

Extending evolution mapping to massive neutrinos with COMET

Andrea Pezzotta^{1,2,*} Alexander Eggemeier³ Giosuè Gambardella^{4,5} Lukas Finkbeiner¹
 Ariel G. Sánchez¹ Benjamin Camacho Quevedo^{4,5,6,7,8} Martin Crocce^{4,5} Nanoom Lee^{9,10}
 Gabriele Parimbelli^{4,7} and Román Scoccimarro⁹

¹*Max Planck Institute for Extraterrestrial Physics, Giessenbachstraße 1, 85748 Garching, Germany*

²*INAF—Osservatorio Astronomico di Brera, via Emilio Bianchi 46, 23807 Merate, Italy*

³*Universität Bonn, Argelander-Institut für Astronomie, Auf dem Hügel 71, 53121 Bonn, Germany*

⁴*Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain*

⁵*Institut d'Estudis Espacials de Catalunya (IEEC), Carrer Gran Capitá 2-4, 08034 Barcelona, Spain*

⁶*Institute for Fundamental Physics of the Universe, Via Beirut 2, 34151 Trieste, Italy*

⁷*Scuola Internazionale di Studi Superiori Avanzati, via Bonomea 265, 34136 Trieste, Italy*

⁸*INAF—Osservatorio Astronomico di Trieste, via Tiepolo 11, 34143 Trieste, Italy*

⁹*Center for Cosmology and Particle Physics, Department of Physics, New York University, New York, New York 10003, USA*

¹⁰*William H. Miller III Department of Physics and Astronomy, Johns Hopkins University, Baltimore, Maryland 21218, USA*



(Received 5 May 2025; accepted 20 June 2025; published 9 July 2025)

We introduce an extension of the evolution mapping framework to cosmological models that include massive neutrinos. The original evolution mapping framework exploits a degeneracy in the linear matter power spectrum when expressed in Mpc units, which compresses its dependence on cosmological parameters into those that affect its shape and a single extra parameter σ_{12} , defined as the rms linear variance in spheres of radius 12 Mpc. We show that by promoting the scalar amplitude of fluctuations, A_s , to a shape parameter, we can additionally describe the suppression due to massive neutrinos at any redshift to sub-0.01% accuracy across a wide range of masses and for different numbers of mass eigenstates. This methodology has been integrated into the public COMET package, enhancing its ability to emulate predictions of state-of-the-art perturbative models for galaxy clustering, such as the effective field theory (EFT) model. Additionally, the updated software now accommodates a broader cosmological parameter space for the emulator, enables the simultaneous generation of multiple predictions to reduce computation time, and incorporates analytic marginalization over nuisance parameters to expedite posterior estimation. Finally, we explore the impact of different infrared resummation techniques on galaxy power spectrum multipoles, demonstrating that any discrepancies can be mitigated by EFT counterterms without impacting the cosmological parameters.

DOI: 10.1103/physrevd.112.023520

I. INTRODUCTION

The advent of Stage IV experiments represents a remarkable milestone in the field of observational cosmology and the study of the large-scale structure (LSS) of the Universe. Observational projects such as the Dark Energy Spectroscopic Instrument (DESI) [1], *Euclid* [2,3], the

Nancy Grace Roman Space Telescope (Roman) [4], and the Rubin Observatory Legacy Survey of Space and Time LSST [5] will collect galaxy samples comprising tens of millions of sources. These datasets will thus enable the reconstruction of a 3D map of the observable Universe to an unprecedented level of detail.

Harnessing the full potential of these datasets requires the development of efficient tools capable of jointly analyzing multiple cosmological probes. By integrating different observables within a unified framework, it becomes possible to break parameter degeneracies and mitigate prior volume effects in parameter estimation [6–8], ultimately enhancing the robustness of cosmological constraints. At the same time, the need for fast and accurate methods to evaluate theoretical models is becoming increasingly

*Contact author: andrea.pezzotta@inaf.it

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

pressing, particularly at nonlinear scales, where computational costs would otherwise be prohibitive.

Over the last decade, the cosmological community has made significant progress in the development of emulators—computational tools designed to provide smart multidimensional interpolations of various statistics across the cosmological parameter space. The primary advantage of emulators is their ability to replicate different clustering statistics as measured from high-precision N -body simulations, while significantly reducing computational time. This approach has been widely applied to the matter power spectrum [9–13], as well as to other statistics beyond dark matter clustering, such as baryonic corrections [14] and the large-scale distribution of galaxies. For the latter, different strategies exist. One approach bypasses the need to emulate galaxy bias parameters and directly emulates the individual blocks of the standard one-loop perturbative expansion of the galaxy power spectrum [15,16]. Alternatively, it is possible to train emulators to reproduce the final galaxy power spectrum on N -body simulations that reproduce different halo occupation distribution (HOD) galaxy samples, either incorporating the HOD parameters in the parameter space of the emulator [10,17] or analytically modeling the galaxy-halo distribution *a posteriori* [18].

An alternative approach focuses on constructing emulators for theory codes that rely on a perturbative expansion of the density and velocity fields [19–25]. This method preserves the accuracy of perturbative methods on mildly nonlinear scales while accelerating the computation of nonlinear models by several orders of magnitude. However, the development of accurate emulators faces significant challenges due to the growing parameter space that must be explored when considering extensions to the standard cosmological model. This issue is particularly important in light of the ambitious goals of Stage IV surveys, which aim to place percent-level constraints on key cosmological parameters, such as the equation of state of dark energy, the curvature of the Universe, and the total mass of neutrinos.

In this context, evolution mapping [26] offers an innovative way of addressing the high dimensionality of the cosmological parameter spaces. This approach captures the complex interplay of various cosmological parameters, referred to as evolution parameters, by capturing their combined effect on the matter density field through a single derived quantity, σ_{12} [27]. This parameter represents the variance of the linear density field on a fixed scale in units of Mpc, as opposed to the more conventional choice of using h^{-1} Mpc units. By parametrizing the evolution of the matter density field in terms of σ_{12} , one can show that different combinations of evolution parameters produce equivalent density fields, provided they share the same value of σ_{12} . Furthermore, for cosmologies with the same power spectrum shape, the evolution of the density field can be consistently rescaled across different sets of evolution parameters by expressing it in terms of σ_{12} .

