Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **530,** 743–760 (2024) https://doi.org/10.1093/mnras/stae783
Advance Access publication 2024 March 18

Cosmic-Eν: An- emulator for the non-linear neutrino power spectrum

Amol Upadhye,[1,2,3]★ Juliana Kwan,[2,4]★ Ian G. McCarthy [2], Jaime Salcido [2], Kelly R. Moran [5], Earl Lawrence[5] and Yvonne Y. Y. Wong[3]

[1]*South-Western Institute for Astronomy Research, Yunnan University, Kunming 650500, People's Republic of China*
[2]*Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK*
[3]*Sydney Consortium for Particle Physics and Cosmology, School of Physics, The University of New South Wales, Sydney NSW 2052, Australia*
[4]*Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, UK*
[5]*Statistical Sciences Group, CCS Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA*

Accepted 2024 March 8. Received 2024 February 9; in original form 2023 November 27

ABSTRACT

Cosmology is poised to measure the neutrino mass sum M_ν and has identified several smaller-scale observables sensitive to neutrinos, necessitating accurate predictions of neutrino clustering over a wide range of length scales. The `FlowsForTheMasses` non-linear perturbation theory for the the massive neutrino power spectrum, $\Delta^2_\nu(k)$, agrees with its companion N-body simulation at the 10 per cent − 15 per cent level for $k \leq 1\ h\,\mathrm{Mpc}^{-1}$. Building upon the Mira-Titan IV emulator for the cold matter, we use `FlowsForTheMasses` to construct an emulator for $\Delta^2_\nu(k)$, `Cosmic-E`ν, which covers a large range of cosmological parameters and neutrino fractions $\Omega_{\nu,0}h^2 \leq 0.01$ ($M_\nu \leq 0.93$ eV). Consistent with `FlowsForTheMasses` at the 3.5 per cent level, it returns a power spectrum in milliseconds. Ranking the neutrinos by initial momenta, we also emulate the power spectra of momentum deciles, providing information about their perturbed distribution function. Comparing a $M_\nu = 0.15$ eV model to a wide range of N-body simulation methods, we find agreement to 3 per cent for $k \leq 3k_{\rm FS} = 0.17\ h\,\mathrm{Mpc}^{-1}$ and to 19 per cent for $k \leq 0.4\ h\,\mathrm{Mpc}^{-1}$. We find that the enhancement factor, the ratio of $\Delta^2_\nu(k)$ to its linear-response equivalent, is most strongly correlated with $\Omega_{\nu,0}h^2$, and also with the clustering amplitude σ_8. Furthermore, non-linearities enhance the free-streaming-limit scaling $\partial \log(\Delta^2_\nu/\Delta^2_{\rm m})/\partial \log(M_\nu)$ beyond its linear value of 4, increasing the M_ν-sensitivity of the small-scale neutrino density.

Key words: cosmology: theory – large-scale structure of Universe – neutrinos.

1 INTRODUCTION

Cosmology over the next several years will measure the neutrino mass sum $M_\nu = \sum m_\nu$, an as-yet-unmeasured fundamental particle physics parameter, assuming that the dark energy is a cosmological constant (Audren et al. 2013; Chudaykin & Ivanov 2019). Future space-based experiments will provide completely independent bounds on M_ν (Petracca et al. 2016; Lin et al. 2022). However, both forecasts and analyses of the current data are consistent with a weakening of the M_ν bound by a factor of ≈3 when the dark energy equation of state is allowed to vary with time (Font-Ribera et al. 2014; Upadhye 2019; Di Valentino, Melchiorri & Silk 2020). Additionally, neutrinos and dark energy may play roles in the resolution of persistent tensions in measurements of the cosmic expansion and clustering amplitude (Böhringer & Chon 2016; Leauthaud et al. 2017; McCarthy et al. 2018; Poulin et al. 2018; Gogoi et al. 2021; Di Valentino & Melchiorri 2022; McCarthy et al. 2023).

On the particle physics side, persistent anomalies in the neutrino sector motivate models containing additional 'sterile' neutrinos as well as non-standard neutrino interactions (Alvarez-Ruso & Saul-Sala 2021; Aguilar-Arevalo et al. 2022; Denton 2022). Furthermore, other hot dark matter species such as axions could mimic the cosmological effects of massive neutrinos (D'Eramo et al. 2022; Giarè et al. 2022; Di Valentino et al. 2023). A quantitative understanding of neutrino clustering in the non-linear regime will prove invaluable for breaking these degeneracies.

In cosmology, several neutrino clustering signatures deep in the non-linear regime have been identified and quantified through N-body simulations. These include 'wakes' of neutrinos streaming coherently past haloes (Inman et al. 2015; Zhu et al. 2016); an odd-parity contribution to the angular momentum field of galaxies (Yu, Pen & Wang 2019); neutrino-dark-matter relative velocities (Zhu et al. 2014; Inman et al. 2017; Zhu & Castorina 2020; Zhou et al. 2022); modifications to the halo mass function (Costanzi et al. 2013; Yu et al. 2017; Biswas et al. 2019; Bocquet et al. 2020; Ryu & Lee 2022); a neutrino contribution to the scale-dependent bias of dark matter haloes (LoVerde 2014; LoVerde & Zaldarriaga 2014; Chiang et al. 2018; Chiang, LoVerde & Villaescusa-Navarro 2019; Banerjee et al. 2020); and inhomogeneities in the cosmic neutrino background detectable in laboratory searches (Baracchini et al. 2018; Betti et al. 2019; Aker et al. 2022). Since the systematic bias associated with these phenomena are substantially different from those of joint analyses of large-scale cosmic surveys, some of them may be decisive to a convincing detection of massive neutrinos.

★ E-mail: a.upadhye@ynu.edu.cn (AU); jk945@cam.ac.uk (JK)

Accurate theoretical modelling of non-linear neutrino clustering effects, possibly through an extension of the halo model to neutrino clustering, will require fast and reliable calculations of the neutrino power spectrum. A recent code comparison by the Euclid collaboration, Adamek et al. (2023), tested a wide range of simulation methods (specifically, Teyssier 2002; Springel 2005; Springel et al. 2008; Beck et al. 2016; Adamek et al. 2016a, b; Adamek, Durrer & Kunz 2017; Dakin et al. 2019; Springel et al. 2021; Dakin, Hannestad & Tram 2022; Marin-Gilabert et al. 2022; Mauland et al. 2023; Schaller et al. 2023) and found agreement at the 30 per cent − 40 per cent level over the range $k \leq 1\ h\,\mathrm{Mpc}^{-1}$ of wave numbers. However, simulations are computationally expensive, motivating an exploration of alternative methods with comparable accuracies.

Linear perturbative calculations of massive neutrino clustering can be carried out to high precision by tracking the evolution of the neutrino distribution function (Ma & Bertschinger 1995) in both position and momentum. An alternative approach introduced by Dupuy & Bernardeau (2014, 2015a, b) discretizes the Fermi-Dirac distribution of initial neutrino velocities. Each neutrino fluid, defined by its initial velocity, can then be treated as a separate fluid individually obeying the continuity and Euler equations of fluid dynamics. This approach is Lagrangian in momentum space, since neutrinos cannot move from one initial-velocity bin to another. Since non-linear cosmological perturbation theory begins with the continuity and Euler equations, this momentum-Lagrangian method is a natural starting point for non-linear neutrino perturbation theories.

`FlowsForTheMasses`, the first non-linear perturbative power spectrum calculation for free-streaming particles such as massive neutrinos, began with precisely this approach (Chen, Upadhye & Wong 2023b). Since a fluid with non-zero initial velocity $\vec{v}$ has a preferred direction $\hat{v}$, Chen et al. (2023b) began by extending the Time-Renormalization Group perturbation theory of Pietroni (2008); Lesgourgues et al. (2009) to fluids with homogeneous initial velocities. Their Fourier-space clustering depends not only upon the magnitude of the Fourier vector $\vec{k}$ but also its angle with $\hat{v}$, whose cosine is $\mu = \hat{k} \cdot \hat{v}$.

Expanding the density contrast and velocity divergence in Legendre polynomials in μ, Chen et al. (2023b) showed that the mode-coupling integrals of non-linear perturbation theory couple different Legendre moments, drastically increasing the computational cost. However, by applying Fast Fourier Transform (FFT) techniques introduced by McEwen et al. (2016), Schmittfull, Vlah & McDonald (2016), and Fang et al. (2017) to its mode-coupling integrals, Chen et al. (2023b) was able to accelerate them by more than two orders of magnitude. The resulting `FlowsForTheMasses` perturbation theory can compute a non-linear massive neutrino power spectrum with reasonable accuracy settings at relatively low computational cost.[1]

Though the computational expense of `FlowsForTheMasses` is much lower than that of N-body neutrino simulations, it remains too high to be included in cosmological data analysis pipelines; the 50-flow production runs of Chen et al. (2023b) each took about a day on a 32-core machine. A machine learning technique known as emulation, introduced into cosmology by Heitmann et al. (2009, 2010) and Lawrence et al. (2010), is ideal for quickly approximating expensive functions. Emulation begins with a training set of evaluations of an expensive function (in our case, the `FlowsForTheMasses` neutrino power spectrum) the size of which is determined by computational budget and required level of accuracy. To mitigate the computational expense associated with emulating multivariate data, the spectra are represented via a principal component (PC) decomposition. Gaussian process (GP) models are then used to model the cosmology-dependent PC weights, enabling fast prediction at new cosmologies.

Our goal in this work is an emulator of the `FlowsForTheMasses` non-linear neutrino power spectrum. Furthermore, since `FlowsForTheMasses` already divides neutrinos by their initial momenta and tracks the evolution of each one separately, our training set includes individual neutrino momentum deciles at no extra computational cost. For both the $z \leq 2$ CDM + baryon power spectra and the emulator design points, we use the Mira-Titan IV (MT4) emulator of Moran et al. (2023). We demonstrate that our emulator, `Cosmic-Eν`, precisely reproduces `FlowsForTheMasses` to < 3.5 per cent for $10^{-3}\ h\,\mathrm{Mpc}^{-1} \leq k \leq 1\ h\,\mathrm{Mpc}^{-1}$ and $0 \leq z \leq 3$.[2]

Next, we compare `Cosmic-Eν` to a range of N-body neutrino simulations. We begin with the Euclid code comparison of Adamek et al. (2023) for $M_\nu = 0.15$ eV, which runs simulations with a variety of resolutions, box sizes, and massive neutrino implementations. Compared with their highest-resolution `SWIFT` simulation of Schaller et al. (), `Cosmic-Eν` is accurate to 3 per cent up to $k = 3k_{\mathrm{FS}} = 0.17\ h\,\mathrm{Mpc}^{-1}$, 19 per cent to $k = 0.4\ h\,\mathrm{Mpc}^{-1}$, and 49 per cent to $k = 1\ h\,\mathrm{Mpc}^{-1}$, all of which are somewhat larger than but comparable to the scatter among very different simulation methods. This qualitative picture is unchanged when M_ν is raised to 0.3 and 0.6 eV in the Adamek et al. (2023) simulations, and when the dark energy is allowed to vary rapidly in one of our own simulations (described below).

Finally, we employ `Cosmic-Eν` to study enhancements to the ratios $\Delta^2_\nu/\Delta^2_\nu[\mathrm{LR}]$ and $\Delta^2_\nu/\Delta^2_\mathrm{m}$ due to the non-linear clustering of massive neutrinos. The first of these, the non-linear enhancement of neutrino clustering relative to linear response (LR), was considered in Chen et al. (2023b). After confirming the accuracy of `Cosmic-Eν` for this quantity, we quantify the sensitivity of $\Delta^2_\nu/\Delta^2_\nu[\mathrm{LR}]$ to each of the eight cosmological parameters, showing that $\Omega_{\nu,0}h^2$ is by far the most significant for determining the neutrino clustering enhancement, followed by $\Omega_{\mathrm{b},0}h^2$, $\Omega_{\mathrm{m},0}h^2$, and σ_8. The second ratio, $\Delta^2_\nu/\Delta^2_\mathrm{m}$, was shown by Ringwald & Wong (2004) and Wong (2008) to scale as the fourth power of M_ν, hence $\Omega_{\nu,0}h^2$, for LR neutrinos in the free-streaming limit. We confirm this result in the linear case, then show that non-linear corrections enhance this scaling relation. For example, at $k = 1\ h\,\mathrm{Mpc}^{-1}$ and $\Omega_{\nu,0}h^2 = 0.002$, $\Delta^2_\nu/\Delta^2_\mathrm{m} \sim M_\nu^{4.5}$.

This study is organized as follows. Section 2 briefly describes the emulation procedure and the `FlowsForTheMasses` perturbation theory. Our emulator training set is assembled in Section 3 after improving the high-M_ν numerical stability of `FlowsForTheMasses`. Section 4 constructs the `Cosmic-Eν` emulator and quantifies its accuracy with respect to `FlowsForTheMasses`. `Cosmic-Eν` is then compared with a wide variety of N-body simulation methods in Section 5. Finally, Section 6 quantifies the non-linear enhancements to the $\Delta^2_\nu/\Delta^2_\nu[\mathrm{LR}]$ and $\Delta^2_\nu/\Delta^2_\mathrm{m}$ ratios, and Section 7 concludes.

[1] The `FlowsForTheMasses` perturbation theory code is publicly available at github.com/upadhye/FlowsForTheMasses.

[2] `Cosmic-Eν` is publicly available at `github.com/upadhye/Cosmic-Enu`.

2 BACKGROUND

2.1 Emulation

A thorough discussion of emulation in cosmology may be found in Heitmann et al. (2009). Here, we briefly summarize their procedure, with slight differences in notation.

Suppose that we wish to approximate a dimensionless function $\mathscr{P}(k, z, \vec{C})$ of the wave number k, redshift z, and cosmological parameters $\vec{C}$. This may be proportional to the power spectrum itself, or a function of the power spectrum chosen to reduce its dynamic range. At each of $m \in [0, N_\mathrm{M})$ cosmological models defined by parameters $\vec{C}^*_m$, we are given $\mathscr{P}$ at each of $i \in [0, N_{kz})$ points (k_i, z_i), that is, $\mathscr{P}^*_{im} = \mathscr{P}(k_i, z_i, \vec{C}^*_m)$. We seek an approximation of the form

$$\mathscr{P}(k_i, z_i, \vec{C}) \approx \mu^*(k_i, z_i) + \sigma^* \sum_{j=0}^{N_\mathrm{PC}-1} w_j(\vec{C})\, \phi_j(k_i, z_i) \tag{1}$$

$$\mu^*(k_i, z_i) := \frac{1}{N_\mathrm{M}} \sum_{m=0}^{N_\mathrm{M}-1} \mathscr{P}^*_{im} \tag{2}$$

$$(\sigma^*)^2 := \frac{1}{N_\mathrm{M} N_{kz}} \sum_{m=0}^{N_\mathrm{M}-1} \sum_{i=0}^{N_{kz}-1} \left(\mathscr{P}^*_{im} - \mu^*(k_i, z_i)\right)^2 \tag{3}$$

where μ^*_i is the mean input $\mathscr{P}^*_{im}$ across cosmologies, the ϕ_j are a set of N_PC orthogonal basis functions to be defined below, and $w_j(\vec{C})$ are the corresponding basis weights. If the number of bases N_PC is chosen equal to the number of training models N_M minus one[3], then the approximation of equation (1) can be made exact for the N_M input models, though we will typically choose N_PC smaller than this.

Let $\mathscr{D}^*_{im} = (\mathscr{P}^*_{im} - \mu^*_i)/\sigma^*$, so $\mathscr{D}^*_{im}$ for fixed i has zero mean by construction, and $\mathscr{D}^*_{im}$ is an $N_{kz} \times N_\mathrm{M}$ matrix. By means of a compact singular value decomposition, we may write

$$\mathscr{D}^*_{im} = \sum_{m'=0}^{N_\mathrm{M}-1} \sum_{m''=0}^{N_\mathrm{M}-1} U_{im'} D_{m'm''} (V^T)_{m''m} \tag{4}$$

where U is an $N_{kz} \times N_\mathrm{M}$ orthogonal matrix, D an $N_\mathrm{M} \times N_\mathrm{M}$ diagonal matrix, and V an $N_\mathrm{M} \times N_\mathrm{M}$ orthonormal matrix. In terms of these three matrices, we may write the PC basis functions $\phi_j(k_i, z_i)$ and the weights of the input data w^*_{jm} as

$$\phi_{ij} := \phi_j(k_i, z_i) = \frac{1}{\sqrt{N_\mathrm{M}}} \sum_{m=0}^{N_\mathrm{M}-1} U_{im} D_{mj} \tag{5}$$

$$w^*_{jm} = \sqrt{N_\mathrm{M}} V_{mj} \tag{6}$$

where the j index may be truncated to $j < N_\mathrm{PC}$ for any chosen $N_\mathrm{PC} < N_\mathrm{M}$. Since the functions $\phi_j(k, z)$ may be interpolated from the ϕ_{ij} using standard methods, our remaining task is to model the weight functions $w_j(\vec{C})$ using w^*_{jm}.

We model each $w_j(\vec{C})$ using a GP. A GP is an infinite dimensional generalization of a multivariate Gaussian distribution, in which any finite set of random variables is defined to follow a multivariate Gaussian distribution specified by a mean function and a covariance function (Williams & Rasmussen 2006). We define the GP over each $w_j(\vec{C})$ to have a mean function of 0 and a Gaussian correlation function $R_j(\vec{C}, \vec{C}', \vec{\beta}_j)$. This correlation function is specified by a set of correlation hyperparameters $\beta_{j\ell}$, one for each of the N_C cosmological parameters of $\vec{C}$:

$$R_j(\vec{C}, \vec{C}', \vec{\beta}_j) = \prod_{\ell=0}^{N_\mathrm{C}-1} \exp\left(-\beta_{j\ell}(C_\ell - C'_\ell)^2\right) \tag{7}$$

The input data weights w^*_{jm} are now assumed to arise from the following hyperparameter-dependent probability distribution:

$$p(w^*_j | \lambda_{\mathrm{W},j}, \lambda_{\mathrm{U},j}, \vec{\beta}_j) = \frac{\exp\left[-\frac{\lambda_{\mathrm{W},j}}{2} \sum_{m=0}^{N_\mathrm{M}-1} \sum_{n=0}^{N_\mathrm{M}-1} w^*_{jm} \mathscr{R}(\vec{\beta})^{-1}_{jmn} w^*_{jn}\right]}{\sqrt{\det \mathscr{R}} (2\pi/\lambda_{\mathrm{W},j})^{N_\mathrm{M}/2}} \tag{8}$$

$$\mathscr{R}_{jmn}(\lambda_{\mathrm{W},j}, \lambda_{\mathrm{U},j}, \vec{\beta}_j) = \lambda^{-1}_{\mathrm{U},j} R_j(\vec{C}^*_m, \vec{C}^*_n, \vec{\beta}_j) + \lambda^{-1}_{\mathrm{W},j} \delta^{(\mathrm{K})}_{mn}, \tag{9}$$

where $\delta^{(\mathrm{K})}_{mn}$ is the Kronecker delta. Heitmann et al. (2009) has included an additional set of hyperparameters, a scaling term $\lambda_{\mathrm{U},j}$ and a 'nugget' term $\lambda_{\mathrm{W},j}$. The former scales the correlation function into a covariance function, while the latter accommodates slight numerical fluctuations in the computation of $\mathscr{P}^*_{im}$. For each PC j, we thus have one $\lambda_{\mathrm{U},j}$, one $\lambda_{\mathrm{W},j}$, and N_C different $\beta_{j\ell}$, for a total of $N_\mathrm{C} + 2$ hyperparameters.

