abarbanel, Biological cybernetics **105**, 217 (2011).
- [43] M. Kostuk, B. A. Toth, C. D. Meliza, D. Margoliash, and H. D. Abarbanel, Biological cybernetics **106**, 155 (2012).
- [44] F. Hamilton, T. Berry, N. Peixoto, and T. Sauer, Physical Review E **88**, 052715 (2013).
- [45] C. D. Meliza, M. Kostuk, H. Huang, A. Nogaret, D. Margoliash, and H. D. Abarbanel, Biological cybernetics **108**, 495 (2014).
- [46] A. Nogaret, C. D. Meliza, D. Margoliash, and H. D. Abarbanel, Scientific reports **6**, 1 (2016).
- [47] E. Armstrong, Physical Review E **101**, 012415 (2020).
- [48] N. Madhusudhan, “Atmospheric retrieval of exoplanets,” in *Handbook of Exoplanets*, edited by H. J. Deeg and J. A. Belmonte (Springer International Publishing, Cham, 2018) pp. 2153–2182.
- [49] I. Kitiashvili and A. G. Kosovichev, The Astrophysical Journal **688**, L49 (2008).
- [50] I. N. Kitiashvili, in *Solar and Stellar Magnetic Fields: Origins and Manifestations*, Vol. 354, edited by A. Kosovichev, S. Strassmeier, and M. Jardine (2020) pp. 147–156, arXiv:2003.04563 [astro-ph.SR].
- [51] E. Armstrong, A. V. Patwardhan, L. Johns, C. T. Kishimoto, H. D. Abarbanel, and G. M. Fuller, Physical Review D **96**, 083008 (2017).
- [52] E. Armstrong, A. V. Patwardhan, E. Rrapaj, S. F. Ardizi, and G. M. Fuller, Physical Review D **102**, 043013 (2020).
- [53] E. Armstrong, arXiv preprint arXiv:2111.07412 (2021).
- [54] G. Raffelt and G. Sigl, Astroparticle Physics **1**, 165 (1993).
- [55] G. Sigl and G. Raffelt, Nuclear Physics B **406**, 423 (1993).
- [56] H. Abarbanel, *Predicting the future: completing models of observed complex systems* (Springer, 2013).
- [57] J. Ye, D. Rey, N. Kadakia, M. Eldridge, U. I. Morone, P. Rozdeba, H. D. Abarbanel, and J. C. Quinn, Physical Review E **92**, 052901 (2015).
- [58] B. Müller, Annual Review of Nuclear and Particle Science **69**, 253 (2019).
- [59] A. Wächter, in *Dagstuhl Seminar Proceedings* (Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2009).
- [60] “minAone interface with Interior-point Optimizer,” <https://github.com/yejingxin/minAone>, accessed: 2020-05-26.
- [61] D. Ruelle, Publications Mathématiques de l’Institut des Hautes Études Scientifiques **50**, 27 (1979).
- [62] J.-P. Eckmann and D. Ruelle, The theory of chaotic attractors, 273 (1985).