The degeneracy experienced by evolution parameters is exact at the linear level, but Ref. [26] demonstrated that this relation holds approximately even in the highly nonlinear regime, on scales around $k \sim 1 \text{ Mpc}^{-1}$ and for redshifts approaching $z = 0$. More recently, Ref. [28] extended evolution mapping by showing that it can also effectively describe the time evolution of the velocity field, quantified through the auto-velocity-velocity and cross-density-velocity power spectra ($P_{\theta\theta}$ and $P_{\delta\theta}$), obtaining consistent accuracy levels. In addition, Refs. [21,29] applied the evolution mapping framework to construct COMET, a theory emulator for redshift-space distortion (RSD) models that also includes the effective field theory (EFT) of large-scale structure (e.g., Refs. [30–32]) and the recently developed framework based on the velocity-difference-generating function (VDG $_{\infty}$), effectively propagating the degeneracy to also describe the clustering of biased tracers in redshift space, and reducing the input parameter space over which the training has to be performed. The emulator, publicly available through PyPI¹ and GitLab,² has already been exploited in the context of prelaunch analyses within the Euclid Consortium for efficiently sampling posterior distributions in galaxy clustering likelihood analyses, as shown in recent works (e.g., Ref. [33]).

However, to meet the stringent precision requirements of Stage IV surveys, a critical limitation of the evolution mapping framework must be addressed: the proper incorporation of massive neutrinos. Unlike cold dark matter or baryons, neutrinos possess significant thermal velocities, which allow them to free-stream out of overdense regions, suppressing the growth of structure on scales below their free-streaming length (see, e.g., Ref. [34], for a comprehensive review on the subject). This suppression introduces a scale-dependent growth factor, fundamentally complicating the rescaling of the amplitude of density fluctuations across redshift. Moreover, the magnitude of this effect depends on the total neutrino mass and the number of massive species, adding further complexity to the problem. As such, the standard evolution mapping framework, which assumes a scale-independent growth factor, is inadequate for cosmologies including massive neutrinos. Extending the framework to incorporate these effects is not just an improvement but an essential step toward ensuring its applicability to the broader parameter spaces probed by next-generation cosmological surveys.

This work proposes a novel mapping applicable to the suppression factor induced by massive neutrinos on the matter power spectrum, showing that it relies on the relative growth factor across cosmologies compared to initial density fluctuation amplitudes. With this method, cosmologies can still be mapped using σ_{12} as a time variable, maintaining consistency with the original evolution

¹<https://pypi.org/project/comet-emu>

²<https://gitlab.com/aegge/comet-emu>

mapping approach. We implemented this extension in COMET, achieving evaluation speeds for nonlinear models with massive neutrinos that are 3 orders of magnitude faster than direct theory code evaluations [from $\mathcal{O}(10\text{s})$ to $\mathcal{O}(10\text{ ms})$ on a single CPU]. The new version software, labeled v2, has been trained to reproduce theory predictions for the two aforementioned RSD frameworks, EFT and VDG_∞ , whose relative performances have been recently benchmarked in Ref. [29]. Additional features include a broadening of the input cosmological parameter space used to train the emulator, the simultaneous generation of multiple theory predictions for reduced evaluation time, the analytical marginalization of linear parameters that facilitates a faster posterior estimation, and finally the use of a more robust algorithm to account for the impact of infrared modes on the amplitude of the baryon acoustic peaks in the galaxy power spectrum. With this article, we officially make this second version of COMET publicly available.

This paper is structured in the following way: In Sec. II, we first revisit the original formulation of evolution mapping and then proceed to extend it to cosmologies including massive neutrinos. In Sec. III, we provide an overview of the EFT and VDG_∞ models, focusing on the different treatment of RSD at separations approaching the nonlinear scale. In Sec. IV, we present the design and extra features of the new version of COMET, while in Sec. V, we discuss the validation and performance of the code and examine how different techniques that resum infrared modes to account for the damping of baryon acoustic oscillations (BAOs) affect the recovery of cosmological parameters. Finally, in Sec. VI, we give our conclusions.

II. EVOLUTION MAPPING

In the following subsections, we summarize the baseline formulation of evolution mapping in the context of a massive-neutrino-free cosmology, focusing on the mapping between different cosmologies at the level of the linear power spectrum [26]. We then proceed to describe the minimal setup required to extend this framework to include the effects of massive neutrinos, which is one of the main results of this paper.

A. Original formulation

The main aspect behind evolution mapping relies on the distinction between two different subsets of cosmological parameters, *shape* and *evolution*, respectively. The shape parameters,

$$\Theta_s = \left\{ \omega_b, \omega_c, \omega_\gamma, n_s, \frac{dn_s}{d \log k}, \dots \right\}, \quad (1)$$

directly impact the shape of the matter transfer function $T(k)$ and the postinflationary primordial power spectrum $P_{\mathcal{R}}(k)$. They are comprised of the physical density

parameters, $\omega_X \equiv \Omega_X h^2$, of different components, such as baryons (ω_b), cold dark matter (ω_c), and radiation (ω_γ), in addition to the scalar primordial index (n_s), its running with the wave mode k , and potentially other parameters from extensions to the standard cosmological model. On the other side, the evolution parameters,

$$\Theta_e = \{\omega_K, \omega_{\text{DE}}, w_0, w_a, A_s, \dots\}, \quad (2)$$

only affect the overall amplitude of the matter power spectrum at any later (than the matter-radiation equality) epoch. They include the physical density parameters of curvature (ω_K) and dark energy (ω_{DE}), together with the parameters describing its equation of state³ (w_0, w_a), the scalar primordial amplitude (A_s), and similarly other parameters that do not modify the shape of the power spectrum. In this context, the Hubble expansion rate [$h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$] is a mixture of shape and evolution parameters, since from the first Friedmann equation it follows that

$$h^2 = \sum_i \omega_i, \quad (3)$$

with the sum running over all the components of the Universe. However, for fixed shape parameters, any change in h effectively leads to a modification in the total dark energy density ω_{DE} . Likewise, the fractional density parameters ($\Omega_i \equiv \omega_i/h^2$) are also a combination of shape and evolution parameters.

For fixed shape parameters, the effect of changing the evolution parameters can then be described using a single quantity—or, in different words, evolution parameters are degenerate with each other in terms of the actual amplitude of linear density fluctuations.⁴ Traditionally, the amplitude of the linear power spectrum is described in terms of the rms fluctuations of the density field within some fixed physical scale R , as

³This is based on the well-known CLP parametrization [35,36], for which the time evolution of the dark energy equation of state is described as

$$w(z) = w_0 + w_a \frac{z}{1+z}.$$

However, any other standard parametrization may be added as a further set of evolution parameters.