Since we are given w^*_{jm} but not the hyperparameter values, our next step is to find the hyperparameter values most consistent with w^*_{jm}. We do so using a 50 000-step Markov chain Monte Carlo sampling through the `SEPIA` code of Gattiker et al. (2020)[4], using the default hyperparameter priors and bounds defined in `SEPIA`. Let $\hat{\lambda}_{\mathrm{U},j}$, $\hat{\lambda}_{\mathrm{W},j}$, and $\hat{\vec{\beta}}_j$ be the posterior mean values of these hyperparameters.

We have now arrived at our goal, a predictive model for the weights $w_j(\vec{C})$ in equation (1) for a given cosmology $\vec{C}$. Let $\hat{\Theta} = \{\hat{\vec{\beta}}_j, \hat{\lambda}_{\mathrm{U},j}, \hat{\lambda}_{\mathrm{W},j}\}$. Using these optimal hyperparameter values, we specify the conditional Gaussian probability distribution of each of the weights as follows:

$$w_j(\vec{C} \mid w^*_j, \hat{\Theta}) \sim N(\bar{W}_j(\vec{C}), \bar{\Sigma}_j(\vec{C})), \tag{10}$$

$$\bar{W}_j(\vec{C} \mid w^*_j, \hat{\Theta}) = [\vec{r}^*_j(\vec{C})]^T [\hat{\mathscr{R}}^*_j]^{-1} \vec{w}^*_j, \tag{11}$$

$$\bar{\Sigma}_j(\vec{C} \mid \hat{\Theta}) = r_j(\vec{C}) - [\vec{r}^*_j(\vec{C})]^T [\hat{\mathscr{R}}^*_j]^{-1} \vec{r}^*_j(\vec{C}), \tag{12}$$

$$r^*_{jm}(\vec{C} \mid \hat{\Theta}) = \hat{\lambda}^{-1}_{\mathrm{U},j} R_j(\vec{C}^*_m, \vec{C}, \hat{\vec{\beta}}_j). \tag{13}$$

Here, $\vec{w}^*_j$ is the $N_M \times 1$ vector of observed weights, $\hat{\mathscr{R}}^*_j$ is the $N_M \times N_M$ matrix of realizations of equation (9) evaluated at the optimal values of the hyperparameters, $\vec{r}^*_j(\vec{C})$ is an $N_M \times 1$ vector having mth entry $r^*_{jm}(\vec{C})$, and $r_j(\vec{C}) = \hat{\lambda}^{-1}_{\mathrm{U},j} R_j(\vec{C}, \vec{C}, \hat{\vec{\beta}}_j) + \hat{\lambda}_{\mathrm{W},j}$. Equation (10) indicates that w_j is drawn from a normalized Gaussian distribution of mean $\bar{W}_j(\vec{C})$ and standard deviation $\bar{\Sigma}_j(\vec{C})$. As in Heitmann et al. (2009), we use $\bar{W}_j(\vec{C})$ as our emulated prediction of $w_j(\vec{C})$. We see from equation (11) that this mean value is a weighted average of the observed weights $\vec{w}^*_j$, with weights determined by the covariance between the observed and predictive cosmological parameters.

Thus far, we have assumed the initial data $\mathscr{P}^*_{im}$ to be given. Choosing the N_M input models used to train the emulator is a separate problem known as emulator design. The goal of emulator design is to cover the given parameter space to a specified accuracy using the smallest number N_M of input cosmological models, since each model is computationally expensive. A simple grid in parameter space is one of the least efficient designs for this purpose. Heitmann et al. (2009, 2016), Lawrence et al. (2017), and Moran et al. (2023)

[3]Including the mean term reduces the remaining degrees of freedom in the model by one.

[4]`SEPIA` is publicly available at github.com/lanl/SEPIA.

employ efficient space-filling Latin hypercube designs or similar nested, space-filling lattices; the latter is useful when runs are to be done in 'batches' so that interim analyses may be performed as partial results are available. The training sets we construct, described in Section 3.1, are built on the design choices made by these authors.

The batched MT4 emulator design allows for the addition of more design points to improve its accuracy. We leave this possibility open for future work. However, we will see in Section 4.2 that even with the existing MT4 design, `Cosmic-Eν` attains an accuracy of ≈ 3.5 per cent, which is subdominant to the 14 per cent error in `FlowsForTheMasses` itself, as measured by Chen et al. (2023b). Thus we do not pursue here the possibility of including additional design points.

2.2 Non-linear perturbation theory for neutrinos

We are particularly interested in the non-linear growth of the neutrino density perturbations, which occurs at late times and at scales well within the Hubble horizon. Thus, to excellent approximation, we assume that all matter in the Universe obeys the scalarized non-relativistic continuity, Euler, and Poisson equations in a box expanding uniformly at a rate given by the time-dependent Hubble parameter. General Relativistic clustering including vector and tensor perturbations, as well as multiple fluids, has been considered previously (e.g. Hwang & Noh 2006a, b, 2007; Jeong et al. 2011; Hwang & Noh 2013a, b; Adamek, Durrer & Kunz 2014; Yoo 2014; Yoo & Zaldarriaga 2014; Fidler et al. 2015, 2016; Hwang, Noh & Park 2016; Adamek et al. 2016a; Adamek et al. 2017; Fidler et al. 2017; Gong et al. 2017; Fidler et al. 2019; Magi & Yoo 2022). Generalizing these results to multiple fluids with different non-zero initial velocities is beyond the scope of the present study, as well as unnecessary to our goal of accuracy in the non-linear regime. Furthermore, following (Moran et al. 2023), we restrict our consideration to spatially flat cosmologies.

The chief difficulty in applying the continuity and Euler fluid equations to neutrinos is that, in the Eulerian fluid description, neutrinos are not fluids. At each spatial point, neutrinos have a velocity dispersion arising from their initial Fermi–Dirac distribution; the number of degrees of freedom required to describe neutrino perturbations is therefore technically infinite. Führer & Wong (2015) devised a Eulerian fluid-like perturbation theory for the neutrino bispectrum by absorbing the infinite degrees of freedom into temporally non-local couplings, while Garny & Taule (2021, 2022) combined a linear treatment of free streaming with non-linear corrections to the density and velocity monopole perturbations.

On the other hand, the approach of Dupuy & Bernardeau (2014, 2015a, b) and Chen, Upadhye & Wong (2021a, 2023b), which we describe in below, is to formulate a neutrino perturbation theory that is Lagrangian in momentum space. Let $\vec{\tau}$ be the lower-index three-momentum in the limit of an unperturbed universe, $P_i^{(0)}$, which is time-independent. The Fermi–Dirac distribution may be binned by the magnitude $\tau = |\vec{\tau}|$ of this momentum, allowing us to approximate the neutrino population using N_τ momenta, τ_α, for $\alpha \in [0, N_\tau)$.

For a given momentum vector $\vec{\tau}_\alpha$ in bin α, the set of neutrinos with initial momentum $\vec{\tau}_\alpha$ has no thermal velocity dispersion. Thus it behaves as a fluid obeying the continuity and Euler equations. Spatial isotropy implies that these equations depend upon the direction of $\vec{\tau}_\alpha$ only through its angle with respect to the Fourier vector $\vec{k}$, that is, through $\mu_\alpha := \vec{k}\cdot\vec{\tau}_\alpha/(k\tau_\alpha)$. Chen et al. (2021a) demonstrates that this μ_α-dependence can be expanded in N_μ Legendre polynomials $\mathcal{P}_\ell(\mu_\alpha)$, using an appropriate boundary term at $\ell = N_\mu - 1$, with an error of order N_μ^{-2}. Furthermore, for given τ_α, all neutrino fluids with initial momenta $\vec{\tau}$ such that $|\vec{\tau}| = \tau_\alpha$ obey the same fluid equations. We use the term 'flow' for this entire set of fluids.

The subhorizon non-relativistic linear theory of Chen et al. (2021a) began with dimensionless scalar perturbations to the density, $\delta_\alpha(\vec{x}) = (\rho_\alpha(x) - \bar{\rho}_\alpha)/\bar{\rho}_\alpha$, and the momentum divergence, $\theta_\alpha(\vec{x}) = -\vec{\nabla}\cdot\vec{P}/(m_\nu a\mathcal{H})$, with m_ν the neutrino mass and $\mathcal{H}$ the conformal Hubble expansion rate. Since linear theory allows us to choose an arbitrary normalization for the perturbation variables, we normalize them to the square roots of their corresponding power spectra,

$$\delta_\alpha^{\vec{k}} = \sum_\ell (-i)^\ell \mathcal{P}_\ell(\mu_\alpha)\delta_{\alpha\ell}^k \quad \text{and} \quad \theta_\alpha^{\vec{k}} = \sum_\ell (-i)^\ell \mathcal{P}_\ell(\mu_\alpha)\theta_{\alpha\ell}^k \tag{14}$$

$$P_{\alpha 00}^{\vec{k}} = \sum_\ell \mathcal{P}_\ell(\mu_\alpha)^2 \delta_{\alpha\ell}^k \delta_{\alpha\ell}^k =: \sum_\ell \mathcal{P}_\ell(\mu_\alpha)^2 P_{\alpha 00\ell}^k \tag{15}$$

$$P_{\alpha 11}^{\vec{k}} = \sum_\ell \mathcal{P}_\ell(\mu_\alpha)^2 \theta_{\alpha\ell}^k \theta_{\alpha\ell}^k =: \sum_\ell \mathcal{P}_\ell(\mu_\alpha)^2 P_{\alpha 11\ell}^k \tag{16}$$

where $\mathcal{P}_\ell(x)$ is the Legendre polynomial of order ℓ. Above we have employed two conventions we will use henceforth.

(i) Power spectrum indices b and c in $P_{\alpha bc}$ take the value 0 for δ and 1 for θ, so, for example, $P_{\alpha 00} = P_{\alpha\delta\delta}$.
(ii) Wave number superscripts denote a functional dependence, so that $P_{\alpha bc}^{\vec{k}} = P_{\alpha bc}(\vec{k})$ and $\delta_{\alpha\ell}^k = \delta_{\alpha\ell}(k)$.

The perfect correlation between the random variables corresponding to $\delta_\alpha(\vec{k})$ and $\theta_\alpha(\vec{k})$ in linear theory breaks down beyond the linear order, leading the `FlowsForTheMasses` perturbation theory of Chen et al. (2023b) to introduce the quantities

$$\chi_{\alpha\ell}^k := 1 - \frac{P_{\alpha 01\ell}^k}{\sqrt{P_{\alpha 00\ell}^k P_{\alpha 11\ell}^k}} \Rightarrow P_{\alpha 01\ell}^k = (1-\chi_{\alpha\ell}^k)\delta_{\alpha\ell}^k\theta_{\alpha\ell}^k. \tag{17}$$

Since $P_{\alpha 01} = P_{\alpha 10}$, equations (14–17) completely specify the power spectra $P_{\alpha bc}^{\vec{k}}$ in terms of the perturbation variables.

In terms of the bispectrum integrals $I_{\alpha,acd,bef,\ell}^k$ of Chen et al. (2023b), the evolution of $\delta_{\alpha\ell}$, $\theta_{\alpha\ell}$, and $\chi_{\alpha\ell}$ is given by

$$\begin{aligned}(\delta_{\alpha\ell}^k)' = {} & \frac{kv_\alpha}{\mathcal{H}}\left(\frac{\ell}{2\ell-1}\delta_{\alpha,\ell-1}^k - \frac{\ell+1}{2\ell+3}\delta_{\alpha,\ell+1}^k\right) \\ & + \theta_{\alpha\ell}^k + \frac{2}{\delta_{\alpha\ell}^k} I_{\alpha,001,001,\ell}^k\end{aligned} \tag{18}$$

$$\begin{aligned}(\theta_{\alpha\ell}^k)' = {} & -\left(1+\frac{\mathcal{H}'}{\mathcal{H}}\right)\theta_{\alpha\ell}^k + \frac{kv_\alpha}{\mathcal{H}} \\ & \times\left(\frac{\ell}{2\ell-1}\theta_{\alpha,\ell-1}^k - \frac{\ell+1}{2\ell+3}\theta_{\alpha,\ell+1}^k\right)\cdots \\ & - \delta_{\ell 0}^{(\mathrm{K})}\frac{k^2\Phi^k}{\mathcal{H}^2} + \frac{1}{\theta_{\alpha\ell}^k} I_{\alpha,111,111,\ell}^k\end{aligned} \tag{19}$$

$$\begin{aligned}(\chi_{\alpha\ell}^k)' = {} & \frac{2(1-\chi_{\alpha\ell}^k)}{(\delta_{\alpha\ell}^k)^2} I_{\alpha,001,001,\ell}^k + \frac{1-\chi_{\alpha\ell}^k}{(\theta_{\alpha\ell}^k)^2} I_{\alpha,111,111,\ell}^k\cdots \\ & - \frac{2}{\delta_{\alpha\ell}^k\theta_{\alpha\ell}^k} I_{\alpha,001,101,\ell}^k - \frac{1}{\delta_{\alpha\ell}^k\theta_{\alpha\ell}^k} I_{\alpha,111,011,\ell}^k.\end{aligned} \tag{20}$$

where $v_\alpha := \tau_\alpha/(m_\nu a)$ is the flow velocity, and primes denote derivatives with respect to $\eta := \log(a/a_{\mathrm{in}})$ for given initial scale factor a_{in}. The gravitational potential Φ is given by the Poisson equation

$$k^2\Phi^k = -\frac{3}{2}\mathcal{H}^2\left(\Omega_{\mathrm{cb}}(\eta)\delta_{\mathrm{cb}}^k + \sum_{\alpha=0}^{N_\tau-1}\Omega_\alpha(\eta)\delta_{\alpha 0}^k\right) \tag{21}$$

where δ_{cb} is the density contrast of the CDM and baryons, treated as a single fluid, labelled cb; $\Omega_{\mathrm{cb}}(\eta) = \Omega_{\mathrm{cb},0}\mathcal{H}_0^2/(a\mathcal{H}^2)$ is the time-dependent density fraction of this cb fluid; $\Omega_{\mathrm{cb},0}$ its density fraction

today; $\Omega_\alpha(\eta) = \Omega_{\alpha,0}\mathcal{H}_0^2/(a\mathcal{H}^2(1-v_\alpha^2)^{1/2})$ the time-dependent density fraction of neutrino flow α; and $\Omega_{\alpha,0}$ its density fraction today.

The bispectrum integrals $I^k_{\alpha,acd,bef,\ell}$ are defined and thoroughly studied in Chen et al. (2023b). For our purposes, we may define them by their evolution equations

$$(I^k_{\alpha,acd,bef,\ell})' = -\Xi^k_{\alpha bg\ell} I^k_{\alpha,acd,gef,\ell} - \tilde{\Xi}^k_{\alpha eg\ell} I^k_{\alpha,acd,bgf,\ell} \cdots - \tilde{\Xi}^k_{\alpha fg\ell} I^k_{\alpha,acd,beg,\ell} + 2A^k_{\alpha,acd,bef,\ell} \tag{22}$$

$$\Xi^k_{\alpha ab\ell} = \begin{bmatrix} 0 & -1 \\ \frac{k^2\Phi^k}{\mathcal{H}^2\delta^k_{\alpha 0}}\delta^{(\mathrm{K})}_{\ell 0} & 1+\frac{\mathcal{H}'}{\mathcal{H}} \end{bmatrix} \cdots - \delta^{(\mathrm{K})}_{a0}\delta^{(\mathrm{K})}_{b0}\frac{kv_\alpha}{\mathcal{H}}\left(\frac{\ell}{2\ell-1}\frac{\delta^k_{\alpha,\ell-1}}{\delta^k_{\alpha\ell}} - \frac{\ell+1}{2\ell+3}\frac{\delta^k_{\alpha,\ell+1}}{\delta^k_{\alpha\ell}}\right)\cdots - \delta^{(\mathrm{K})}_{a1}\delta^{(\mathrm{K})}_{b1}\frac{kv_\alpha}{\mathcal{H}}\left(\frac{\ell}{2\ell-1}\frac{\theta^k_{\alpha,\ell-1}}{\theta^k_{\alpha\ell}} - \frac{\ell+1}{2\ell+3}\frac{\theta^k_{\alpha,\ell+1}}{\theta^k_{\alpha\ell}}\right) \tag{23}$$

$$\tilde{\Xi}^k_{\alpha ab\ell} = \begin{bmatrix} 0 & -1 \\ 0 & 1+\frac{\mathcal{H}'}{\mathcal{H}} \end{bmatrix} \tag{24}$$

with initial conditions $I^k_{\alpha,acd,bef,\ell} = 2A^k_{\alpha,acd,bef,\ell}$ at $\eta = 0$. Here, the mode-coupling integrals are given by

$$A^{\vec{k}}_{\alpha,acd,bef} := \int \frac{d^3q}{(2\pi)^3}\frac{d^3p}{(2\pi)^3}(2\pi)^3\delta^{(\mathrm{D})}(\vec{k}-\vec{p}-\vec{q})\gamma^{\vec{k}\vec{q}\vec{p}}_{acd}\left[\gamma^{\vec{k}\vec{q}\vec{p}}_{bgh}P^{\vec{q}}_{\alpha ge}P^{\vec{p}}_{\alpha hf}\cdots + \gamma^{\vec{q},-\vec{p},\vec{k}}_{egh}P^{\vec{p}}_{\alpha gf}P^{\vec{k}}_{\alpha hb} + \gamma^{\vec{p},\vec{k},-\vec{q}}_{fgh}P^{\vec{k}}_{\alpha gb}P^{\vec{q}}_{\alpha he}\right]. \tag{25}$$

$$=: \sum_\ell \mathcal{P}_\ell(\mu_\alpha)^2 A^k_{\alpha,acd,bef,\ell} \tag{26}$$

$$\gamma^{\vec{k}\vec{q}\vec{p}}_{001} = \frac{(\vec{q}+\vec{p})\cdot\vec{p}}{2p^2}, \quad \gamma^{\vec{k}\vec{q}\vec{p}}_{010} = \gamma^{\vec{k}\vec{p}\vec{q}}_{001}, \quad \gamma^{\vec{k}\vec{q}\vec{p}}_{111} = \frac{(\vec{q}+\vec{p})^2\vec{q}\cdot\vec{p}}{2q^2p^2} \tag{27}$$

with all other γ_{abc} vanishing. In equations (22,25), we have assumed implicit summation over the indices g and h for compactness. Numerical computation of these mode-coupling integrals $A^k_{\alpha,acd,bef,\ell}$ is the main computational expense of `FlowsForTheMasses` perturbation theory. Chen et al. (2023b) shows that they may be reduced to FFTs and then computed using the methods of Hamilton (2000), McEwen et al. (2016), Schmittfull et al. (2016), Fang et al. (2017), and Upadhye (2019).

3 TRAINING DATA SET

3.1 Emulator design

We begin by constructing the training set $\mathscr{P}^*_{im}$ upon which the emulator is built. Emulator design is described broadly in Heitmann et al. (2009), and the particular design of the Mira-Titan IV (MT4) emulator upon which `Cosmic-Eν` is built is described in Heitmann et al. (2016), Lawrence et al. (2017), and Moran et al. (2023). The 111 design points in cosmological parameter space, $\vec{C}^*_m$, are chosen to strike a balance between broad parameter coverage and a high density of points, necessary for achieving high accuracy. Out of these, $N_\mathrm{M} = 101$ design points have non-zero neutrino masses, with physical density fractions $\Omega_{\nu,0}h^2$ ranging from 0.00017 to 0.01, corresponding to M_ν from 0.0158 to 0.931 eV.

Our strategy is to build upon the MT4 emulator design. Since massless neutrinos are already described well by linear theory, `Cosmic-Eν` uses only these $N_\mathrm{M} = 101$ massive-neutrino points. We build our emulator upon MT4 for two reasons. Firstly, the MT4 design has already been optimized and thoroughly tested for a balance between breadth and accuracy. Secondly, the MT4 CDM + baryon power spectrum is used as an input to the `FlowsForTheMasses` neutrino perturbation theory. Since MT4 is most accurate at its own design points, choosing this same design for `Cosmic-Eν` avoids the compounding of errors that would result from using the outputs of one emulator as the inputs for another. Another potential error is the back reaction of enhanced neutrino clustering on the CDM + baryon power. Chen et al. (2021a) quantified the LR back reaction to be 0.05 per cent for $\Omega_{\nu,0}h^2 = 0.01$, the largest value considered here, while Section 6.3 argues that non-linear clustering only increases this by a factor of $\sim$3. Thus, this back reaction is a negligible source of error for `Cosmic-Eν`.

Table 1. Allowed ranges of each cosmological parameter in `Cosmic-Eν`. As with the MT4 emulator, we have assumed three degenerate-mass neutrinos.

Parameter	Minimum	Maximum
$\Omega_{\mathrm{m},0}h^2$	0.12	0.155
$\Omega_{\mathrm{b},0}h^2$	0.0215	0.0235
$\Omega_{\nu,0}h^2$	0.00017	0.01
σ_8	0.7	0.9
h	0.55	0.85
n_s	0.85	1.05
w_0	−1.3	−0.7
$(-w_0-w_a)^{1/4}$	0.3	1.29

Table 1 lists the parameter ranges covered by the `Cosmic-Eν` emulator. The allowed range of $\Omega_{\nu,0}h^2$ values is determined by the set of 101 massive-neutrino MT4 models. Its lower bound of 0.00017 is over three times smaller than the lower bound imposed by laboratory oscillation experiments (Capozzi et al. 2018; de Salas et al. 2018; Esteban et al. 2020). In Section 4.2, we will quantify `Cosmic-Eν` errors near this low-$\Omega_{\nu,0}h^2$ boundary. For still smaller values, the non-relativistic-neutrino approximation made by `FlowsForTheMasses` becomes increasingly inaccurate, and we instead recommend the use of relativistic linear perturbation theories such as `CLASS` (Blas, Lesgourgues & Tram 2011; Lesgourgues 2011; Lesgourgues & Tram 2011) and `CAMB` (Lewis, Challinor & Lasenby 2000; Lewis & Bridle 2002). The upper bound on $\Omega_{\nu,0}h^2$, consistent with $M_\nu = 0.931$ eV, allows for a broad exploration of the parameter space, which may be useful for finding solutions to the Hubble and σ_8 tensions (McCarthy et al. 2018; Di Valentino & Melchiorri 2022; McCarthy et al. 2023). As with the MT4 emulator, we have assumed three degenerate-mass neutrinos.

Allowed ranges on the remaining seven parameters are taken directly from the MT4 emulator. We parametrize the dark energy equation of state as $w(z) = w_0 + w_a z/(1+z)$, following Chevallier & Polarski (2001); Linder (2003). The final range in Table 1 implies that if $w_0 = -1$, then $-1.77 \le w_a \le 0.99$. Allowing significant variation in the dark energy equation of state is important for cosmological constraints, as Upadhye (2019) showed that this variation weakens the neutrino mass bound by a factor of $\approx$3.

`Cosmic-Eν` is most useful at large neutrino masses, for which non-linear corrections to the neutrino power spectra are significant compared with its error. Though it accurately reproduces its input data even at small M_ν, to a level which we will quantify in subsequent sections, these training data assume three non-relativistic degenerate-mass neutrinos. Before proceeding to construct our training set, we briefly discuss the applicability of `Cosmic-Eν` to realistic low-M_ν models, for which we may use linear theory. For $\Omega_{\nu,0}h^2 = 0.00017$ ($M_\nu = 0.016$ eV), the smallest value in our training set, we find large-scale differences between our training data and the relativistic

CAMB power spectrum to be 24 per cent, 17 per cent, 8 per cent, and 0.7 per cent at redshifts 3, 2, 1, and 0, respectively. For $\Omega_{\nu,0}h^2 = 0.00067$ ($M_\nu = 0.063$ eV), near the lower bound of Salas et al. (2018), errors at these same redshifts are 6 per cent, 5 per cent, 3 per cent, and 0.1 per cent, respectively.

Power spectrum differences for normal (NO), inverted (IO), and degenerate (DO) neutrino mass orderings are more significant. The power spectra are nearly identical at large scales, but well below the free-streaming scales the linear power of each neutrino species scales as the fourth power of its mass (Ringwald & Wong 2004). Consider models with $M_\nu = 0.06$ eV, which for NO implies masses of ≈0.05 eV, 0.01 eV, and 0. The small-scale power spectrum will be dominated by the heavy neutrino, comprising 5/6 of the neutrino mass density. Hence, the small-scale neutrino power will exceed its DO counterpart by a factor of $(0.05\ \text{eV}/0.02\ \text{eV})^4(5/6)^2 \approx 27$. A similar calculation for the minimal-mass IO case, $M_\nu \approx 0.1$ eV, shows the ratio of small-scale IO to DO power spectra to be $(0.05\ \text{eV}/0.0333\ \text{eV})^4 \approx 5$. By $M_\nu = 0.2$ eV, small-scale differences with DO have fallen to 10 per cent for NO and 12 per cent for IO.

3.2 Stabilizing perturbation theory

Our next challenge is the high-k numerical instability of the FlowsForTheMasses perturbation theory. The mode-coupling integral $A^k_{\alpha,acd,bef,\ell}$ at low k and large ℓ rises sharply with k, increasing its dynamic range. Since FFTs spread errors across the entire range, small numerical errors near the peak of $A_{\alpha,acd,bef,\ell}$ lead to large fractional errors where $A_{\alpha,acd,bef,\ell}$ is small. These lead to instabilities, preventing integration of the equations of motion at high k. Additionally, the computational cost of the full set of mode-coupling integrals $A^k_{\alpha,acd,bef,\ell}$ of equations (25–26) was shown by Chen et al. (2023b) to scale as N_μ^6, further motivating a truncation in the range of ℓ.

In practice, FlowsForTheMasses integrates 128 values of the wave number, logarithmically distributed between $10^{-4}\,h\,\text{Mpc}^{-1}$ and $10\,h\,\text{Mpc}^{-1}$, and sets a stability threshold $k_{\rm st}$ to $10\,h\,\text{Mpc}^{-1}$ at the beginning of integration. The integration step size in η is chosen dynamically. Each time that high-k numerical instabilities drive this step size below 10^{-6}, the integrator discards the highest k value, effectively lowering $k_{\rm st}$ by 9.1 per cent, and then resumes integration for $k \le k_{\rm st}$. Since non-linear physics causes power to flow from low to high k, we may safely discard wave numbers above this stability threshold without affecting those below it. We will see below that no significant noise or discontinuities affect the power spectrum in the range $k \le k_{\rm st}$. As an added precaution, we will require stability up to a threshold $k_{\rm st}$ that is 20 per cent larger than our largest wave number of interest, $1\,h\,\text{Mpc}^{-1}$.

In order to stabilize FlowsForTheMasses for its fiducial model, with $\Omega_{\nu,0}h^2 = 0.005$, Chen et al. (2023b) truncated the power spectra $P^k_{abc\ell}$ used to compute $A^{\vec{k}}_{\alpha,acd,bef}$ to $\ell < N_{\mu,\rm NL}$, while allowing $N_{\mu,\rm NL} \le \ell < N_\mu$ elsewhere. They demonstrated that $N_{\mu,\rm NL}$ of 6, 7, and 8, respectively agreed with their N-body neutrino simulations to 14 per cent, 12 per cent, and 10 per cent for $k \le 1\,h\,\text{Mpc}^{-1}$. However, $N_{\mu,\rm NL}$ of 9 suffered from severe numerical instabilities preventing the equations of motion from being integrated to the present time. Henceforth, we fix $N_{\mu,\rm NL} = 6$ since its substantial reduction in computational expense relative to 7 and 8 only modestly decreases its accuracy.

Although the $\ell < N_{\mu,\rm NL}$ truncation of Chen et al. (2023b) sufficed to stabilize their $\Omega_{\nu,0}h^2 = 0.005$ model over the range $k \le 3\,h\,\text{Mpc}^{-1}$, we find that increasing the neutrino density fraction tends to exacerbate the numerical instabilities in FlowsForTheMasses. The MT4 emulator of Moran et al. (2023) allows $\Omega_{\nu,0}h^2$ to be twice as high as the fiducial model of Chen et al. (2023b). Requiring stability up to $k = 1.2\,h\,\text{Mpc}^{-1}$, so as to allow for a buffer around our desired range $k \le 1\,h\,\text{Mpc}^{-1}$, we find that 19 of the 101 design models are numerically unstable (that is, have $k_{\rm st} < 1.2\,h\,\text{Mpc}^{-1}$). The mean and minimum $\Omega_{\nu,0}h^2$ values for these models are 0.0083 and 0.0063, respectively, so this instability is a high-$\Omega_{\nu,0}h^2$ problem.

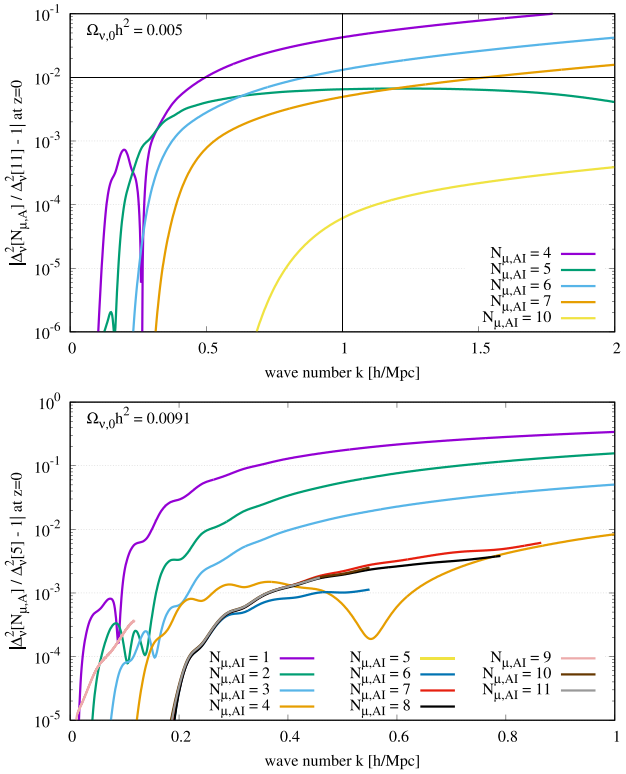

Figure 1. Accuracy of truncation $\ell < N_{\mu,AI}$ in mode-coupling and bispectrum integrals. (Top) $\Omega_{\nu,0}h^2 = 0.005$. All $N_{\mu,AI}$ are stable up to $k = 2\,h\,\text{Mpc}^{-1}$, so the power spectrum for each has been divided by that for $N_{\mu,AI} = 11$, the maximum value. (Bottom) MT4 design model with $\Omega_{\nu,0}h^2 = 0.0091$. Each power spectrum has been divided by that for $N_{\mu,AI} = 5$, the largest value for which the calculation is stable up to $k = 1.2\,h\,\text{Mpc}^{-1}$.

As $\Omega_{\nu,0}h^2$, hence the neutrino mass sum, increases, flow velocities v_α decrease, gradually decoupling the neutrino density and velocity monopoles from higher Legendre moments ℓ. Thus we employ a more aggressive high-ℓ truncation in order to stabilize these high-$\Omega_{\nu,0}h^2$ models. In addition to truncating the power spectrum used to compute mode-coupling integrals, we truncate the ℓ expansions of the mode-coupling integrals $A^k_{\alpha,acd,bef,\ell}$ and bispectrum integrals $I^k_{\alpha,acd,bef,\ell}$ themselves, $\ell < N_{\mu,AI}$. The power truncation $\ell < N_{\mu,\rm NL}$ itself implies an $N_{\mu,AI}$ of $2N_{\mu,\rm NL} - 1$, or 11 for $N_{\mu,\rm NL} = 6$, so we allow $N_{\mu,AI}$ to be reduced below this number.

As an upper bound on the error associated with this truncation, consider the $\Omega_{\nu,0}h^2 = 0.005$ fiducial model of Chen et al. (2023b). Fig. 1 (Top) compares neutrino power spectra for several $N_{\mu,AI}$ values to the maximum value $2N_{\mu,\rm NL} - 1 = 11$. In the range $k \le 1\,h\,\text{Mpc}^{-1}$, $N_{\mu,AI} = 4$ is accurate to 4 per cent, while the higher $N_{\mu,AI}$ considered are accurate to 1.3 per cent or better.

Fig. 1 (Bottom) makes a similar comparison for one of the MT4 design points with $\Omega_{\nu,0}h^2 = 0.0091$. Since power spectra with $N_{\mu,AI} > 5$ cannot be stably integrated to $z = 0$ all the way to $k =$

1.2 $h\,\mathrm{Mpc}^{-1}$, we use the $N_{\mu,AI} = 5$ power spectrum for comparison. Encouragingly, the $N_{\mu,AI} = 4$ power spectrum agrees with this to better than 1 per cent at all $k \leq 1\ h\,\mathrm{Mpc}^{-1}$. Also, even those higher-$N_{\mu,AI}$ power spectra with $k_{\rm st} < 1\ h\,\mathrm{Mpc}^{-1}$ agree with the $N_{\mu,AI} = 5$ power to better than 1 per cent across their entire stable ranges $k \leq k_{\rm st}$. Evidently from the figure, our stabilization procedure discards higher k early enough to prevent them from contaminating $k \leq k_{\rm st}$ at even the per cent level.

Out of the nineteen MT4 design points requiring a mode-coupling truncation $N_{\mu,AI} < 2N_{\mu,\rm NL} - 1$, 15 reach $k = 1.2\ h\,\mathrm{Mpc}^{-1}$ with $N_{\mu,AI} = 5$, and the remaining four reach that wave number with $N_{\mu,AI} = 4$. These four $N_{\mu,AI} = 4$ models all have $\Omega_{\nu,0}h^2 > 0.008$, with a mean $\Omega_{\nu,0}h^2$ value of 0.0092. All have stability thresholds $k_{\rm st} > 0.7\ h\,\mathrm{Mpc}^{-1}$ when run with $N_{\mu,AI} = 5$. By comparison with Fig. 1 (Bottom), we may estimate their truncation error as ~ 1 per cent. Another estimate of the truncation error, for all nineteen stabilized models, is the difference between the $N_{\mu,AI} = 4$ and $N_{\mu,AI} = 5$ power spectra, either to $k = 1\ h\,\mathrm{Mpc}^{-1}$ or to $k_{\rm st}$ if this is less than one. By this measure, we find that one of these nineteen has a 2.5 per cent truncation error, while all others have ≤ 2 per cent and nine of them have ≤ 1 per cent. Since these nineteen represent nearly half of the 40 design models with $\Omega_{\nu,0}h^2 \geq 0.0063$, we estimate that truncation leads to a ≈ 1 per cent power underestimate in that range.

3.3 Reduced power spectrum

The power spectrum of massive neutrinos declines sharply below the free-streaming scale, $k \gg k_{\rm FS}$, giving $\Delta^2_\nu(k)$ a large dynamic range, which in turn makes it difficult to emulate. Our strategy is to divide the neutrino power spectra, the total power as well as the single-decile power spectra, by quickly calculable linear approximations. Further reduction in the dynamic range is achieved by taking the natural logarithm of this power spectrum ratio, for each of the MT4 models, resulting in the training set $\mathscr{P}^*_{im}$.

As a starting point for a fast linear approximation to the neutrino power spectra, we take the matter power spectra of Eisenstein and Hu (Eisenstein & Hu 1997, 1998; Hu & Eisenstein 1998). In the clustering limit, $k \ll k_{\rm FS}$, neutrinos and cold matter cluster very similarly. In the free-streaming limit, $k \gg k_{\rm FS}$, the ratio of the neutrino and total matter density contrasts scales as $k^2_{\rm FS}/k^2$, as shown by Ringwald & Wong (2004) and Wong (2008). Those studies developed and tested an interpolation function, $\delta_\nu/\delta_{\rm m} \approx (1 + k/k_{\rm FS})^{-2}$, that applies to linearly clustering neutrinos, even when the CDM and baryons cluster non-linearly.