⁴This statement is valid only over a restricted range of separations, since on large enough scales, strong deviations from a cosmological constant ($w = -1$) lead to $\mathcal{O}(1\%)$ broadband differences at $k \lesssim 0.01 \text{ Mpc}^{-1}$, in the same way that curvature modifies the amplitude of the linear power spectrum on super-horizon scales. However, for the range of scales that are typically employed in galaxy clustering analyses ($k > 0.01 \text{ Mpc}^{-1}$), this does not represent a limitation to the applicability of evolution mapping.

$$\sigma^2(R) = \frac{1}{2\pi^2} \int k^2 P_{\text{lin}}(k) W_{\text{TH}}^2(kR) dk, \quad (4)$$

where

$$W_{\text{TH}}(kR) = 3 \left[\frac{\sin(kR)}{(kR)^3} - \frac{\cos(kR)}{(kR)^2} \right] \quad (5)$$

is the Fourier transform of a top-hat window function of size R . In this context, selecting a reference value of R in units of Mpc is crucial, rather than the conventional choice of $R = 8h^{-1}$ Mpc. This is necessary because varying the reference value of h would otherwise affect k and $P(k)$ differently and thus break the degeneracy of the evolution parameters [27]. For a *Planck*-like value of $h \approx 0.67$, the previous scale translates approximately to 12 Mpc, making $\sigma_{12} \equiv \sigma(12 \text{ Mpc})$ a reasonable choice to describe the degeneracy experienced by evolution parameters.

With these premises, it is then possible to derive an expression for the linear matter power spectrum in terms of the shape and evolution parameters, as

$$P_{\text{lin}}(k|\Theta_s, \Theta_e, z) = P_{\text{lin}}(k|\Theta_s, \sigma_{12}(\Theta_s, \Theta_e, z)). \quad (6)$$

This means that, once expressed in Mpc units, the dependence of the power spectrum on the whole set of cosmological parameters can be reduced to the dependence on the shape parameters and on σ_{12} , which effectively acts as a proxy for the amplitude of the power spectrum. The dependence on the redshift z can also be absorbed into σ_{12} , since time evolution only induces an overall constant increase of power via the growth factor $D(z)$.

This behavior is illustrated in Fig. 1, where we show the time evolution of the linear power spectrum for three different cosmologies sharing the same set of shape parameters, but with varying evolution parameters. In all cases, we obtain predictions using the public Boltzmann solver CAMB,⁵ and we have verified that these predictions are consistent with those obtained from the CLASS⁶ Boltzmann solver at the level of $\mathcal{O}(0.01\% - 0.1\%)$. Considering configurations with the same redshift ($z = 1$ in this case) leads to different amplitudes for the three spectra, a consequence of the different growth factors $D(z)$ of the three cosmologies. However, by considering individual redshifts resulting in the same value of σ_{12} , it is possible to obtain consistent amplitudes across the entire range of separations considered, as can be observed in the bottom panel.

While the degeneracy between evolution parameters is perfect at linear level, Ref. [26] showed that this can be also propagated to the nonlinear regime, with small residuals appearing at the level of the one-halo term that are sourced




FIG. 1. Evolution of the linear power spectrum for cosmologies with the same set of shape parameters Θ_s . In the top panel, the power spectra are evaluated at a fixed reference redshift ($z = 1$), showing differences in amplitude. In the bottom panel, they are evaluated at redshifts chosen to match the same value of σ_{12} , resulting in identical amplitudes.

by different structure growth histories.⁷ This can be corrected by properly accounting for the different structure formation histories, quantified in terms of the suppression factor $g(z) = (1+z)D(z)$, leading to $\mathcal{O}(0.1\%)$ accuracies up to $k = 1.5 \text{ Mpc}^{-1}$ for a broad range of σ_{12} values.⁸ More recently, Ref. [28] extended the validity of this formalism from the density δ to the velocity field θ , also showing that the velocity power spectra—auto- $P_{\theta\theta}$ and cross- $P_{\delta\theta}$ —can

⁷The residual discrepancy mentioned above is, in any case, irrelevant within the framework of perturbative models of galaxy clustering under the Einstein–de Sitter (EdS) approximation. This is because any remaining dependence on the chosen cosmology vanishes, and the model becomes fully determined by P_{lin} , which can be perfectly mapped with evolution mapping.

⁸The accuracy of the mapping slightly decreases for higher values of σ_{12} , or alternatively, at lower redshifts, due to the stronger influence of nonlinear evolution in this regime. Nevertheless, even when considering snapshots with $\sigma_{12} \sim 0.83$, different cosmologies are in agreement at better than 1%, and most of the discrepancies are concentrated in the strongly nonlinear regime ($k \gtrsim 1 \text{ Mpc}^{-1}$).

⁵<https://github.com/cmbant/CAMB>

⁶https://github.com/lesgourg/class_public

be mapped from one cosmology to the other in terms of σ_{12} with similar levels of accuracy.

B. Extension to massive neutrinos

The framework described in the previous section can be used to characterize the evolution of the density and velocity fields across vastly different cosmologies. However, a key requirement for its validity is the absence of a massive neutrino component contributing to the total energy budget of the Universe. The fundamental issue is that, for a fixed number of effective neutrino species N_{eff} , massive neutrinos—or more specifically, the total neutrino mass M_ν —do not fit neatly into either the shape or evolution parameter categories, but instead act as a mixture of the two. This is because massive neutrinos modify the shape of the linear power spectrum through a scale-dependent growth factor $D(z, k)$, where time dependence can no longer be separated from spatial dependence due to the damping effect caused by neutrino free streaming, which evolves over time.