Chen et al. (2021a, b) generalized the free-streaming scale of Ringwald & Wong (2004) to individual neutrino flows through the replacement of their neutrino sound speed c_ν by the flow velocity v,

$$k_{\rm FS}(a,v)^2 = \frac{3\Omega_{\rm m}(a)\mathcal{H}(a)^2}{2v(a)^2} \tag{28}$$

where, in the non-relativistic approximation, $v(a) = \tau/(am_\nu)$, implying that $k_{\rm FS}(a,v) \propto \sqrt{a}$. Thus our Eisenstein-Hu-like approximations to the total and decile neutrino power spectra are, respectively,

$$\Delta^2_{{\rm EH},\nu}(k,a) = \Delta^2_{\rm EH}(k,a)\left(1 + \frac{k}{k_{\rm FS}(a,c_\nu)}\right)^{-4} \text{ with } c^2_\nu = \frac{3\zeta(3)T^2_{\nu,0}}{2\log(2)m^2_\nu a^2} \tag{29}$$

$$\Delta^2_{{\rm EH},L}(k,a) = \Delta^2_{\rm EH}(k,a)\left(1 + \frac{k}{k_{\rm FS}(a,v_L)}\right)^{-4} \text{ with } v_L = \langle v_\alpha\rangle_L \tag{30}$$

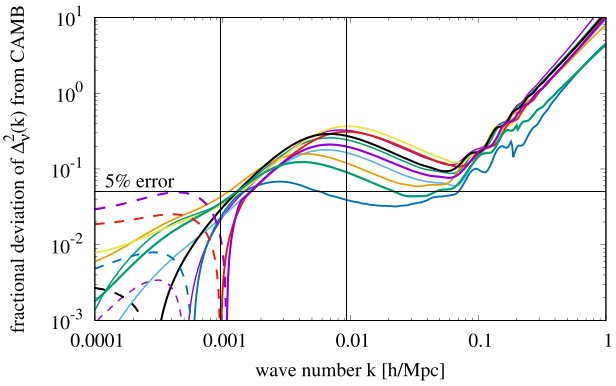

Figure 2. Fractional difference $\Delta^2_{\nu,\rm FlowsForTheMasses}/\Delta^2_{\nu,\rm CAMB} - 1$ between `FlowsForTheMasses` and `CAMB` power spectra for ten cosmologies spanning the parameter space, listed in Table 2. Solid (dashed) lines show positive (negative) differences.

where $<v_\alpha>_L$ in the final line denotes an average over all flows α making up decile L, and $\zeta(x)$ is the Riemann zeta function.

Next, we consider the range of scales over which we emulate the neutrino power spectra. We set the upper limit of our range of wave numbers to $1\ h\,\mathrm{Mpc}^{-1}$, since beyond that we expect non-perturbative effects, such as the capture of neutrinos by haloes, to dominate the power spectrum. Meanwhile, our minimum wave number should be small enough that `FlowsForTheMasses` agrees closely with `CAMB` below it, but no smaller, so as to avoid a tradeoff between high-k and low-k accuracy. Fig. 2 demonstrates that this minimum wave number should not be much larger than $\sim$0.001 $h\,\mathrm{Mpc}^{-1}$. Below that value, all power spectra shown agree with `CAMB` to ≤ 5 per cent, with per cent-level agreement for nearly all models below $k = 0.0002\ h\,\mathrm{Mpc}^{-1}$. Above $k = 0.001\ h\,\mathrm{Mpc}^{-1}$, fractional differences rise rapidly, reaching ~ 20 per cent by $k = 0.01\ h\,\mathrm{Mpc}^{-1}$. Thus we choose to emulate $\log(\Delta^2_\nu/\Delta^2_{{\rm EH},\nu})$ and $\log(\Delta^2_L/\Delta^2_{{\rm EH},L})$ over the range $0.001\ h\,\mathrm{Mpc}^{-1} \leq k \leq 1\ h\,\mathrm{Mpc}^{-1}$.

3.4 Non-linear enhancement: A first look

The chief goal of this study is to quantify the non-linear clustering of neutrinos across a wide range of parameters. One tool we will use for this is the non-linear enhancement ratio, that is, the ratio of the `FlowsForTheMasses` power spectrum to that from the Multi-Fluid Linear Response code `MuFLR` of Chen et al. (2021a):[5]

$$R_\nu(k,z) = \Delta^2_{\nu,FlowsForTheMasses}(k,z)/\Delta^2_{\nu,MuFLR}(k,z). \tag{31}$$

`MuFLR` allows the CDM + baryon fluid to cluster non-linearly while limiting neutrino clustering to the linear terms. That is, the mode-coupling integrals $A_{\alpha,acd,bef}$, hence also the bispectrum integrals $I_{\alpha,acd,bef}$ and the non-linear correlations $\xi_{\alpha\ell}$, are set to zero, leaving only equations (18, 19, and 21) to be solved for the neutrinos.

The enhancement ratio is instructive because it allows us to isolate directly the effects of non-linearity in the neutrino sector. High-quality simulations using neutrino LR include Ali-Haimoud & Bird (2012), McCarthy et al. (2017), and Liu et al. (2018). $R_\nu(k, z)$ quantifies the amount by which these underestimate neutrino clustering. Furthermore, Chen et al. (2023a, b) showed that `FlowsForTheMasses` itself becomes inaccurate in the regime $R_\nu -$

[5] `MuFLR` is publicly available at github.com/upadhye/MuFLR.

$1 \gtrsim 1$, indicating the scales on which a particle neutrino simulation is necessary for accurate predictions of neutrino clustering. Section 6 will use emulation to isolate the impact upon R_ν of individual cosmological parameters.

Fig. 3 provides a glimpse of the parameter-dependence of $R_\nu(k, 0)$. From Fig. 3 (a), we immediately see that $\Omega_{\nu,0}h^2 \propto M_\nu$ has a significant effect upon R_ν. Furthermore, non-linear corrections to the lightest neutrinos reduce R_ν slightly below unity. This is not surprising, as non-linear corrections in the weakly non-linear regime are known to suppress even the CDM clustering (Bernardeau et al. 2002). Evidently from Fig. 3, the clustering amplitude σ_8 also has a discernible effect upon R_ν, while the impact of the total matter and baryon fractions are less obvious. We will revisit the impact of different cosmological parameters upon $R_\nu(k, 0)$ in Section 6.1.

4 EMULATION OF THE NON-LINEAR NEUTRINO POWER SPECTRUM

4.1 Emulation using SEPIA

For each neutrino momentum decile L, we use the SEPIA code of Gattiker et al. (2020) to determine the values of the PC basis functions $\phi_j^{(L)}(k_i, z_i)$ as well as sample from the posteriors of the hyperparameters $\beta_{j\ell}^{(L)}$, $\lambda_{\mathrm{U},j}^{(L)}$, and $\lambda_{\mathrm{W},j}^{(L)}$. We use $N_{\mathrm{PC}} = 50$ PCs for each L, comparable to the 45 used for the MT4 emulator. Hyperparameters are optimized in SEPIA using $N_{\mathrm{MCMC}} = 50\,000$ Markov chain Monte Carlo steps. Since the error on the total neutrino power spectrum for N_{MCMC} of 10 000, 20 000, and 50 000 is, respectively, 3.92 per cent, 3.86 per cent, and 3.48 per cent, measured against ten randomly chosen models outside of the MT4 sample, we regard $N_{\mathrm{MCMC}} = 50\,000$ to have converged.

SEPIA is designed to draw weights from the probability distribution of equation (10). We instead prefer a deterministic emulator such as that described in Heitmann et al. (2009), which uses the mean weights $\bar{W}_j^{(L)}(\vec{C}\,)$ of equation (11) as the emulator prediction. These depend on the hyperparameters' posterior means $\hat{\beta}_{j\ell}^{(L)}$, $\hat{\lambda}_{\mathrm{U},j}^{(L)}$, and $\hat{\lambda}_{\mathrm{W},j}^{(L)}$. Appendix A shows how to obtain these quantities from SEPIA.

Our goal is now in sight: the mean weight $\bar{W}_j^{(L)}(\vec{C}\,)$ of equation (11) for each decile L and PC j, for a given cosmological parameter vector $\vec{C}$. The final ingredient needed to emulate $\bar{W}_j^{(L)}(\vec{C}\,)$ is the so-called Kriging basis $[\hat{\mathscr{R}}_j^{*(L)}]^{-1}\vec{w}_j^{\,*(L)}$. With fixed L and j, the quantity $\hat{\mathscr{R}}_j^{*(L)}$ is an $N_{\mathrm{M}} \times N_{\mathrm{M}}$ matrix, $[\hat{\mathscr{R}}_j^{*(L)}]^{-1}$ is its matrix inverse, and $\vec{w}_j^{\,*(L)}$ and $[\hat{\mathscr{R}}_j^{*(L)}]^{-1}\vec{w}_j^{\,*(L)}$ are vectors of length N_{M}. Since $\hat{\mathscr{R}}_j^{*(L)}$ and $\vec{w}_j^{\,*(L)}$ are both known, we find the Kriging basis by solving the linear system:

$$\hat{\mathscr{R}}_j^{*(L)} X_j^{(L)} = \vec{w}_j^{\,*(L)} \tag{32}$$

for $X_j^{(L)}$, which equals $[\hat{\mathscr{R}}_j^{*(L)}]^{-1}\vec{w}_j^{\,*(L)}$.

Since the Kriging basis is independent of $\vec{C}$, we compute it once and save the result. Now, given $\vec{C}$, we may readily compute $r_{jm}^*(\vec{C}\,)$ of equation (13) for each $m = 0, \ldots, N_{\mathrm{M}} - 1$. The dot product of the vector $\vec{r}_j^{\,*}(\vec{C}\,)$ with the Kriging basis gives $\bar{W}_j^{(L)}(\vec{C}\,)$, as in equation (11). The emulated reduced neutrino power $\mathscr{P}^{(L)}(k_i, z_i, \vec{C}) = \log(\Delta_L^2/\Delta_{\mathrm{EH},L}^2)$ for decile L is given by:

$$\mathscr{P}^{(L)}(k_i, z_i, \vec{C}) = \mu_i^{*(L)} + \sigma^{*(L)} \sum_{j=0}^{N_{\mathrm{PC}}-1} \bar{W}_j^{(L)}(\vec{C}\,)\phi_j^{(L)}(k_i, z_i), \tag{33}$$

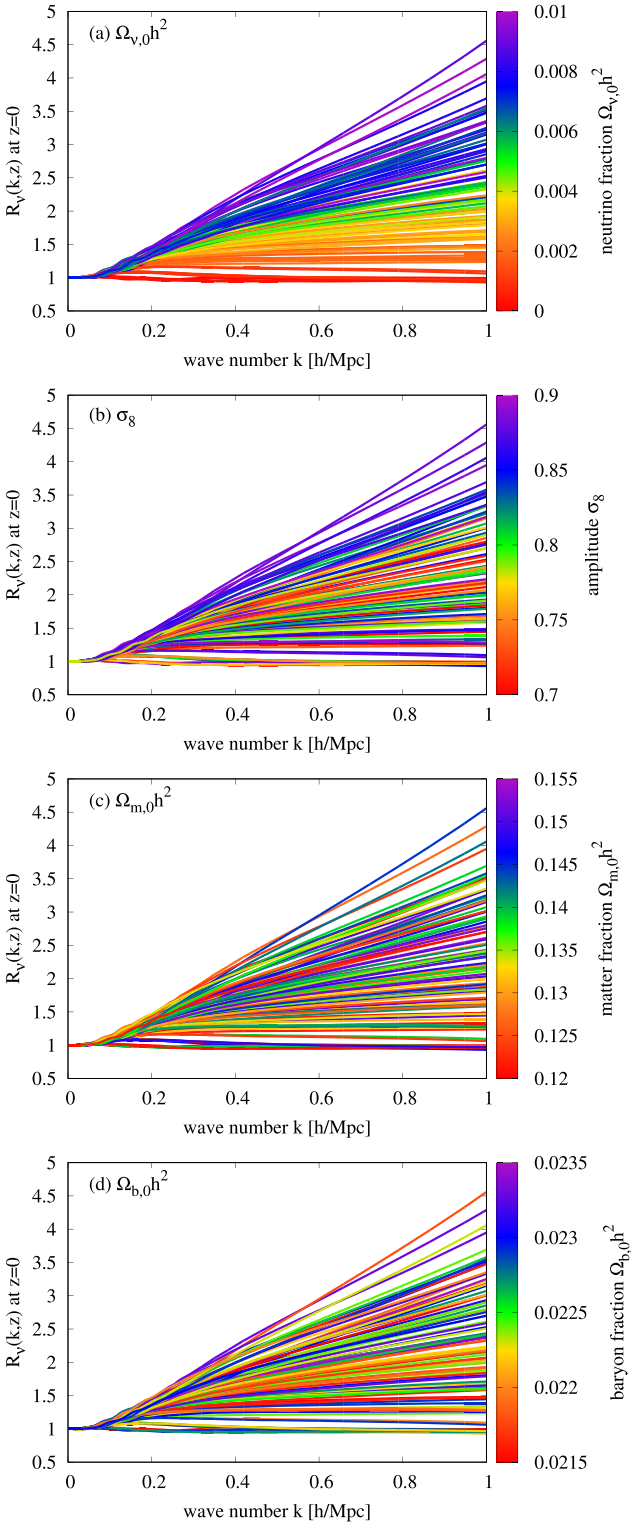

Figure 3. Neutrino non-linear enhancement ratio of equation (31) for each of the $N_{\mathrm{M}} = 101$ MT4 $M_\nu > 0$ design points at redshift $z = 0$. Each curve is color-coded by (a) $\Omega_{\nu,0}h^2$, (b) σ_8, (c) $\Omega_{\mathrm{m},0}h^2$, or (d) $\Omega_{\mathrm{b},0}h^2$.

Table 2. Cosmological parameters for the ten out-of-sample test models. Each is a spatially flat νwCDM model with $w(a) = w_0 + (1 - a)w_a$.

Model	$\Omega_{\mathrm{m},0}h^2$	$\Omega_{\mathrm{b},0}h^2$	$\Omega_{\nu,0}h^2$	σ_8	h	n_s	w_0	w_a
E001	0.1433	0.02228	0.008078	0.8389	0.7822	0.9667	−0.8000	−0.0111
E002	0.1333	0.02170	0.005311	0.8233	0.7444	0.9778	−1.1560	−1.1220
E003	0.1450	0.02184	0.003467	0.8078	0.6689	0.9000	−0.9333	−0.5667
E004	0.1367	0.02271	0.002544	0.8544	0.8200	0.9444	−0.8889	−1.4000
E005	0.1400	0.02257	0.009000	0.7300	0.7067	0.9889	−0.9778	−0.8444
E006	0.1350	0.02213	0.000700	0.8700	0.7633	0.9111	−1.0220	0.5444
E007	0.1383	0.02199	0.007156	0.7456	0.6500	0.9556	−1.1110	1.1000
E008	0.1300	0.02286	0.006233	0.7922	0.8011	1.0000	−1.0670	0.2667
E009	0.1417	0.02300	0.004389	0.7767	0.7256	0.9222	−0.8444	0.8222
E010	0.1317	0.02242	0.001622	0.7611	0.6878	0.9333	−1.2000	−0.2889

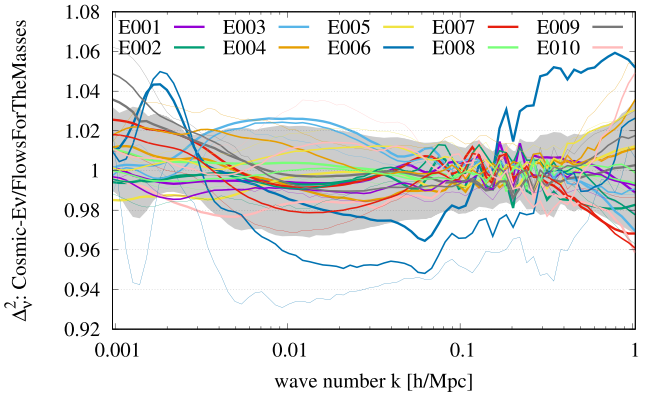

Figure 4. `Cosmic-Eν` versus `FlowsForTheMasses` for the total neutrino power spectra $\Delta^2_\nu(k, z)$ of the out-of-sample models of Table 2. For each model, lines show the redshifts 3, 1, and 0, in order of increasing thickness. The grey shaded region is the mean plus-or-minus one standard deviation, maximized over all emulated redshifts. Across all k and z, this error is less than 3.5 per cent.

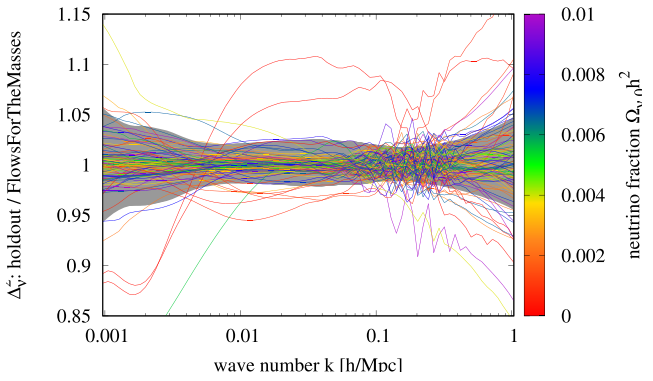

Figure 5. Leave-one-out holdout tests of `Cosmic-Eν`. Lines show the holdout-to-`FlowsForTheMasses` power spectrum ratios at $z = 0$, with colors corresponding to $\Omega_{\nu,0}h^2$. The dark grey shaded region is the mean plus-or-minus one standard deviation, maximized over all emulated redshifts; the maximum 1σ error across all k and z is 5.7 per cent.

and the total neutrino power spectrum $\Delta^2_\nu(k_i, z_i)$ is formed by averaging over the individual decile powers $\Delta^2_L(k_i, z_i)$:

$$\Delta^2_\nu(k_i, z_i) = \left[\frac{1}{10}\sum_{L=0}^{9}\sqrt{\Delta^2_L(k_i, z_i)}\right]^2 . \tag{34}$$

As an alternative, we could have emulated $\Delta^2_\nu(k_i, z_i)$ separately. However, we find averaging over decile powers to be more accurate. Further, averaging ensures consistency between $\Delta^2_\nu(k_i, z_i)$ and the individual $\Delta^2_L(k_i, z_i)$, which should obey equation (34).

4.2 Tests of `Cosmic-Eν`

Now that our `Cosmic-Eν` emulator is complete, we quantify its accuracy. We first consider the total neutrino power spectrum $\Delta^2_\nu(k, z)$, the main goal of this study. Power spectra of individual momentum deciles, $\Delta^2_L(k, z)$, are covered at the end of this section.

We begin with the most accurate error measurement, which compares emulator predictions to `FlowsForTheMasses` computations for a set of ten test models outside of the MT4 design set. Table 2 lists the cosmological parameters of these out-of-sample models, E001 through E010. They cover a large range of $\Omega_{\nu,0}h^2$ from 0.0007, for E006, to 0.009, for E005, and allow for a substantial variation in the dark energy equation of state.

Fig. 4 compares the total neutrino power spectra of `Cosmic-Eν` and `FlowsForTheMasses` for the 10 out-of-sample models of Table 2. For each of the 27 redshifts emulated and at each k, the mean and standard deviation of the `Cosmic-Eν`-to-`FlowsForTheMasses` ratio are computed. The 1σ error interval is the region within one standard deviation of the mean ratio at each k and z emulated, i.e., the standard deviation across cosmologies for a single redshift and wave number. At each k, the grey shaded region shows the 1σ error interval maximized over redshift. The 1σ error, the difference between this region and unity, is everywhere less than 3.5 per cent. The maximum error across all 10 models, over all k and z, is 6.9 per cent. The largest error is for model E006, which has a small neutrino density.

Another emulator error estimate, the leave-one-out holdout test, is performed directly through `SEPIA`. For each model m of the N_M models in the training set, `SEPIA` builds an emulator with all $N_\mathrm{M} - 1$ models excluding m, then compares the resulting emulator prediction $\mathscr{P}_{im}$ at $\vec{C}^*_m$ with the excluded training data $\mathscr{P}^*_{im}$ for that model. Since leaving out model m creates a gap in the training data at $\vec{C}^*_m$, holdout tests tend to overestimate the emulator error, particularly in the context of a space-filling design.

Fig. 5 shows the results of leave-one-out holdout tests of `Cosmic-Eν`. The maximum 1σ error is now 5.7 per cent, about 60 per cent higher than the out-of-sample test. Nevertheless, at all but the largest scales, $k < 0.002\,h\,\mathrm{Mpc}^{-1}$, and the smallest scales, $k > 0.8\,h\,\mathrm{Mpc}^{-1}$, the error is 4 per cent or less. Meanwhile, the largest holdout test error across all k and z, and all 101 models, is 25.4 per cent.

The two models with the largest holdout test errors are further studied in Fig. 6. Out of the models in the MT4 training set, these two have the smallest $\Omega_{\nu,0}h^2$; they are the only models allowing $\Omega_{\nu,0}h^2$ to fall below half of its lower bound ≈0.00064 from neutrino oscillation

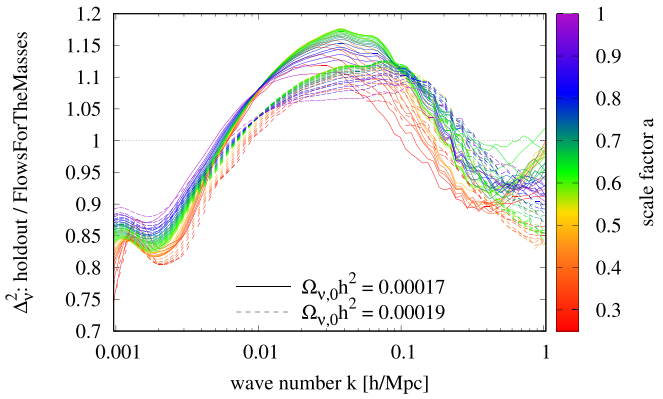

Figure 6. Holdout tests for the two lowest-neutrino-mass models, which also have the largest errors in Fig. 5, colour-coded by the scale factor a. Solid (dashed) curves correspond to $\Omega_{\nu,0}h^2 = 0.00017$ ($\Omega_{\nu,0}h^2 = 0.00019$).

experiments. Evidently, the 25 per cent errors noted above are due to sharp falls in the Cosmic-Eν predictions at the earliest times and largest scales. Errors are under 20 per cent for all $k \geq 0.0012\,h\,\mathrm{Mpc}^{-1}$ and all z. Since the holdout test errors of Fig. 5 overestimate the more accurate out-of-sample errors of Fig. 4 by about 60 per cent, we estimate Cosmic-Eν errors of 10 per cent − 12 per cent at the lowest neutrino masses, in the range $k \geq 0.0012\,h\,\mathrm{Mpc}^{-1}$. Meanwhile, the model with the third-largest holdout error in Fig. 5, which has a 22 per cent low-k holdout error giving way to ≤ 20 per cent holdout errors at $k \geq 0.002\,h\,\mathrm{Mpc}^{-1}$ and ≤ 5 per cent errors at $k \gtrsim 0.01\,h\,\mathrm{Mpc}^{-1}$, has $M_\nu \approx 0.5$ eV but near-maximal values of $\Omega_{\mathrm{b},0}h^2$ and $w_0 + w_a$.

Finally, we compare individual neutrino momentum deciles between Cosmic-Eν and FlowsForTheMasses. In multifluid perturbation theories such as ours of Section 2.2, the neutrino flow speed v_α behaves similarly to a sound speed, and the resulting power spectrum exhibits oscillatory behaviour on sufficiently small scales. Averaging over a large number of flows eliminates these oscillations. However, each decile averages over only 5 flows, rather than 50 for the total neutrino power, making the small-scale decile powers noisier and more difficult to emulate. This is especially true for the higher-momentum deciles for which individual-flow oscillations are more prominent. Thus, we expect less accuracy in $\Delta^2_L(k, z)$, particularly for high L, than in the $\Delta^2_\nu(k, z)$ shown above.

Fig. 7 combines out-of-sample tests and leave-one-out holdout tests for each of the ten momentum deciles. The 1σ errors shown are in line with our expectations above. For deciles $L = 0$ and 1, errors from out-of-sample tests are < 7 per cent, or about twice the maximum error on $\Delta^2_\nu(k, z)$. For higher L, errors remain at this level for $k \leq 0.1\,h\,\mathrm{Mpc}^{-1}$ but grow substantially at large k, rising to nearly four times the $\Delta^2_\nu(k, z)$ emulator error at $k = 1\,h\,\mathrm{Mpc}^{-1}$.

5 COMPARISON TO N-BODY SIMULATIONS

5.1 Variation of the ν implementation

The previous section quantified the precision with which Cosmic-Eν reproduced its underlying FlowsForTheMasses perturbation theory. Next, we consider its accuracy relative to numerical simulations of the massive neutrino power spectrum. Adamek et al. (2023) carried out an extensive comparison of neutrino simulation methods for a spatially flat νΛCDM model with $M_\nu = 0.15$ eV ($\Omega_{\nu,0}h^2 = 0.00161$), $\Omega_{\mathrm{m},0}h^2 = 0.1432$, $\Omega_{\mathrm{b},0}h^2 = 0.022$, $A_\mathrm{s} = 2.215 \times 10^{-9}$ ($\sigma_8 = 0.815$), $h = 0.67$, and $n_\mathrm{s} = 0.9619$. Though all of the neutrino power spectra agree at about the per cent level in the linear regime, $k \lesssim 0.1\,h\,\mathrm{Mpc}^{-1}$, the different methods are discrepant at the 30 per cent − 40 per cent level by $k = 1\,h\,\mathrm{Mpc}^{-1}$.

Even within a given simulation method, choices of initial conditions and simulation parameters have a substantial impact on the massive neutrino power spectrum. Sullivan et al. (2023) implemented the 'tiling' method of Banerjee et al. (2018). They find that the number of allowed neutrino momentum directions is the most important parameter for determining convergence, and they judge their highest-direction-number run to have converged at the ∼ 10 per cent level up to half of the neutrino Nyquist frequency, or $\sim 1\,h\,\mathrm{Mpc}^{-1}$ in their standard run. Combining this with the discrepancies among the different simulation methods, we may regard simulation uncertainty in the neutrino power at $k \lesssim 1\,h\,\mathrm{Mpc}^{-1}$ to be in the 30 per cent − 50 per cent range.

Fig. 8 compares Cosmic-Eν and MuFLR LR to the results of Adamek et al. (2023). Each power spectrum is compared to that computed using the SWIFT simulation of Schaller et al. (2023), with their ratio smoothed using a centered 10-point moving average. Compared with SWIFT, the CO$\mathcal{N}$CEPT code of Dakin et al. (2019) predicts 22 per cent more power at $k = 1\,h\,\mathrm{Mpc}^{-1}$, while gevolution (Adamek et al. 2016a) predicts ≈ 15 per cent − 20 per cent less.

Cosmic-Eν agrees closely with all of the simulated Δ^2_ν at low k. Its predicted power spectrum falls below that of SWIFT by ≤ 3 per cent up to $k = 0.17\,h\,\mathrm{Mpc}^{-1}$, or three times the free-streaming wave number for this neutrino mass. This power deficit grows rapidly, rising to 19 per cent at $k = 0.4\,h\,\mathrm{Mpc}^{-1}$ and 49 per cent at $k = 1\,h\,\mathrm{Mpc}^{-1}$, somewhat greater than, but comparable to, the scatter between different N-body simulation methods. Thus, FlowsForTheMasses and the Cosmic-Eν emulator appear to provide accurate computations of the non-linear neutrino power, given the current level of simulation uncertainty.

Most of the scatter among the N-body methods at wave numbers $k \approx 0.5\,h\,\mathrm{Mpc}^{-1}$ is due to a systematic difference between particle-based methods, such as SWIFT, and CO$\mathcal{N}$CEPT, which integrates the massive neutrino fluid equations on a grid. *A priori* we have no reason to consider one of these more accurate. However, if we exclude CO$\mathcal{N}$CEPT as an outlier among the N-body simulations, then the range spanned by the remaining simulations drops significantly, and the Cosmic-Eν power deficit exceeds this range by a factor of a few for $k \gtrsim 0.5\,h\,\mathrm{Mpc}^{-1}$. Thus, it is not clear whether the Cosmic-Eν small-scale power deficit is due to systematic uncertainties among the different non-linear methods or to a genuine non-perturbative effect such as the capture of neutrinos by CDM + baryon haloes.

Also shown in Fig. 8 is the MuFLR LR power spectrum. As with Cosmic-Eν, its power is less than that of the SWIFT simulation used as a reference. Its power deficits relative to SWIFT are significantly larger than those of Cosmic-Eν, 6.1 times larger at $k = 0.1\,h\,\mathrm{Mpc}^{-1}$ and 1.8 times larger at $k = 0.4\,h\,\mathrm{Mpc}^{-1}$. Thus Cosmic-Eν represents a significant accuracy improvement over LR approximations.

5.2 Variation of the ν mass

Next, we consider variations in M_ν. Adamek et al. (2023) varied M_ν from 0.15 to 0.6 eV, in their flat νΛCDM model with $\Omega_{\mathrm{m},0}h^2$, $\Omega_{\mathrm{b},0}h^2$, h, A_s, and n_s fixed to the values of Section 5.1. Fixing the initial power spectrum amplitude $A_\mathrm{s} = 2.215 \times 10^{-9}$ implies $\sigma_8 = 0.815$ for $M_\nu = 0.15$ eV, $\sigma_8 = 0.776$ for $M_\nu = 0.3$ eV, and $\sigma_8 = 0.731$

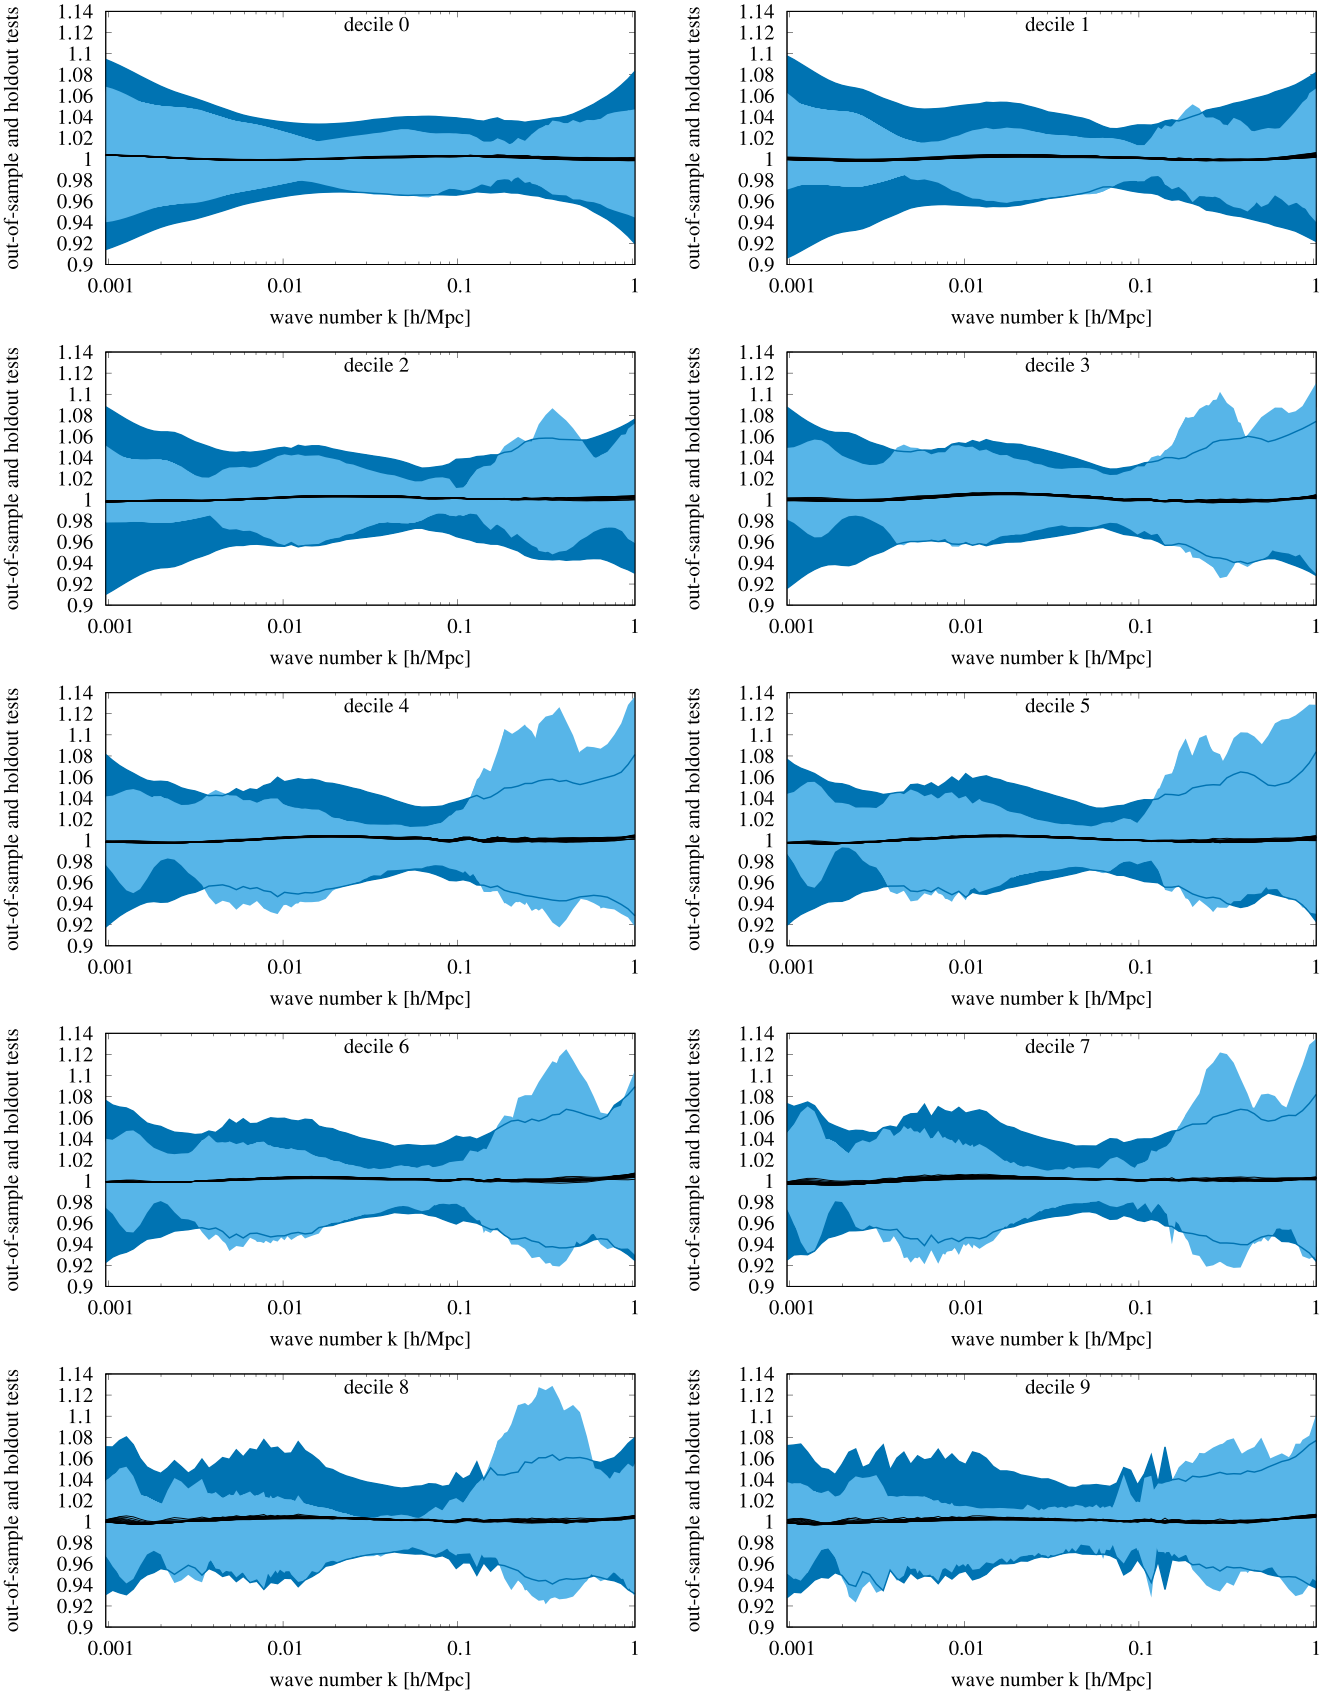

Figure 7. Out-of-sample (light shaded regions) and holdout (dark shaded regions) tests for the individual neutrino decile power spectra $\Delta^2_L(k, z)$. Shaded regions show the mean plus-or-minus one standard deviation, maximized over all emulated redshifts.