This can be easily observed in Fig. 2, where we plot the time evolution of the suppression factor,

$$S(k) \equiv \frac{P(k)}{P_\star(k)}, \quad (7)$$

for three different values of the total neutrino mass M_ν . In this case as well, predictions have been obtained using CAMB, but we have explicitly verified that consistent results (within 0.1%) can be obtained with CLASS, provided the default fluid approximation for massive neutrinos is disabled.⁹

In Eq. (7), we use P to denote the linear power spectrum of a cosmology that includes massive neutrinos, and $P_\star$ to represent the corresponding spectrum of an equivalent cosmology with the same total dark matter density, but composed entirely of cold dark matter. The latter is defined by $\Theta_{s,\star}$, where $\omega_{c,\star} = \omega_c + \omega_\nu$ and $\omega_{\nu,\star} = 0$. In this comparison, the cosmology with massive neutrinos assumes one massive species and two massless ones, and it adopts an effective number of relativistic species $N_{\text{ur}} = 2.029$, so that the total energy density in relativistic species corresponds to an effective neutrino number $N_{\text{eff}} = 3.044$. Two main effects can be seen in Fig. 2. First, since the total matter density is conserved, the presence of massive neutrinos delays the time of matter-radiation equality, as indicated by the small excess of power at $k \sim 0.005 \text{ Mpc}^{-1}$. Second, the growth of linear perturbations is suppressed on scales much smaller than the free-




FIG. 2. Evolution of the neutrino suppression factor as a function of the different total neutrino mass M_ν , redshift z , and wave number k . Different variations of the same color correspond to different values of z , with the darkest shade corresponding to the case at $z = 0$, for which the neutrino damping is strongest.

streaming length of any massive neutrino species [34]. For an individual species with mass m_ν , the comoving free-streaming scale can be approximated as

$$\lambda_s(z) = 2\pi \sqrt{\frac{2}{3}} \frac{c_\nu(z)}{H(z)} \simeq 7.7 \sqrt{\frac{\omega_m(z)}{\omega_{m,0}(1+z)}} \left(\frac{1 \text{ eV}}{m_\nu} \right) \text{ Mpc}, \quad (8)$$

where $\omega_m(z)$ and $\omega_{m,0}$ are the physical matter density parameter at the considered redshift and at present time, respectively; $H(z)$ is the Hubble expansion rate; and c_ν is the neutrino thermal velocity. As time passes, the amplitude of the damping gets larger, as massive neutrinos suppress the growth of density fluctuations below the free-streaming scale. This implies that the evolution mapping principle breaks down in this scenario, as redshift can no longer be used to track the evolution of the density field in a parameter-independent way. This is especially noticeable when exploring cosmologies with a sufficiently large value of M_ν , like $M_\nu \sim 0.5\text{--}1 \text{ eV}$, which have not been completely discarded by galaxy clustering constraints from Stage III experiments [37,38].

To address the residual effect at scales smaller than the neutrino free-streaming scale within the context of evolution mapping, we can rearrange Eq. (7) and express the linear power spectrum for a generic cosmology (including massive neutrinos) as the product of the neutrino-free power spectrum and the suppression factor,

$$P(k) = P_\star(k)S(k). \quad (9)$$

As discussed in Sec. II A, $P_\star$ can be mapped between different cosmologies solely based on the shape parameters

⁹Specifically, in this comparison we adopted `ncdm_fluid_approximation = 3` and `l_max_ncdm = 40`, which correspond to a full Boltzmann hierarchy integration for massive neutrinos up to multipole $\ell = 40$, without invoking any fluid approximation.




FIG. 3. Mapping between the suppression factors $S(k)$ across different cosmologies including massive neutrinos. The left, middle, and right columns correspond to different values of the total neutrino mass: $M_\nu = 0.1, 0.5, 1$ eV, respectively. *Top*: suppression factor for the cosmologies C0, C1, and C2, at eight different values of the ratio $\sigma_{12,*}/\sqrt{A_s}$, specifically the ones corresponding to the C0 cosmology at redshift from $z = 0.5$ (largest suppression) to $z = 4$ (smallest suppression) with a step of $\Delta z = 0.5$. *Bottom*: fractional percent difference between the suppression factors of cosmology C1 and C2 with respect to C0.

$\Theta_{s,*}$ and the total amplitude $\sigma_{12,*}$. The key question is then whether it is feasible to develop a new mapping for the suppression factor S that would allow evolution mapping to effectively apply to P as well.

Before proceeding, a few considerations are necessary. First, the value of σ_{12} used to track the evolution of the density field should be the one corresponding to the cosmology $\Theta_{s,*}$ —i.e., $\sigma_{12,*}$, regardless of the actual neutrino mass. This ensures that the large-scale limit of P_* matches that of the cosmology including massive neutrinos, as can be seen in Fig. 2. Second, the value of $\sigma_{12,*}$ alone does not fully determine the severity of neutrino damping on scales $\lambda \lesssim \lambda_s$. This becomes evident when comparing two cosmologies for which P_* can be mapped from one to the other through evolution mapping (and thus they share the same shape parameters, including M_ν , and the same $\sigma_{12,*}$) and observing that the amplitude of the damping generally differs between them. This indicates that a subset of evolution parameters carries additional information beyond $\sigma_{12,*}$. Considering the suppression factor $S(k)$, we note that it is independent from the scalar amplitude A_s , since the latter has the same effect on both P and P_* . This suggests that a more natural parameter to describe the evolution of the suppression factor is the ratio $\sigma_{12,*}/\sqrt{A_s}$, which removes the explicit dependence of $\sigma_{12,*}$ on A_s . This choice can be further justified by the fact that fixing $\sigma_{12,*}$ while varying A_s effectively requires adjusting the redshift to match the large-scale amplitude. Consequently, a lower (higher) value of A_s corresponds to a lower (higher) redshift, which in turn leads to stronger (weaker) suppression, as neutrino free-streaming has had more (less) time to dampen structure formation.

To validate the preceding arguments, Fig. 3 illustrates the suppression factor for three distinct cosmologies (C0, C1, C2), with their respective parameters detailed in Table I. As with the original formulation excluding massive neutrinos,

the shape parameters, this time also including the total neutrino mass, are kept constant to maintain the shape of the power spectrum. For the evolution parameters, we adopt a *Planck*-like Λ CDM cosmology [39] as the reference model (C0), and we consider two additional cases (C1 and C2) that feature significantly different values of A_s , and the dark energy equation of state w and energy density ω_{DE} . By evaluating the rescaling of the suppression factor across three different values of M_ν , this approach effectively explores a total of nine distinct cosmologies.

The top panels of Fig. 3 show the suppression factors of the cosmologies (C0, C1, C2) for the three selected values of M_ν and for eight different values of the ratio $\sigma_{12,*}/\sqrt{A_s}$.

TABLE I. Cosmological parameters assumed for the reference cosmology C0 and two other configurations (C1, C2). In order, different rows show the value of the baryon ω_b and the total matter ω_m density parameter, the scalar index n_s , the dark energy density ω_{DE} , the expansion rate h [this is a derived parameter following Eq. (3)], the scalar amplitude A_s [normalized by the value assumed in C0], the equation-of-state parameter (w) for dark energy, and the values of the total neutrino mass (M_ν). Given the requirements that the total matter density stay the same, the use of three different values for M_ν implies the effective use of nine different cosmologies.