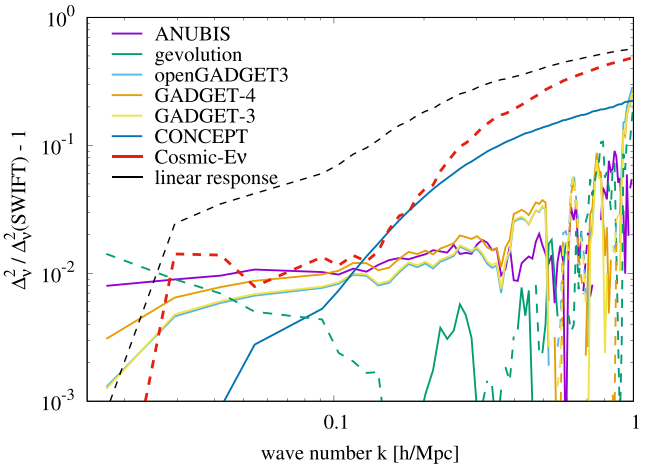

Figure 8. Comparison of `Cosmic-Eν` and `MuFLR` (LR) to N-body neutrino power spectra, for $M_\nu = 0.15$ eV, at $z = 0$, computed using a variety of methods by Adamek et al. (2023). Solid (dashed) lines represent power spectra that are greater (less) than that of `SWIFT`. Ratios have been smoothed using a centred 10-point moving average.

for $M_\nu = 0.6$ eV. Their simulations for $M_\nu = 0.3$ eV and 0.6 eV use 512^3 neutrino particles in a $(512\ \mathrm{Mpc})^3$ box.

Fig. 9 compares `Cosmic-Eν` and `MuFLR` to the Adamek et al. (2023) simulations with varying M_ν. For $k \sim 0.1\,h\,\mathrm{Mpc}^{-1}$, `Cosmic-Eν` agrees with the simulations at the 2 per cent − 4 per cent level, while `MuFLR` underpredicts power by 7 per cent for the lower masses. At larger k, both errors quickly increase. `Cosmic-Eν` errors at $k = 0.4\,h\,\mathrm{Mpc}^{-1}$ are 17 per cent and 14 per cent for $M_\nu = 0.3$ and 0.6 eV, respectively, while those for `MuFLR` are respectively 46 per cent and 49 per cent. Above $k = 0.4\,h\,\mathrm{Mpc}^{-1}$, particularly for $M_\nu = 0.6$ eV, the slopes of the N-body power spectra in Fig. 9 (Top) flatten in a manner not captured by either `Cosmic-Eν` or `MuFLR`.

5.3 Rapidly evolving dark energy

Since neutrino mass bounds are dependent upon constraints on the growth factor of large-scale structure, they are degenerate with variations in the dark energy equation of state. For example, Upadhye (2019) found a factor-of-three degradation in the 95 per cent-confidence M_ν bound when w_0 and w_a were allowed to vary. In recognition of this degeneracy, MT4 and `Cosmic-Eν` allow for substantial variations in w_0 and $w_0 + w_a$. Here, we test the accuracy of `Cosmic-Eν` for such a rapidly varying equation of state by comparison to a νwCDM N-body simulation.

Our νwCDM simulation is a part of a forthcoming suite of simulations designed to investigate the simultaneous variation of parameters describing physics beyond the ΛCDM model. Its cosmological parameters are $\Omega_{\mathrm{m},0}h^2 = 0.1429476$, $\Omega_{\mathrm{b},0}h^2 = 0.0240724$, $\Omega_{\nu,0}h^2 = 0.0015$, $\sigma_8 = 0.841333$, $h = 0.613333$, $n_s = 0.9544$, $w_0 = -1.036$, and $w_a = -0.872$. It tracked $N = 1260^3$ particles in a cubic volume with box length 1400 Mpc and periodic boundary conditions. Its 3LPT initial conditions were produced at $z = 31$ using `Monofonic` (Hahn, Rampf & Uhlemann 2021; Rampf, Uhlemann & Hahn 2021), and it tracked massive neutrinos using the δf method of Elbers et al. (2021, 2022). The simulation was evolved to $z = 0$ using `SWIFT` (Schaller et al.).

Fig. 10 tests `Cosmic-Eν` for this νwCDM simulation, with a rapidly varying equation of state, at redshifts 0 and 1. The accuracy of `Cosmic-Eν` at $z = 0$ is in line with our previous

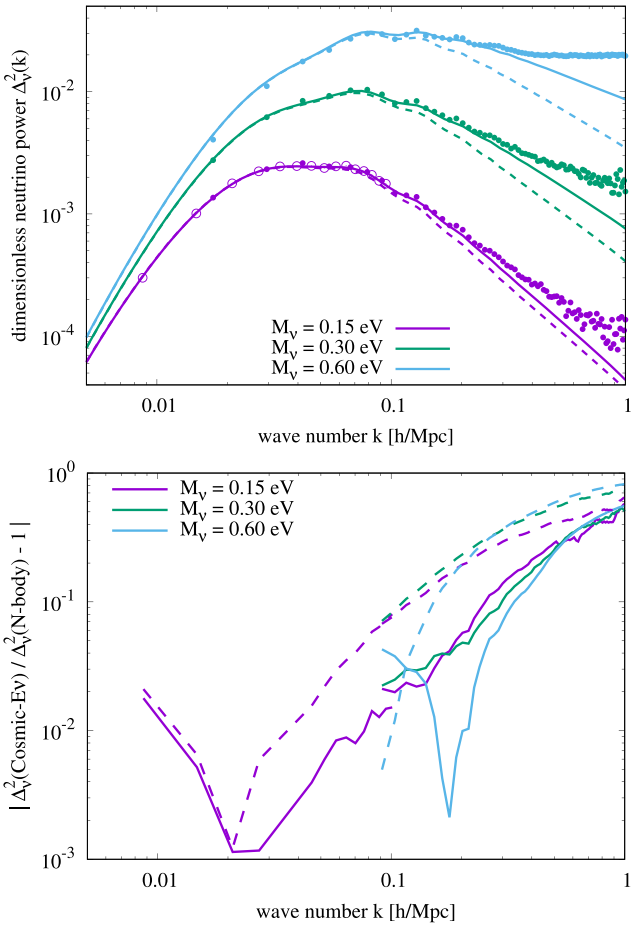

Figure 9. Comparison of `Cosmic-Eν` (solid) and `MuFLR` (dashed) to `GADGET-3` N-body neutrino power spectra of Adamek et al. (2023), at $z = 0$, for neutrino masses M_ν ranging from 0.15 to 0.6 eV. For $M_\nu = 0.15$ eV, the high-resolution 1024^3-particle simulations of that reference are used. (Top) N-body power spectra are shown as filled circles for $(512\ \mathrm{Mpc})^3$ simulation volumes and open circles for the $(1024\ \mathrm{Mpc})^3$ $M_\nu = 0.15$ eV simulation. (Bottom) Fractional errors in `Cosmic-Eν` and `MuFLR` compared with the $(512\ \mathrm{Mpc})^3$-box simulations at $k \geq 0.1\,h\,\mathrm{Mpc}^{-1}$ and the $(1024\ \mathrm{Mpc})^3$-box simulation below that wave number.

$\nu\Lambda$CDM comparisons to the `SWIFT` simulations: ≈ 20 per cent up to $k = 0.3\,h\,\mathrm{Mpc}^{-1} - 0.4\,h\,\mathrm{Mpc}^{-1}$, and ≈ 50 per cent up to $k = 1\,h\,\mathrm{Mpc}^{-1}$. At $z = 1$, `Cosmic-Eν` is slightly more accurate up to $k \approx 0.7\,h\,\mathrm{Mpc}^{-1}$, above which quantifying its accuracy becomes difficult due to residual shot noise. Compared with `MuFLR`, `Cosmic-Eν` is significantly more accurate at $z = 0$ and $k \geq 0.1\,h\,\mathrm{Mpc}^{-1}$, while at $z = 1$, `Cosmic-Eν` is somewhat more accurate at $k \gtrsim 0.1\,h\,\mathrm{Mpc}^{-1}$ Thus we conclude that even $|w_a| \sim 1$ does not diminish the accuracy of `Cosmic-Eν`.

This section has quantified the accuracy of the `Cosmic-Eν` $\Delta^2_\nu(k, z)$ emulator across a wide range of M_ν, for a cosmological constant as well as a rapidly evolving equation of state, by comparison to N-body simulations using a few very different massive neutrino simulation methods. Its error at $z = 0$ is fairly consistent across a wide range of methods: a few per cent up to $k \approx 0.15\,h\,\mathrm{Mpc}^{-1}$, an ≈ 20 per cent power underestimate at $k = 0.4\,h\,\mathrm{Mpc}^{-1}$, and an ≈ 50 per cent underestimate at $k = 1\,h\,\mathrm{Mpc}^{-1}$. Since $\Delta^2_\nu(k, z)$ for $k \gg k_{\mathrm{FS}}$ scales approximately as M_ν^4 (Ringwald & Wong 2004; Wong 2008), these underestimates at $k = 0.4\,h\,\mathrm{Mpc}^{-1}$ and $k = 1\,h\,\mathrm{Mpc}^{-1}$ are consistent with 5 per cent and 13 per cent biases in M_ν, respectively.

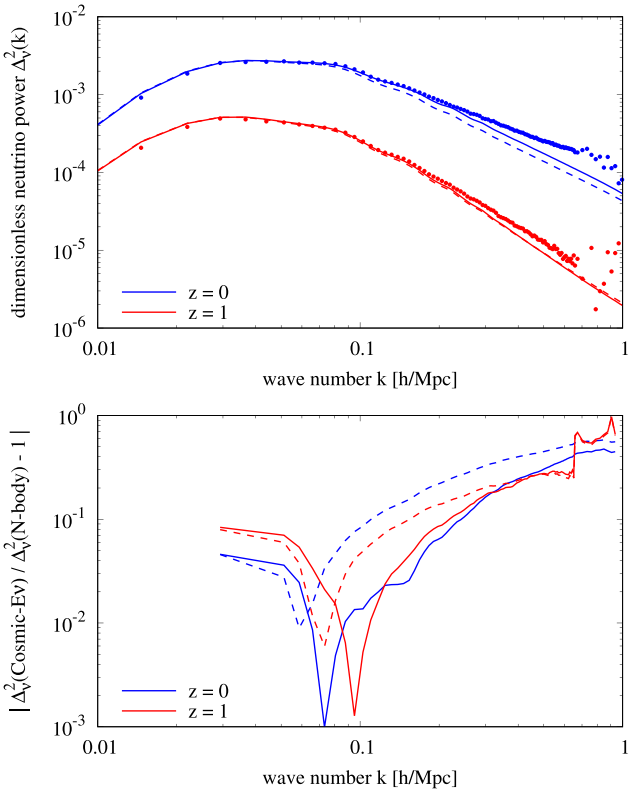

Figure 10. Comparison of `Cosmic-Eν` (solid) and `MuFLR` (dashed) to our νwCDM N-body simulation for a model with rapidly-evolving dark energy, $w_0 = -1.036$ and $w_a = -0.872$, as well as $\Omega_{\nu,0}h^2 = 0.0015$ ($M_\nu = 0.14$ eV), at $z = 0$ and 1. (Top) N-body power spectra are shown as points. (Bottom) Fractional errors in `Cosmic-Eν` and `MuFLR` compared with the N-body simulation, shown using a centred 10-point moving average.

Before proceeding, we comment upon the discrepancy between the 14 per cent error in `FlowsForTheMasses` reported in Chen et al. (2023b), over the entire range $k \leq 1\ h\,\mathrm{Mpc}^{-1}$, and the larger differences with N-body simulations evident in Figs 8, 9, and 10 for $k \approx 1\ h\,\mathrm{Mpc}^{-1}$. There are two possibilities: errors in the hybrid simulations of Chen et al. (2023a) used to test `FlowsForTheMasses`, and a systematic error causing the discrepancies between differing N-body implementations of neutrinos.

Errors in Chen et al. (2023a) may be due to a finite number of neutrino flows, residual shot noise, and a finite simulation volume. Finite-flow-number errors in the LR calculations of Chen et al. (2021a) were found to be ≈ 10 per cent, and non-linear neutrino clustering likely increases them somewhat. Of course, increasing the number of flows, or sampling the Fermi-Dirac distribution more efficiently, will improve the accuracy of `FlowsForTheMasses` as well as the simulations. While the estimated simulation shot noise $(kL_{\rm sim})^3/(2\pi^2 N_{\rm sim})$, for $N_{\rm sim}$ particles in a volume $L_{\rm sim}^3$, was subtracted from $\Delta_\nu^2(k)$, residual shot noise remains. In an effort to mitigate shot noise, Chen et al. (2023a) chose a small box, $L_{\rm sim} = 128\ \mathrm{Mpc}\,h^{-1}$, at the cost of neglecting the contributions of larger modes to small-scale non-linear growth.

Meanwhile, the 30 per cent − 40 per cent spread among the different simulations in Fig. 8 at $k = 1\ h\,\mathrm{Mpc}^{-1}$ suggests small-scale systematic errors in some of these methods. The `gevolution` power spectrum of Adamek et al. (2016a) is about 15 per cent − 20 per cent lower than `SWIFT` at $k = 1\ h\,\mathrm{Mpc}^{-1}$, meaning that correcting the 14 per cent underestimate of `Cosmic-Eν` relative to Chen et al. (2023a), as well as the ≈ 10 per cent underestimate due to a finite number of flows, would put `Cosmic-Eν` within 5 per cent − 10 per cent of `gevolution`. Moreover, in the $k \lesssim 0.2\ h\,\mathrm{Mpc}^{-1}$ range where the N-body methods of Fig. 8 agree with one another to a few per cent, `Cosmic-Eν` also agrees with them at that level. The `SWIFT-Cosmic-Eν` difference rises along with the `CO`$\mathcal{N}$`CEPT-gevolution` difference in Fig. 8. Thus we cannot conclusively attribute the discrepancy between the 14 per cent `FlowsForTheMasses` error estimate of Chen et al. (2023b) and the 49 per cent `SWIFT-Cosmic-Eν` difference to errors in Chen et al. (2023a).

6 NON-LINEAR ENHANCEMENT

6.1 Parameter-sensitivity of the enhancement ratio

Now that we have quantified the accuracy of `Cosmic-Eν`, we may use it to study the non-linear clustering of massive neutrinos. We focus here on the non-linear enhancement ratio $R_\nu(k, z)$ of equation (31), that is, the ratio of the neutrino power spectra using `FlowsForTheMasses` and `MuFLR`, with the CDM + baryon treatment held fixed. We emulate R_ν by taking the ratio of `Cosmic-Eν` to a `MuFLR` emulator.

Fig. 11 compares perturbative (dashed) and emulated (solid) calculations of R_ν at $z = 0$ for the out-of-sample models of Table 2. A couple trends are evident. First, the total neutrino power is typically more accurate at high k than individual decile powers. At $k = 1\ h\,\mathrm{Mpc}^{-1}$, the emulated R_ν agrees with the `FlowsForTheMasses` computation to better than 2 per cent for eight of the 10 models. For decile 0, this error rises to 3.2 per cent, and for decile 1–4.8 per cent.

Secondly, the lower deciles and higher neutrino masses tend to have smaller errors. This is due to the fact that larger L and smaller masses lead to larger average velocities, hence more prominent oscillatory behaviour in the free-streaming limit, making these flows difficult to emulate. Both of these trends are consistent with the individual-decile out-of-sample and holdout tests of Section 4.2.

Next, we consider the sensitivity of R_ν to the cosmological parameters. As a fiducial model at which to test this sensitivity, we choose a ΛCDM model in which each parameter except for w_0 and w_a is set to the midpoint of its range in Table 1. Fig. 12 shows the derivative of $\log R_\nu$ with respect to each parameter about this fiducial model. We have checked that the results are qualitatively similar for $\Omega_{\nu,0}h^2 = 0.002$.

Above $k \approx k_{\rm FS} = 0.16\ h\,\mathrm{Mpc}^{-1}$, the dominant effect is a rise in R_ν with $\Omega_{\nu,0}h^2 \propto M_\nu$. This can be understood by noting that, with all parameters in Table 1 other than $\Omega_{\nu,0}h^2$ held fixed, the small-scale linear and linear-response neutrino power scales as $(\Omega_{\nu,0}h^2)^4$, while the non-linear power rises relative to the LR power. We can estimate the non-linear enhancement to the linear scaling law from Fig. 12:

$$\frac{\partial \log(\Delta_\nu^2[\mathrm{nonlin}])}{\partial \log(\Omega_{\nu,0}h^2)} = \frac{\partial \log(\Delta_\nu^2[\mathrm{LR}])}{\partial \log(\Omega_{\nu,0}h^2)} + \frac{\partial \log(R_\nu)}{\partial \log(\Omega_{\nu,0}h^2)}. \tag{35}$$

The final term on the right, the logarithmic derivative of R_ν, is $\Omega_{\nu,0}h^2$ times $d\ln(R_\nu)/d(\Omega_{\nu,0}h^2)$ from the figure, or about 0.57 at $k = 1\ h\,\mathrm{Mpc}^{-1} \approx 6k_{\rm FS}$. A similar analysis for $\Omega_{\nu,0}h^2 = 0.002$ finds this scaling enhancement to be 0.35 at $k = 0.4\ h\,\mathrm{Mpc}^{-1} \approx 6k_{\rm FS}$ and 0.59 at $k = 1\ h\,\mathrm{Mpc}^{-1} \approx 15k_{\rm FS}$, suggesting a rise in this scaling enhancement with both $\Omega_{\nu,0}h^2$ and $k/k_{\rm FS}$. Section 6.2 will explore this enhancement further.