	C0	C1	C2
ω_b	0.0219961	0.0219961	0.0219961
ω_m	0.1431991	0.1431991	0.1431991
n_s	0.97	0.97	0.97
ω_{DE}	0.3057009	0.6668009	0.1068009
h	0.67	0.90	0.50
$(A_s/2.1)[10^{-9}]$	1.00	3.00	0.33
w	-1.0	-1.1	-0.9
M_ν [eV]	(0.1, 0.5, 1.0)	(0.1, 0.5, 1.0)	(0.1, 0.5, 1.0)

For all of these values, we find an exquisite match among the suppression factors of the three different cosmologies, despite the extremely different evolution parameters that they feature. To better appreciate the level of agreement, in the bottom panels we show the fractional percentage difference of C1 and C2 with respect to the reference case C0. For all the considered values of M_ν , the residuals are well below the 0.01% threshold, with the only exception represented by the case at $M_\nu = 1$ eV and for the largest value of $\sigma_{12,\star}/\sqrt{A_s}$. Looking into the details, we can observe a trend of increasing residuals for larger values of $\sigma_{12,\star}/\sqrt{A_s}$, as highlighted by the lighter color shades. This suggests that, while the overall agreement is strong, there may be subtle effects from variations in the remaining evolution parameters that are not fully captured by the ratio $\sigma_{12,\star}/\sqrt{A_s}$. Nevertheless, the amplitude of the residual is still extremely small, and it never exceeds the order of magnitude of $\mathcal{O}(0.01\%)$.

In conclusion, these observations on the suppression factor $S(k)$ enable us to extend the original formulation of evolution mapping, as given by Eq. (6), to scenarios that include massive neutrinos. The updated equation can then be expressed as

$$P_{\text{lin}}(k|\Theta_s, \Theta_e, z) = P_{\text{lin}}(k|\Theta_s, A_s, \sigma_{12,\star}(\Theta_{s,\star}, \Theta_e, z)). \quad (10)$$

In this new formulation, M_ν is now incorporated into the shape parameters Θ_s , the scalar amplitude A_s must be explicitly specified and can no longer be treated as part of the remaining evolution parameters Θ_e , and the time variable σ_{12} must be calculated based on the cosmology with $\Theta_{s,\star}$ rather than Θ_s . The requirement to treat A_s and $\sigma_{12,\star}$ independently arises because, while $S(k)$ can be mapped across different cosmologies using the ratio $\sigma_{12,\star}/\sqrt{A_s}$, $P_\star(k)$ requires a specific value of $\sigma_{12,\star}$ to correctly recover the large-scale amplitude of density fluctuations, effectively making necessary the specification of both parameters.

C. Effect on the total matter power spectrum

In the previous subsection, the validation of the mapping across different cosmologies has been conducted under the implicit assumption that the power spectrum quantifies the amplitude of perturbations of cold dark matter (CDM) and baryons. This assumption stems from the conventional approach of using a two-fluid model [40–42] to describe the matter density field, δ_{cb} and δ_ν . Within this framework, it is typically assumed that the clustering of biased tracers can be fully characterized by δ_{cb} alone [43–45]. As a result, the relevant matter power spectrum to start from in the modeling of the Galactic power spectrum is $P_{\text{cb}} \propto \langle \delta_{\text{cb}}^2 \rangle$.

Even if the previous assumption is well justified in the limit of a small neutrino mass and for typical survey volumes from Stage III experiments [38,46], its applicability to future




FIG. 4. Same as in Fig. 3, but displaying the comparison between the suppression factor of the c + b power spectrum (faint) and that of the total matter component—i.e., c + b + ν (bright). In this case, we only show the configuration with the largest value of the total neutrino mass—i.e., $M_\nu = 1$ eV—and only three values for the ratio $\sigma_{12,\star}/\sqrt{A_s}$.

surveys must be carefully assessed. For this reason, we are interested in understanding how the evolution mapping of the suppression factor behaves when replacing P_{cb} with the total matter power spectrum $P_{\text{cb}\nu}$. In terms of the transfer function, this can be written as

$$T_{\text{cb}\nu}(k) = f_{\text{cb}}T_{\text{cb}}(k) + f_\nu T_\nu(k), \quad (11)$$

where $f_{\text{cb}} \equiv (\omega_c + \omega_b)/\omega_m$ and $f_\nu \equiv \omega_\nu/\omega_m$ are the fractions of CDM plus baryons, and neutrinos, respectively. In Fig. 4, we show a comparison between the suppression factor in the two cases, still making use of the same set of cosmologies considered in Fig. 3. Since we only consider the case with $M_\nu = 1$ eV, we find a significantly stronger ($\sim 10\%$) suppression when considering fluctuations of the total matter density field. Nevertheless, the mapping across different sets of evolution parameters still follows degeneracies for constant values of the ratio $\sigma_{12,\star}/\sqrt{A_s}$, as highlighted in the bottom panel. We observe a larger mismatch at the transition scale, where the damping factor starts to flatten out, at approximately $k \sim 0.5 \text{ Mpc}^{-1}$, but still find the different curves to be consistent at better than 0.02% for all the considered cases. We therefore conclude that evolution mapping can also be applied when considering the total matter power spectrum $P_{\text{cb}\nu}$, even if for the remainder of this article we will still reference only the CDM-plus-baryons component.

D. Effect on the number of massive neutrinos

One of the key insights in neutrino cosmology is that the total mass M_ν of the different neutrino flavors is the




FIG. 5. Same as in Fig. 3, but displaying the comparison between the suppression factor when considering one massive neutrino flavor (faint) or three individual neutrino species with the same total neutrino mass (bright). In this case, we only show the configuration with the largest value of the total neutrino mass—i.e., $M_\nu = 1$ eV—and only three values for the ratio $\sigma_{12,\star}/\sqrt{A_s}$.

primary parameter influencing the suppression of power below the free-streaming scale. However, subtle yet significant differences can arise in the matter power spectrum depending on the individual mass fractions $m_{\nu,i}$ of each neutrino eigenstate. This occurs because, although the total mass M_ν determines the composition and expansion rate of the Universe, variations in the mass distribution among the three neutrino flavors can cause shifts in the epoch at which they transition from a relativistic to a nonrelativistic behavior. In the previous section, we assumed a single massive eigenstate encompassing the total mass M_ν —i.e., $N_\nu = 1$. Here, we aim to investigate whether the mapping of the suppression factor is sensitive to this assumption or not.

In Fig. 5, we compare the suppression factor as we did in Fig. 4, but this time we examine two cases with the same total mass $M_\nu = 1$ eV and different numbers of massive neutrino species: $N_\nu = 1$ and $N_\nu = 3$. In the second case, the mass is equally distributed among the three different eigenstates, with $m_{\nu,i} \simeq 0.33$ eV for $i = 1, 2, 3$. In both configurations, we set the number of massless neutrino species such that the total effective number of neutrino species remains fixed at $N_{\text{eff}} = 3.044$. This corresponds to $N_{\text{ur}} = 2.029$ for $N_\nu = 1$ and $N_{\text{ur}} = 0$ for $N_\nu = 3$. As anticipated, the latter configuration exhibits a larger free-streaming scale, since the individual neutrino masses are smaller and thus remain relativistic for a longer period. However, this also means that they begin to suppress the growth of structure later in cosmic history and for a shorter duration, resulting in a milder suppression of small-scale

power compared to the configuration with a single, more massive species.

When assessing the accuracy of the evolution mapping, we observe that the three different cosmologies (C0, C1, C2) can still be effectively mapped onto each other, maintaining the high level of agreement previously seen in the $N_\nu = 1$ configuration. The most notable difference is that discrepancies now begin to appear at lower k values, although they eventually converge with the original differences as we approach the small-scale plateau of the suppression factor. Even in this scenario, the mapping accuracy remains better than $\mathcal{O}(0.02\%)$.

We explicitly tested that these results remain consistent with more realistic neutrino mass hierarchies—normal ($m_{\nu,1} \approx m_{\nu,2} < m_{\nu,3}$) and inverted ($m_{\nu,1} < m_{\nu,2} \approx m_{\nu,3}$)—in addition to the degenerate case shown in Fig. 5. For these tests, we assumed a minimum total neutrino mass of $M_\nu = 0.06$ eV, consistent with current neutrino oscillation experiments, and considered variations of this value up to $M_\nu = 0.2$ eV. The individual mass splits are chosen according to the assumed hierarchy: for a normal hierarchy, we adopt mass fractions of (0.1, 0.1, 0.8), while for an inverted hierarchy, we use (0.1, 0.45, 0.45). In all cases, we find that the residuals in the mapping do not exceed the 0.005% limit, essentially resulting in the same matter power spectrum. These findings underscore the robustness of the suppression factor mapping across different neutrino mass distributions, confirming that, even with variations in the number of massive neutrino species and different neutrino hierarchies, the mapping achieves remarkable precision, ensuring its applicability across a wide range of cosmological models.

III. THEORY MODELS

To introduce some of the theoretical quantities that are going to be relevant later in Sec. IV, in this section we provide a brief overview of the two models implemented in the current version of COMET, the EFT and VDG $_\infty$ models (see Ref. [29] for a more detailed discussion of their formulation and parametrization).

In both cases, we model the overdensity of biased tracers, δ_g , using a standard galaxy bias expansion up to third order in the matter density δ [47–50]:

$$\delta_g = b_1 \delta + b_{\nabla^2} \nabla^2 \delta + \varepsilon_g + \frac{b_2}{2} \delta^2 + \gamma_2 \mathcal{G}_2 + \gamma_{21} \mathcal{G}_{21} + \dots \quad (12)$$

This expansion incorporates both local (δ, δ^2) and nonlocal ($\mathcal{G}_2, \mathcal{G}_{21}$) operators [50,51], parametrized by four bias parameters: the linear bias b_1 , the quadratic bias b_2 , and the leading- and next-to-leading-order tidal biases, γ_2 and γ_{21} . Additionally, the model includes a higher-derivative contribution $\nabla^2 \delta$, modeling spatial nonlocalities in the

formation of galaxies, and a stochastic correction ε_g to account for the effects of short-wavelength fluctuations on galaxy clustering at larger scales. As already discussed in Sec. II C, in cosmological models that contain massive neutrinos, we define all operators in the bias relation with respect to fluctuations in the *cold* matter density—i.e., the combination of the CDM and baryon components. In contrast to a definition based on the total matter density, which includes the neutrino fluctuations, this prescription has been observed to yield a halo mass function that is close to universal and a linear bias parameter that is only weakly scale dependent in simulations [43,44,52]. The scale-dependent growth in massive neutrino cosmologies generally introduces a scale dependence in the galaxy bias at the few-percent level [53,54]. Since this scale-dependent signal would be difficult to disentangle from the nonlinear and higher-derivative contributions to the bias relation, we do not consider it further in our modeling.

In redshift space, under the assumptions of mass conservation in the mapping from real- to redshift-space coordinates and the distant-observer approximation, the anisotropic galaxy power spectrum for a wave mode k and an orientation to the line of sight θ [with $\mu \equiv \cos(\theta)$] can be written as [55]

$$P_{\text{gg}}^s(k, \mu) = \int \frac{d\mathbf{r}}{(2\pi)^3} e^{ik \cdot \mathbf{r}} \langle e^{i\lambda \Delta u_{\parallel}} [\delta_g(\mathbf{x}) + f \nabla_{\parallel} u_{\parallel}(\mathbf{x})] \times [\delta_g(\mathbf{x}') + f \nabla_{\parallel} u_{\parallel}(\mathbf{x}')] \rangle, \quad (13)$$

where $\mathbf{r} \equiv \mathbf{x} - \mathbf{x}'$, $\Delta u_{\parallel} \equiv u_{\parallel}(\mathbf{x}) - u_{\parallel}(\mathbf{x}')$, and $\lambda = -ifk\mu$. The quantity $\mathbf{u}$ represents the normalized velocity field, defined as $\mathbf{u} \equiv -\mathbf{v}/[faH(a)]$, such that $\nabla \cdot \mathbf{u} = \delta$ in linear perturbation theory. In addition, a and $H(a)$ denote the scale factor and the Hubble parameter, respectively, the subscript “ $\parallel$ ” refers to components along the line of sight, and the growth rate f is given by the logarithmic derivative of the linear growth factor $D(a)$,

$$f(a) \equiv \frac{d \log D(a)}{d \log a}. \quad (14)$$

In cosmologies including massive neutrinos, we assume that galaxy velocities, like the galaxy density, are sourced by the CDM and baryon components, consistent with findings from simulations [56,57]. The growth rate is accordingly computed for the CDM-plus-baryon fluid, but we additionally neglect any scale dependence induced by neutrinos and account for their suppression effect only through the linear power spectrum.¹⁰ This approximation has been found to be accurate for sufficiently small neutrino

masses, $M_{\nu} \lesssim 0.1$ eV [58,59], but an exact treatment may be required when exploring larger M_{ν} values, as allowed by analyses that consider galaxy clustering data without the addition of cosmic microwave background (CMB) information [60].