At $k \gtrsim k_{\rm FS} = 0.16\ h\,\mathrm{Mpc}^{-1}$, the next most significant parameter for determining R_ν is the physical baryon fraction $\Omega_{\rm b,0}h^2$. The oscillatory nature of $\partial \ln R_\nu/\partial(\Omega_{\rm b,0}h^2)$ for $k \lesssim 0.8\ h\,\mathrm{Mpc}^{-1}$, and its decline for

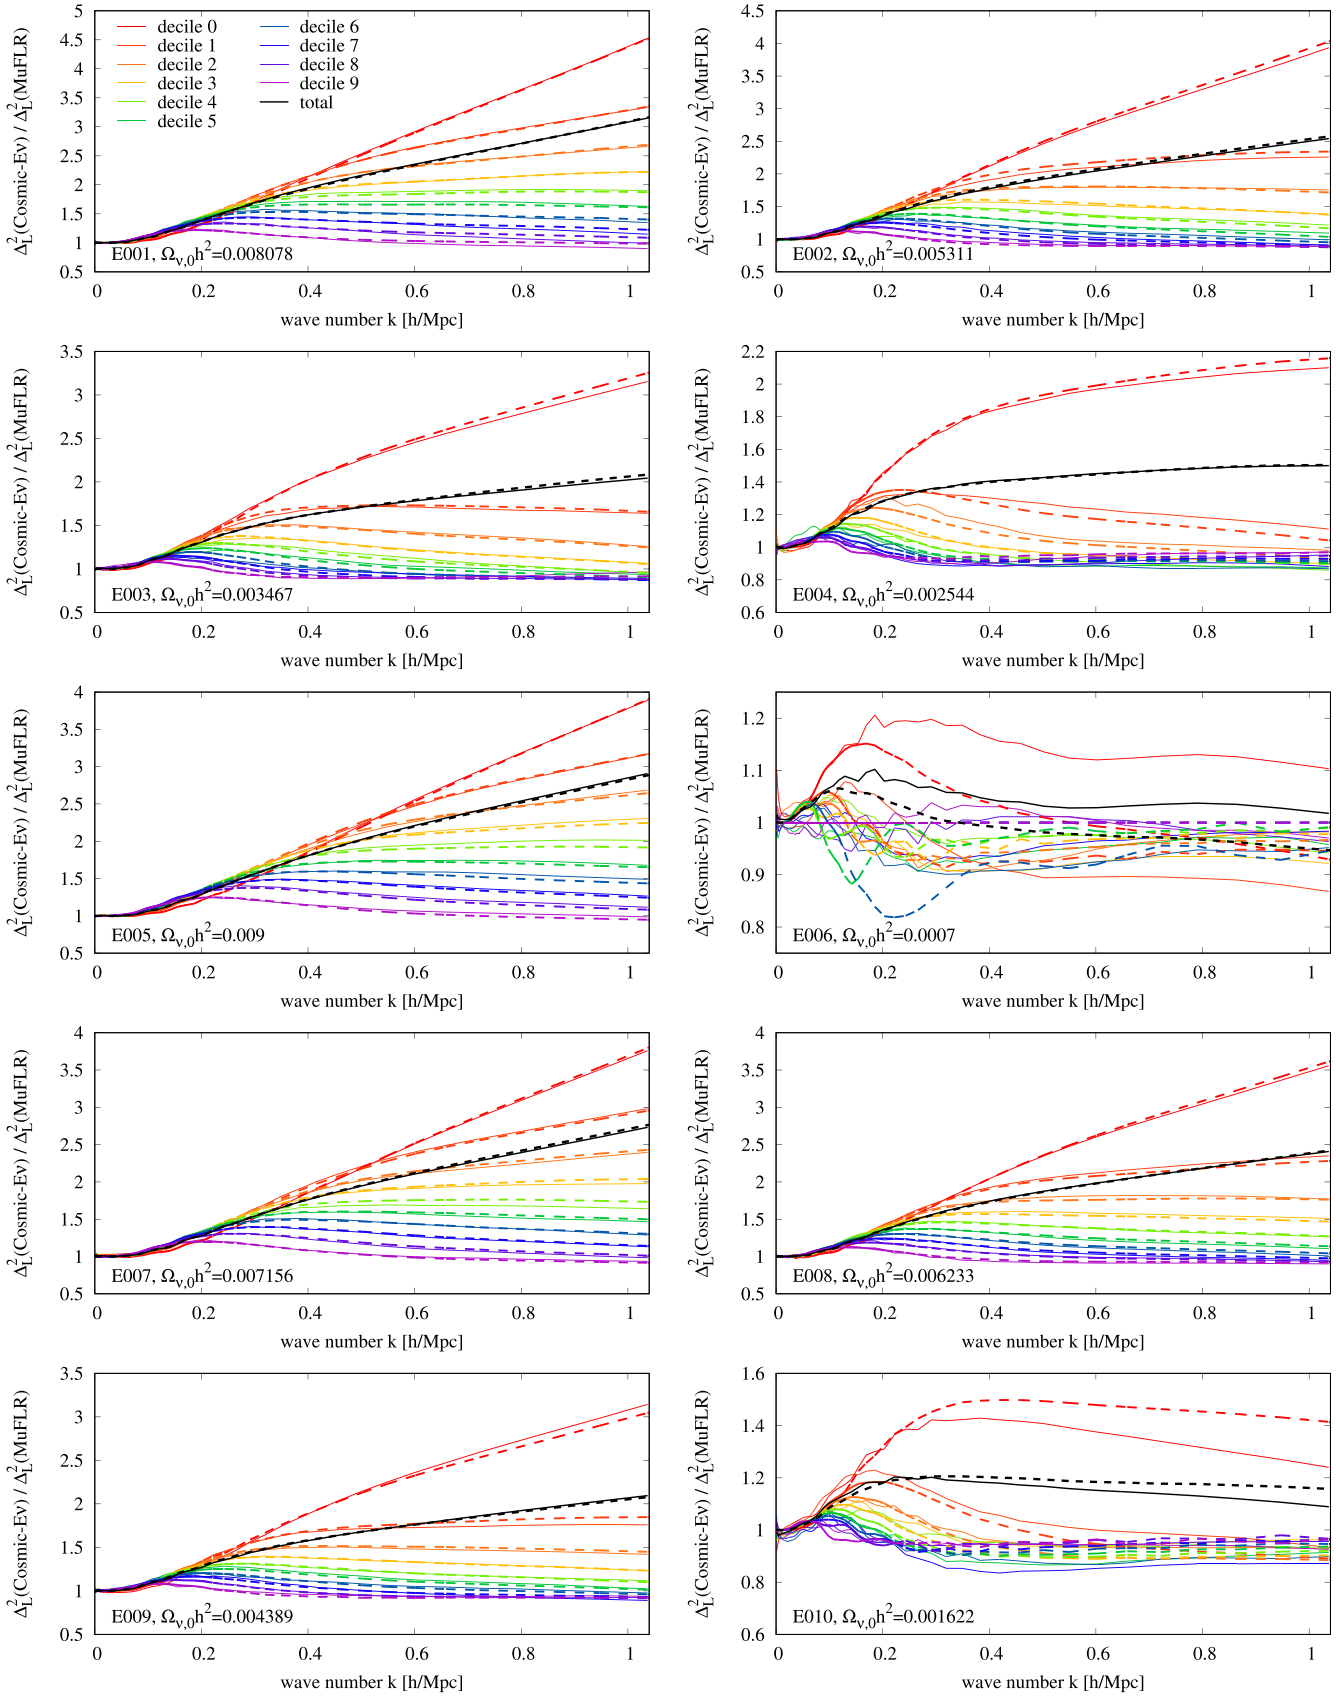

Figure 11. $R_\nu(k, z)$ for each L, at $z = 0$, for the out-of-sample cosmologies of Table 2, computed using `Cosmic-Eν` (solid) and `FlowsForTheMasses` (dashed).

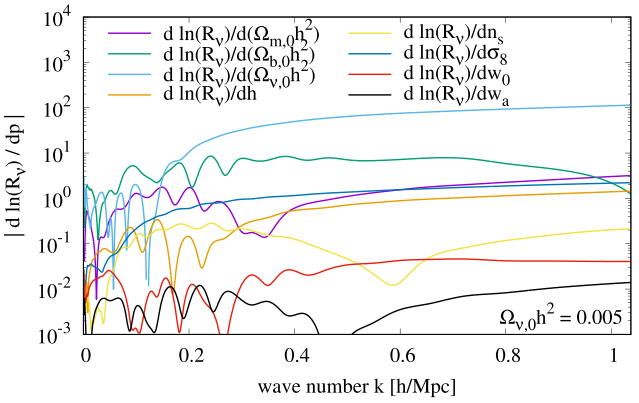

Figure 12. Sensitivity of $R_\nu(k,\ z)$ at $z\ =\ 0$ to variations of the eight cosmological parameters shown in Table 1.

Table 3. Parameter sensitivities from Fig. 12 multiplying the 95 per cent CL intervals Δp for the νwCDM model of Upadhye (2019), analyzed using the combination of Planck, BOSS, and JLA supernova data, marginalized over a five-parameter bias model. Third and fourth columns, respectively, use sensitivities at $k = 0.4\ h\,\mathrm{Mpc}^{-1}$ and $k = 1\ h\,\mathrm{Mpc}^{-1}$. We assume $\Delta\Omega_{\mathrm{m},0}h^2 \approx \Delta\Omega_{\mathrm{c},0}h^2$.

Parameter	Δp[95 %]	$\left\lvert\frac{\partial \log R_\nu(0.4h/\mathrm{Mpc})}{\partial p}\right\rvert \Delta p$	$\left\lvert\frac{\partial \log R_\nu(1.0h/\mathrm{Mpc})}{\partial p}\right\rvert \Delta p$
$\Omega_{\nu,0}h^2$	0.0061	0.32	0.69
σ_8	0.087	0.10	0.19
n_s	0.181	0.019	0.039
h	0.028	0.016	0.037
w_0	0.64	0.013	0.026
w_a	3.0	0.0086	0.042
$\Omega_{\mathrm{b},0}h^2$	0.00059	0.0041	0.0007
$\Omega_{\mathrm{m},0}h^2$	0.0053	0.0039	0.017

$k \gtrsim 0.9\ h\,\mathrm{Mpc}^{-1}$, suggests that $\Omega_{\mathrm{b},0}h^2$ affects R_ν primarily through the baryon acoustic oscillations (BAO). Modifications to the BAO, in turn, are amplified by the non-linear clustering of neutrinos.

After $\Omega_{\nu,0}h^2$ and $\Omega_{\mathrm{b},0}h^2$, the next parameters to which R_ν is most sensitive at $k \gtrsim k_\mathrm{FS}$ are σ_8 and $\Omega_{\mathrm{m},0}h^2$. At first glance, the relative sensitivities of R_ν to σ_8, $\Omega_{\mathrm{m},0}h^2$, and $\Omega_{\mathrm{b},0}h^2$ appear to contradict Fig. 3. However, the range of R_ν associated with each parameter in that figure is the derivative in Fig. 12 times the parameter range in Table 1. This range for σ_8 is about six times larger than for $\Omega_{\mathrm{m},0}h^2$, and 100 times larger than for $\Omega_{\mathrm{b},0}h^2$.

Thus, given a particular data combination, we may define an alternative sensitivity measure for each parameter p by multiplying $\partial \log R_\nu/\partial p$ from Fig. 12 by the range Δp allowed by the data. As an example, we choose the νwCDM analysis of Upadhye (2019), constrained using a combination of CMB, galaxy, and supernova data, and marginalized over a five-parameter model of scale-dependent galaxy bias. We approximate the 95 per cent CL interval of $\Omega_{\mathrm{m},0}h^2$ by the corresponding one for the CDM alone, its dominant component. Table 3 shows the result at two wave numbers. While $\Omega_{\nu,0}h^2$ remains by far the most significant parameter for determining R_ν, σ_8 is also important. In summary, while R_ν is most sensitive to $\Omega_{\nu,0}h^2$, σ_8, $\Omega_{\mathrm{m},0}h^2$, and $\Omega_{\mathrm{b},0}h^2$, the first two of these are the most important given current parameter constraints.

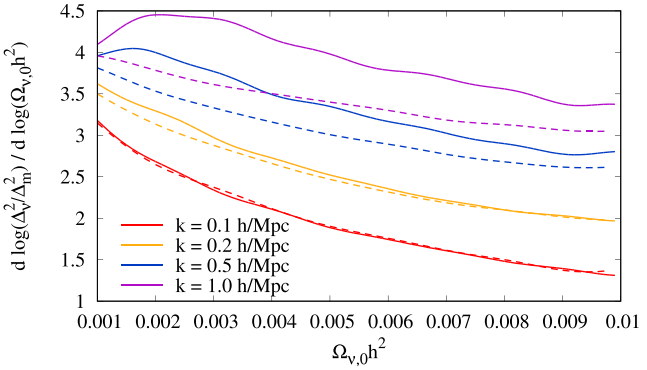

Figure 13. Logarithmic derivative of $\Delta_\nu^2(k,\ z)/\Delta_\mathrm{m}^2(k,\ z)$ with respect to $\Omega_{\nu,\,0}h^2$ at several wave numbers k, at $z = 0$. Solid lines use `Cosmic-Eν` for Δ_ν^2, and dashed lines use `MuFLR`.

6.2 Relative clustering in the free-streaming limit

Next, we consider further the mass scaling of the neutrino power spectrum in the free-streaming limit, raised in the previous subsection. Ringwald & Wong (2004) argues that neutrino LR to non-linear CDM + baryon growth results in a scaling $\Delta_\nu^2/\Delta_\mathrm{m}^2 \propto M_\nu^4$, while the Tremaine-Gunn bound of Shu (1978), Tremaine & Gunn (1979), Shu (1987), and Kull, Treumann & Bohringer (1996) implies $\partial \log(\Delta_\nu^2/\Delta_\mathrm{m}^2)/\partial \log(M_\nu) < 6$. Ringwald & Wong (2004) demonstrate using N-body simulations that haloes approach this latter bound, and, further, that the bound can be exceeded, especially in the case of small M_ν, if one includes all neutrinos present, rather than only those captured by the halo's gravitational potential.

Thus, we may expect $\partial \log(\Delta_\nu^2/\Delta_\mathrm{m}^2)/\partial \log(M_\nu)$ to rise above four while remaining below six. Fig. 13 numerically differentiates the ratio of the neutrino power spectrum to the MT4 total-matter power spectrum with respect to $\Omega_{\nu,\,0}h^2 \propto M_\nu$ with a step size of 1 per cent in $\Omega_{\nu,\,0}h^2 \propto M_\nu$. Solid and dashed lines, respectively, use non-linear (`Cosmic-Eν`) and LR (`MuFLR`) neutrino power spectra.

Consider first the larger wave numbers. The `MuFLR` curves approach 4 from below but never exceed it, as expected. Meanwhile, the `Cosmic-Eν` non-linear logarithmic derivatives for $k = 0.5\ h\,\mathrm{Mpc}^{-1}$ and $k = 1\ h\,\mathrm{Mpc}^{-1}$ both exceed 4 for small $\Omega_{\nu,\,0}h^2$, where these wave numbers are many times the free-streaming scale. The non-linear enhancement to the mass scaling in equation (35) is the difference between the solid and dashed lines. Focusing on $k\ =\ 1\ h\,\mathrm{Mpc}^{-1}$, we find this to be 0.67 for $\Omega_{\nu,\,0}h^2\ =\ 0.002$ and 0.58 for $\Omega_{\nu,\,0}h^2\ =\ 0.005$. Since a 3.5 per cent emulator error implies an error of $\sim$0.1 in the logarithmic derivative, these are consistent with the results of Section 6.1.

This emulator error means that the difference between the linear and non-linear curves for $k = 0.1\ h\,\mathrm{Mpc}^{-1}$ is consistent with zero. The same is true for $k\ =\ 0.2\ h\,\mathrm{Mpc}^{-1}$ for $\Omega_{\nu,\,0}h^2 \gtrsim 0.003$. Further, the small oscillations observed in some of the logarithmic derivatives are consistent with emulator fluctuations. Thus Sections 6.1–6.2 consistently demonstrate a non-linear enhancement of $\gtrsim$0.5 to the power law scaling $\partial \log(\Delta_\nu^2)/\partial \log(\Omega_{\nu,0}h^2)$ at $k\ =\ 1\ h\,\mathrm{Mpc}^{-1}$ for $0.002 \lesssim \Omega_{\nu,\,0}h^2 \lesssim 0.005$.

6.3 Neutrino contribution to the matter power

The MT4 emulator includes fully linear neutrinos, as implemented in the `CAMB` code of Lewis et al. (2000) and Lewis & Bridle (2002), in their CDM + baryon and total matter power spectra, as described in Saito, Takada & Taruya (2008), Agarwal & Feldman (2011),

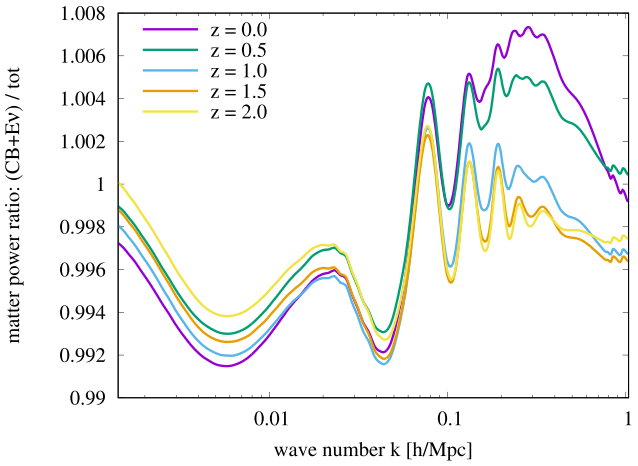

Figure 14. Direct non-linear neutrino contribution to the matter power spectrum at several z for $\Omega_{\nu,0}h^2 = 0.01$. $\Delta^2_{\rm m}(k,z)$, computed using the MT4 CDM + baryon power and the `Cosmic-Eν` neutrino power as in equation (36), is divided by the MT4 emulated total matter power $\Delta^2_{\rm m}(k,z)$.

and Upadhye et al. (2014, 2016). Since the non-linear clustering of neutrinos increases their power by an order-of-magnitude relative to linear theory, as shown in Fig. 2, we quantify here the impact of neutrino non-linearity on the matter power spectrum. Neutrinos will affect $\Delta^2_{\rm m}(k,z)$ in two ways: indirectly, by adding to the gravitational potential, hence enhancing CDM + baryon clustering; and directly, through their inclusion in the total matter power.

Quantifying the indirect effect precisely, by incorporating `FlowsForTheMasses` into an N-body simulation, is beyond the scope of this study. However, we may bound this effect. Chen et al. (2021a) carries out an N-body simulation with neutrino LR through the `MuFLR` code. For the largest neutrino fraction considered here, $\Omega_{\nu,0}h^2 = 0.01$, that study finds an indirect enhancement of 0.05 per cent to the CDM + baryon power spectrum. Since the non-linear enhancement ratios R_ν are less than 5 in Figs 3 and 11, the indirect enhancement is < 0.25 per cent. A more accurate estimate directly multiplying the LR enhancement of Chen et al. (2021a) by R_ν for their model finds an indirect enhancement of 0.16 per cent.

Fig. 14 quantifies the direct effect, which is everywhere less than 1 per cent. Even this is a slight overestimate, as `FlowsForTheMasses`, hence `Cosmic-Eν`, assume that CDM, baryon, and neutrino density-contrast monopoles are perfectly correlated. Under this approximation, the matter power spectrum is

$$\Delta^2_{\rm m}(k,z) = \left(\frac{\Omega_{\rm cb}}{\Omega_{\rm m}}\sqrt{\Delta^2_{\rm cb}(k,z)} + \frac{\Omega_\nu}{\Omega_{\rm m}}\sqrt{\Delta^2_\nu(k,z)}\right)^2. \quad (36)$$

Bird et al. (2018) shows that the actual neutrino-CDM correlation function drops below unity by ≤ 4 per cent for $k \leq 0.5\,h\,{\rm Mpc}^{-1}$ and ≤ 8 per cent for $k \leq 1\,h\,{\rm Mpc}^{-1}$. The smallness of this deviation is due to the fact that the initially-slowest neutrinos, which contribute the most to small-scale clustering, are also closely correlated with the CDM. This correlation being slightly less than one implies that the actual direct contribution of neutrinos to $\Delta^2_{\rm m}$ is slightly smaller than in Fig. 14.

7 CONCLUSIONS

`FlowsForTheMasses`, the first non-linear perturbative power spectrum calculation for free-streaming particles such as massive neutrinos, provides detailed information on the clustering of neutrinos of different initial momenta. We have emulated the total non-linear neutrino power spectrum as well as separate power spectra for the 10 momentum deciles, each representing a 10th of the neutrino number density. Our emulated $\Delta^2_\nu(k,z)$ agrees precisely with `FlowsForTheMasses` to < 3.5 per cent for $10^{-3}\,h\,{\rm Mpc}^{-1} \leq k \leq 1\,h\,{\rm Mpc}^{-1}$ and $0 \leq z \leq 3$, as shown in Fig. 4. Individual-decile errors range from about twice as large for the lowest momenta to four times as large for the fastest-moving neutrinos with highly oscillatory density contrasts; see Fig. 7. We have released our emulator as `Cosmic-Eν`.

Comparing `Cosmic-Eν` to the highest-resolution simulations of Adamek et al. (2023) in Fig. 9, we found agreement to 3 per cent up to $k = 3k_{\rm FS} = 0.17\,h\,{\rm Mpc}^{-1}$ and 19 per cent to $k = 0.4\,h\,{\rm Mpc}^{-1}$. Above this wave number, `Cosmic-Eν` increasingly underpredicts the simulations of Adamek et al. (2023), with this underprediction reaching nearly 50 per cent by $k = 1\,h\,{\rm Mpc}^{-1}$. Even this error is not substantially larger than the 30 per cent − 40 per cent scatter between different simulation methods seen in Fig. 8, so we cannot definitively attribute it either to a non-perturbative effect beyond the capabilities of `FlowsForTheMasses` or to a systematic error in the simulations. Importantly, `Cosmic-Eν` provides a neutrino power spectrum in about ten milliseconds on a standard desktop machine, and we have confirmed that its accuracy is unaffected by rapid variations in the dark energy equation of state.

One strength of the emulation technique is our ability to differentiate numerically the emulated function without the result being dominated by the shot noise and sample variance affecting N-body power spectra. Section 6 took full advantage of this capability by studying the non-linear enhancement ratio $R_\nu(k,z)$ of equation (31) and the neutrino-to-matter ratio $\Delta^2_\nu/\Delta^2_{\rm m}$. Differentiating R_ν with respect to each of the cosmological parameters, we find that it is most sensitive to the physical neutrino density $\Omega_{\nu,0}h^2$, but also to $\Omega_{{\rm b},0}h^2$, σ_8, and $\Omega_{{\rm m},0}h^2$. Furthermore, we demonstrated a non-linear enhancement of $\approx$0.5 to the free-streaming-limit scaling $\partial\log(\Delta^2_\nu/\Delta^2_{\rm m})/\partial\log(M_\nu) \to 4$, meaning that non-linear clustering makes the small-scale density of neutrinos even more sensitive to their mass. Our results demonstrate the speed and efficacy of the emulation technique in neutrino cosmology.

ACKNOWLEDGEMENTS

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 769130). YYYW is supported in part by the Australian Research Council's Future Fellowship (project FT180100031). This research is enabled by the Australian Research Council's Discovery Project (project DP170102382) funding scheme, and includes computations using the computational cluster Katana supported by Research Technology Services at UNSW Sydney. The authors are grateful to J. Conley and S. Habib for insightful conversations.

DATA AVAILABILITY

No new experimental data were generated or analysed in support of this research. The `Cosmic-Eν` code underlying this article is available on Github at https://github.com/upadhye/Cosmic-Enu.

REFERENCES

Adamek J., Durrer R., Kunz M., 2014, Class. Quant. Grav., 31, 234006
Adamek J., Daverio D., Durrer R., Kunz M., 2016a, J. Cosmol. Astropart. Phys., 2016, 053
Adamek J., Daverio D., Durrer R., Kunz M., 2016b, Nature Phys., 12, 346
Adamek J., Durrer R., Kunz M., 2017, J. Cosmol. Astropart. Phys., 11, 004
Adamek J. et al., 2023, J. Cosmol. Astropart. Phys., 06, 035
Agarwal S., Feldman H. A., 2011, MNRAS, 410, 1647
Aguilar-Arevalo A. A. et al., 2022, Phys. Rev. Lett., 129, 201801
Aker M. et al., 2022, Phys. Rev. Lett., 129, 011806
Ali-Haimoud Y., Bird S., 2012, MNRAS, 428, 3375
Alvarez-Ruso L., Saul-Sala E., 2021, Eur. Phys. J. ST, 230, 4373
Audren B., Lesgourgues J., Bird S., Haehnelt M. G., Viel M., 2013, J. Cosmol. Astropart. Phys., 01, 026
Banerjee A., Powell D., Abel T., Villaescusa-Navarro F., 2018, J. Cosmol. Astropart. Phys., 09, 028
Banerjee A., Castorina E., Villaescusa-Navarro F. a., Court T., Viel M., 2020, J. Cosmol. Astropart. Phys., 06, 032
Baracchini E. et al., 2018, preprint (arXiv:1808.01892)
Beck A. M. et al., 2016, MNRAS, 455, 2110
Bernardeau F., Colombi S., Gaztanaga E., Scoccimarro R., 2002, Phys. Rept., 367, 1
Betti M. G. et al., 2019, J. Cosmol. Astropart. Phys., 07, 047
Bird S., Ali-Haïmoud Y., Feng Y., Liu J., 2018, MNRAS, 481, 1486
Biswas R., Heitmann K., Habib S., Upadhye A., Pope A., Frontiere N., 2019, preprint (arXiv:1901.10690)
Blas D., Lesgourgues J., Tram T., 2011, J. Cosmol. Astropart. Phys., 07, 034
Bocquet S., Heitmann K., Habib S., Lawrence E., Uram T., Frontiere N., Pope A., Finkel H., 2020, ApJ, 901, 5
Böhringer H., Chon G., 2016, Mod. Phys. Lett. A, 31, 1640008
Capozzi F., Lisi E., Marrone A., Palazzo A., 2018, Prog. Part. Nucl. Phys., 102, 48
Chen J. Z., Upadhye A., Wong Y. Y. Y., 2021a, J. Cosmol. Astropart. Phys., 03, 065
Chen J. Z., Upadhye A., Wong Y. Y. Y., 2021b, J. Cosmol. Astropart. Phys., 2021, 078
Chen J. Z., Mosbech M. R., Upadhye A., Wong Y. Y. Y., 2023a, J. Cosmol. Astropart. Phys., 03, 012
Chen J. Z., Upadhye A., Wong Y. Y. Y., 2023b, J. Cosmol. Astropart. Phys., 05, 046
Chevallier M., Polarski D., 2001, Int. J. Mod. Phys. D, 10, 213
Chiang C.-T., Hu W., Li Y., Loverde M., 2018, Phys. Rev. D, 97, 123526
Chiang C.-T., LoVerde M., Villaescusa-Navarro F., 2019, Phys. Rev. Lett., 122, 041302
Chudaykin A., Ivanov M. M., 2019, J. Cosmol. Astropart. Phys., 11, 034
Costanzi M., Villaescusa-Navarro F., Viel M., Xia J.-Q., Borgani S., Castorina E., Sefusatti E., 2013, J. Cosmol. Astropart. Phys., 12, 012
D'Eramo F., Di Valentino E., Giarè W., Hajkarim F., Melchiorri A., Mena O., Renzi F., Yun S., 2022, J. Cosmol. Astropart. Phys., 09, 022
Dakin J., Brandbyge J., Hannestad S., Haugbølle T., Tram T., 2019, J. Cosmol. Astropart. Phys., 02, 052
Dakin J., Hannestad S., Tram T., 2022, MNRAS, 513, 991
de Salas P. F., Forero D. V., Ternes C. A., Tortola M., Valle J. W. F., 2018, Phys. Lett. B, 782, 633
Denton P. B., 2022, Phys. Rev. Lett., 129, 061801
Di Valentino E., Melchiorri A., 2022, ApJ, 931, L18
Di Valentino E., Melchiorri A., Silk J., 2020, J. Cosmol. Astropart. Phys., 01, 013
Di Valentino E., Gariazzo S., Giarè W., Melchiorri A., Mena O., Renzi F., 2023, Phys. Rev. D, 107, 103528
Dupuy H., Bernardeau F., 2014, J. Cosmol. Astropart. Phys., 01, 030
Dupuy H., Bernardeau F., 2015a, J. Cosmol. Astropart. Phys., 03, 030
Dupuy H., Bernardeau F., 2015b, J. Cosmol. Astropart. Phys., 08, 053
Eisenstein D. J., Hu W., 1997, ApJ, 511, 5
Eisenstein D. J., Hu W., 1998, ApJ, 496, 605
Elbers W., Frenk C. S., Jenkins A., Li B. o., Pascoli S., 2021, MNRAS, 507, 2614
Elbers W., Frenk C. S., Jenkins A., Li B., Pascoli S., 2022, MNRAS, 516, 3821
Esteban I., Gonzalez-Garcia M. C., Maltoni M., Schwetz T., Zhou A., 2020, J. High Energy Phys., 09, 178
Fang X., Blazek J. A., McEwen J. E., Hirata C. M., 2017, J. Cosmol. Astropart. Phys., 02, 030
Fidler C., Rampf C., Tram T., Crittenden R., Koyama K., Wands D., 2015, Phys. Rev. D, 92, 123517
Fidler C., Tram T., Rampf C., Crittenden R., Koyama K., Wands D., 2016, J. Cosmol. Astropart. Phys., 09, 031
Fidler C., Tram T., Rampf C., Crittenden R., Koyama K., Wands D., 2017, J. Cosmol. Astropart. Phys., 12, 022
Fidler C., Kleinjohann A., Tram T., Rampf C., Koyama K., 2019, J. Cosmol. Astropart. Phys., 01, 025
Font-Ribera A., McDonald P., Mostek N., Reid B. A., Seo H.-J., Slosar A., 2014, J. Cosmol. Astropart. Phys., 05, 023
Führer F., Wong Y. Y. Y., 2015, J. Cosmol. Astropart. Phys., 03, 046
Garny M., Taule P., 2021, J. Cosmol. Astropart. Phys., 01, 020
Garny M., Taule P., 2022, J. Cosmol. Astropart. Phys., 09, 054
Gattiker J., Klein N., Lawrence E., Hutchings G., 2020, lanl/SEPIA, Zenodo
Giarè W., Renzi F., Melchiorri A., Mena O., Di Valentino E., 2022, MNRAS, 511, 1373
Gogoi A., Sharma R. K., Chanda P., Das S., 2021, ApJ, 915, 132
Gong J.-O., Hwang J.-c., Noh H., Wu D. C. L., Yoo J., 2017, J. Cosmol. Astropart. Phys., 10, 027
Hahn O., Rampf C., Uhlemann C., 2021, MNRAS, 503, 426
Hamilton A. J. S., 2000, MNRAS, 312, 257
Heitmann K., Higdon D., White M., Habib S., Williams B. J., Wagner C., 2009, ApJ, 705, 156
Heitmann K., White M., Wagner C., Habib S., Higdon D., 2010, ApJ, 715, 104
Heitmann K. et al., 2016, ApJ, 820, 108
Hu W., Eisenstein D. J., 1998, ApJ, 498, 497
Hwang J.-c., Noh H., 2006a, Gen. Rel. Grav., 38, 703
Hwang J.-C., Noh H., 2006b, MNRAS, 367, 1515
Hwang J.-c., Noh H., 2007, Phys. Rev. D, 76, 103527
Hwang J.-c., Noh H., 2013a, J. Cosmol. Astropart. Phys., 04, 035
Hwang J.-c., Noh H., 2013b, MNRAS, 433, 3472
Hwang J.-c., Noh H., Park C.-G., 2016, MNRAS, 461, 3239
Inman D., Emberson J. D., Pen U.-L., Farchi A., Yu H.-R., Harnois-Déraps J., 2015, Phys. Rev. D, 92, 023502
Inman D. et al., 2017, Phys. Rev. D, 95, 083518
Jeong D., Gong J.-O., Noh H., Hwang J.-c., 2011, ApJ, 727, 22
Kull A., Treumann R. A., Bohringer H., 1996, ApJ, 466, L1
Lawrence E., Heitmann K., White M., Higdon D., Wagner C., Habib S., Williams B., 2010, ApJ, 713, 1322
Lawrence E. et al., 2017, ApJ, 847, 50
Leauthaud A. et al., 2017, MNRAS, 467, 3024
Lesgourgues J., 2011, preprint (arXiv:1104.2932)
Lesgourgues J., Tram T., 2011, J. Cosmol. Astropart. Phys., 09, 032
Lesgourgues J., Matarrese S., Pietroni M., Riotto A., 2009, J. Cosmol. Astropart. Phys., 06, 017
Lewis A., Bridle S., 2002, Phys. Rev. D, 66, 103511
Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
Lin H., Gong Y., Chen X., Chan K. C., Fan Z., Zhan H., 2022, MNRAS, 515, 5743
Linder E. V., 2003, Phys. Rev. Lett., 90, 091301
Liu J., Bird S., Matilla J. M. Z., Hill J. C., Haiman Z., Madhavacheril M. S., Petri A., Spergel D. N., 2018, J. Cosmol. Astropart. Phys., 03, 049
LoVerde M., 2014, Phys. Rev. D, 90, 083530
LoVerde M., Zaldarriaga M., 2014, Phys. Rev. D, 89, 063502
Ma C.-P., Bertschinger E., 1995, ApJ, 455, 7
Magi M., Yoo J., 2022, J. Cosmol. Astropart. Phys., 09, 071
Marin-Gilabert T., Valentini M., Steinwandel U. P., Dolag K., 2022, MNRAS, 517, 5971
Mauland R., Elgarøy O., Mota D. F., Winther H. A., 2023, A&A, 674, A185
McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, MNRAS, 465, 2936

McCarthy I. G., Bird S., Schaye J., Harnois-Deraps J., Font A. S., Van Waerbeke L., 2018, MNRAS, 476, 2999
McCarthy I. G. et al., 2023, MNRAS, 526, 5494
McEwen J. E., Fang X., Hirata C. M., Blazek J. A., 2016, J. Cosmol. Astropart. Phys., 09, 015
Moran K. R. et al., 2023, MNRAS, 520, 3443
Petracca F., Marulli F., Moscardini L., Cimatti A., Carbone C., Angulo R. E., 2016, MNRAS, 462, 4208
Pietroni M., 2008, J. Cosmol. Astropart. Phys., 10, 036
Poulin V., Boddy K. K., Bird S., Kamionkowski M., 2018, Phys. Rev. D, 97, 123504
Rampf C., Uhlemann C., Hahn O., 2021, MNRAS, 503, 406
Ringwald A., Wong Y. Y., 2004, J. Cosmol. Astropart. Phys., 12, 005
Ryu S., Lee J., 2022, ApJ, 933, 189
Saito S., Takada M., Taruya A., 2008, Phys. Rev. Lett., 100, 191301
Schaller M. et al., 2023, MNRAS, preprint (arXiv:2305.13380)
Schmittfull M., Vlah Z., McDonald P., 2016, Phys. Rev. D, 93, 103528
Shu F. H., 1978, ApJ, 225, 83
Shu F. H., 1987, ApJ, 316, 502
Springel V., 2005, MNRAS, 364, 1105
Springel V. et al., 2008, MNRAS, 391, 1685
Springel V., Pakmor R., Zier O., Reinec ke M., 2021, MNRAS, 506, 2871
Sullivan J. M., Emberson J. D., Habib S., Frontiere N., 2023, J. Cosmol. Astropart. Phys., 06, 003
Teyssier R., 2002, A&A, 385, 337
Tremaine S., Gunn J. E., 1979, Phys. Rev. Lett., 42, 407
Upadhye A., 2019, J. Cosmol. Astropart. Phys., 05, 041
Upadhye A., Biswas R., Pope A., Heitmann K., Habib S., Finkel H., Frontiere N., 2014, Phys. Rev. D, 89, 103515
Upadhye A., Kwan J., Pope A., Heitmann K., Habib S., Finkel H., Frontiere N., 2016, Phys. Rev. D, 93, 063515
Williams C. K., Rasmussen C. E., 2006, Gaussian processes for machine learning. Vol. 2, MIT press Cambridge, MA
Wong Y. Y., 2008, J. Cosmol. Astropart. Phys., 10, 035
Yoo J., 2014, Class. Quant. Grav., 31, 234001
Yoo J., Zaldarriaga M., 2014, Phys. Rev. D, 90, 023513
Yu H.-R. et al., 2017, Nature Astron., 1, 0143
Yu H.-R., Pen U.-L., Wang X., 2019, Phys. Rev. D, 99, 123532
Zhou S. et al., 2022, MNRAS, 512, 3319
Zhu H.-M., Castorina E., 2020, Phys. Rev. D, 101, 023525
Zhu H.-M., Pen U.-L., Chen X., Inman Derek an d Yu Y., 2014, Phys. Rev. Lett., 113, 131301
Zhu H.-M., Pen U.-L., Chen X., Inman D., 2016, Phys. Rev. Lett., 116, 141301

APPENDIX A: IMPLEMENTATION OF `Cosmic-Eν`

We implement the emulator described in Sections 2.1 and 4.1 by extracting optimized hyperparameters from `SEPIA`. Following Heitmann et al. (2009), we construct a deterministic emulator which uses the mean weights $\bar{W}_j^{(L)}(\vec{C})$ of equation (11) as the emulated weights.

The hyperparameters upon which $\bar{W}_j^{(L)}(\vec{C})$ depends may be extracted from `SEPIA`. For each decile L, and for a `SepiaData` object called `data` and a `SepiaModel` object called `model`, the hyperparameter means basis weights, and basis functions are stored within `SEPIA` as follows:

(i) $w_{jm}^{*(L)}$ in `model.nu.w`;
(ii) $\hat{\beta}_{j\ell}^{(L)}$ in `model.params.betaU.val`;
(iii) $\hat{\lambda}_{\mathrm{U},j}^{(L)}$ in `model.params.lamUz.val`;
(iv) $\hat{\lambda}_{\mathrm{W},j}^{(L)}$ in `model.params.lamWs.val`;
(v) $\mu_i^{*(L)}$ in `data.sim_data.orig_y_mean`;
(vi) $\sigma^{*(L)}$ in `data.sim_data.orig_y_sd`; and
(vii) $\phi_j^{(L)}(k_i, z_i)$ in the jth row, ith column of `data.sim_data.K`.

This paper has been typeset from a TEX/LATEX file prepared by the author.