The expression for the redshift-space power spectrum in Eq. (13) reveals two distinct regimes of distortions. The terms in square brackets describe large-scale power enhancement due to the coherent infall of galaxies into overdense regions [61,62], while the velocity-difference-generating function (VDG), defined as

$$W(\lambda, \mathbf{r}) \equiv \langle e^{i\lambda \Delta u_{\parallel}} \rangle, \quad (15)$$

suppresses small-scale power via the finger-of-God (FoG) effect [63], caused by velocity dispersion inside virialized structures. In practice, these regimes are tightly intertwined, with no sharp transition scale.

In Ref. [64], it was shown that the expectation value in Eq. (13) can be factorized, with the VDG being a prefactor that controls the small-scale damping of the power spectrum multipoles, and a second factor that depends on cross-correlations with the (large-scale) density perturbations, which can be expanded in perturbation theory up to, e.g., one loop [64–66]. A more refined version of this approach is represented by the VDG_{∞} model (see Ref. [29] for recent detailed comparison against simulations), which resums quadratic nonlinearities and, in the large-separation limit, $\mathbf{r} \rightarrow \infty$, results in the following damping factor:

$$W_{\infty}(\lambda) = \frac{1}{\sqrt{1 - \lambda^2 a_{\text{vir}}^2}} \exp\left(\frac{\lambda^2 \sigma_v^2}{1 - \lambda^2 a_{\text{vir}}^2}\right). \quad (16)$$

Here, a_{vir} is a free parameter that accounts for nonlinear contributions to velocity dispersion and controls the extended tails of the pairwise velocity distribution (i.e., its kurtosis—see, e.g., Ref. [67]), and σ_v is the bulk velocity in linear theory,

$$\sigma_v^2 \equiv \frac{1}{6\pi^2} \int P_{\text{lin}}(k) dk, \quad (17)$$

where $P_{\text{lin}}(k)$ is the linear matter power spectrum. Equation (16) has been used already in the analysis of galaxy redshift surveys—see, e.g., Refs. [37,66,68,69].

On the other side, the second class of RSD models describe the velocity-difference-generating function by expanding it in a power series, aligning with the recently developed EFT framework [30,70]. This approach introduces effective speed-of-sound parameters to account for the small-scale RSD damping, but they are also meant to describe the impact of shell crossing and the breakdown of perturbation theory for ultraviolet modes. At leading order, the nonvirial power spectrum includes terms of the form $k^2 P_{\text{lin}}(k)$ and even powers of μ ,

¹⁰We make use of the total matter density (CDM + b + ν) to predict the large-scale limit of the growth rate, as this quantity correctly reproduces the large-scale limit of the ratios $P_{\delta\theta}/P_{\delta\delta}$ and $P_{\theta\theta}/P_{\delta\theta}$.

$$P_{\text{ctr}}^{\text{LO}}(k, \mu) = -2k^2 P_{\text{lin}}(k) \sum_{i=0,2,4} c_i \mathcal{L}_i(\mu), \quad (18)$$

where $\mathcal{L}_\ell$ is the ℓ -th-order Legendre polynomial, and c_0, c_2, c_4 parametrize the leading RSD contributions to the multipoles $\ell = 0, 2, 4$. To account for stronger nonlinearities in the real-to-redshift mapping, next-to-leading-order corrections are often included,

$$P_{\text{ctr}}^{\text{NLO}}(k, \mu) = c_{\text{nlo}} k^4 \mu^4 f^4 (b_1 + f\mu^2)^2 P_{\text{lin}}(k), \quad (19)$$

where c_{nlo} is an additional nuisance parameter to be marginalized over in likelihood analyses.

The final expressions for the models discussed above, incorporating biased tracers, can be summarized as follows. In the EFT framework, the redshift-space galaxy power spectrum is written as

$$P_{\text{gg}}^{\text{EFT}}(k, \mu) = P_{\text{gg}}^{\text{tree}}(k, \mu) + P_{\text{gg}}^{1\text{-loop}}(k, \mu) + P_{\text{gg}}^{\text{ctr}}(k, \mu) + P_{\text{gg}}^{\text{stoch}}(k, \mu), \quad (20)$$

where tree-level and one-loop terms arise from the coupling of the density and velocity fields, $P_{\text{gg}}^{\text{ctr}}$ includes the counter-terms from Eqs. (18) and (19), and $P_{\text{gg}}^{\text{stoch}}$ accounts for stochastic contributions. The latter, modeled perturbatively and truncated at next-to-leading order, can be written as

$$P_{\text{gg}}^{\text{stoch}}(k, \mu) = \frac{1}{\bar{n}} [N_0^P + k^2 (N_{20}^P + N_{22}^P \mathcal{L}_2(\mu))], \quad (21)$$

where $\bar{n}$ is the mean number density of the sample, and $N_0^P, N_{20}^P, N_{22}^P$ are free parameters describing deviations from Poisson noise.

Using the same formalism, the VDG_∞ model can be written as

$$P_{\text{gg}}^{\text{VDG}_\infty}(k, \mu) = W_\infty(\lambda) [P_{\text{gg}}^{\text{tree}}(k, \mu) + P_{\text{gg}}^{1\text{-loop}}(k, \mu) - \Delta P(k, \mu) + P_{\text{ctr}}^{\text{LO}}(k, \mu)] + P_{\text{gg}}^{\text{stoch}}(k, \mu). \quad (22)$$

A few important clarifications are needed here. First, in the VDG_∞ model, certain additional contributions that appear at one loop in the EFT model are not required. This is because they are implicitly accounted for by treating the velocity-difference-generating function analytically, without expanding it in power series as in the EFT model. Consequently, the one-loop term in the VDG_∞ model can be recast in terms of the EFT one by subtracting a correction of the form [29]

$$\Delta P(k, \mu) = \int \frac{d\mathbf{r}}{(2\pi)^3} e^{i\mathbf{k}\cdot\mathbf{r}} \frac{\lambda^2}{2} \langle \Delta u_{\parallel}^2 \rangle [\delta(\mathbf{x}) + f\nabla_{\parallel} u_{\parallel}(\mathbf{x})] \times [\delta(\mathbf{x}') + f\nabla_{\parallel} u_{\parallel}(\mathbf{x}')]. \quad (23)$$

Second, the RSD damping described by the analytical function W_∞ does not apply to the stochastic term $P_{\text{gg}}^{\text{stoch}}$ [29], such that stochastic corrections are equivalent across the two RSD models. Finally, we notice that despite the small-scale RSD signal being modeled through W_∞ , the VDG_∞ model still features the leading-order counterterms present in the EFT formalism. This choice is motivated to correctly capture residual effects that are equally shared with the EFT model, such as higher derivatives in the galaxy bias expansion, velocity bias, and residual breaking of perturbation theory for ultraviolet modes.

A final common feature in both models is the smearing of the BAO features caused by large-scale bulk motions [71–75]. A common approach to deal with this [76–78] involves decomposing the power spectrum into smooth (P_{nw}) and wiggly (P_{w}) components:

$$P_{\text{lin}}(k) = P_{\text{nw}}(k) + P_{\text{w}}(k). \quad (24)$$

The BAO damping is then modeled by applying a damping factor to P_{w} , yielding the infrared (IR) resummed power spectrum at leading order,

$$P_{\text{lin,IR}}(k) = P_{\text{nw}}(k) + e^{-k^2 \Sigma^2} P_{\text{w}}(k), \quad (25)$$

where Σ^2 represents the variance of the large-scale displacement field, which, in redshift space, receives additional anisotropic contributions. To incorporate IR effects into the galaxy power spectrum, Eq. (25) can be used in place of P_{lin} in all the relevant terms at both leading and next-to-leading order.

IV. DESIGN OF THE EMULATOR

In this section, we proceed with the description of how COMET v2 extends the design originally implemented in the first public release of COMET [21]. The basic structure of the code has been maintained to allow for a fast and accurate prediction of the galaxy power spectrum multipoles $P_\ell(k, z)$ according to different modeling frameworks. In the next few paragraphs, we describe the main modeling choices of COMET v2, explicitly pointing out the extra features that have been added with respect to COMET.

A. Flowchart

Similarly to the original version, COMET v2 builds upon the concept of evolution mapping, now extended to account for the influence of relic neutrinos with nonzero mass, as detailed in Sec. II. Since the mapping between different evolution parameters can be applied without introducing significant inaccuracies, it remains valid for a broad class of nonlinear models that depend solely on the shape and amplitude of the linear power spectrum. This is typically the case for perturbative approaches, such as the EFT and

VDG_∞ models, where nonlinear corrections P_{loop} can be computed via integrals of the form

$$P_{\text{loop}}(\mathbf{k}) \propto \int_{\mathbf{q}} K(\mathbf{q}, \mathbf{k} - \mathbf{q}) P_{\text{lin}}(q) P_{\text{lin}}(|\mathbf{k} - \mathbf{q}|), \quad (26)$$

with $\mathbf{q}$ being a loop variable to be integrated over, and $K(\mathbf{q}_1, \mathbf{q}_2)$ representing a generic kernel describing the nonlinear coupling between modes $\mathbf{q}_1$ and $\mathbf{q}_2$. Assuming that the kernel K is independent of cosmology and that it can be accurately modeled under the EdS approximation (see, e.g., Refs. [79–82] for works aimed at improving over this ansatz), any mapping applied to P_{lin} will naturally propagate to P_{loop} . The only additional components required by RSD models are the growth rate f and the bulk velocity dispersion σ_v (specifically for the VDG_∞ model), both of which can still be derived in terms of P_{lin} .

For the implementation of the two RSD models developed in the COMET package, the impact of massive neutrinos is only considered at the linear level. This is done by including damping above the neutrino suppression scale in the linear matter power spectrum, while assuming that the functional form of the mode-coupling kernels from perturbation theory remains consistent with the massless neutrino scenario. Although there are studies that investigate the theoretical systematics introduced by this approximation (see, e.g., Refs. [40,52,83,84]), addressing these effects is not the focus of this paper and will be considered in a future version of the code.

The approach of COMET v2 remains consistent with the methodology employed in the previous version. To avoid training the emulator directly on the bias and RSD parameters (introducing 11 additional dimensions in the input parameter space of the emulator), we take the approach of emulating the individual terms of the one-loop expansion, and we reconstruct the final power spectrum *a posteriori*. Specifically, all contributions to the power spectrum multipoles, $P_{B,\ell}$, arising from the nonlinear expansion, as described in Eqs. (18), (19), and (12), are organized according to their distinct combinations of bias and counterterm coefficients $\mathcal{B}$. This categorization allows for the emulation of a reduced set of independent terms, which can subsequently be analytically recombined to construct a model for $P_\ell(k, z)$. To reduce the dynamical range of the training data, therefore improving the accuracy of the emulation, the raw $P_{B,\ell}$ are not directly fed into the training routines. Instead, they are first rescaled by the corresponding linear power spectrum, so that the emulator outputs the normalized quantity $\beta_{B,\ell} \equiv P_{B,\ell}/P_{\text{lin}}$.

Altogether, there are 19 individual terms per multipole, which are outlined in Table II as products of the model parameters described in Sec. III. The list does not explicitly mention any of the stochastic parameters from Eq. (21), because the terms associated with the stochastic field are not emulated; instead, they are generated analytically and

TABLE II. Complete list of coefficients $\mathcal{B}$ for the terms emulated by COMET. This is split into those terms that enter at leading (tree) or one-loop order, and further, that are necessary for the real- or redshift-space galaxy power spectrum (in the latter, all the terms are required). The b_1^2 contribution is present both at leading and next-to-leading order, because we emulate the two contributions separately, to allow us to retrieve the leading-order IR-resummed power spectrum. The 1 in the redshift-space one-loop section represents the term sourced by the coupling of the velocity field with itself (for which there are no free parameters in the models reproduced by COMET). The terms marked in magenta are necessary only for the EFT model.

	Real space	Redshift space
Tree level	b_1^2, c_0	$c_2, c_4,$ $b_1^2 c_{\text{nlo}}, b_1 c_{\text{nlo}}, c_{\text{nlo}}$
One loop	$b_1^2, b_1 b_2, b_1 \gamma_2, b_1 \gamma_{21},$ $b_2^2, b_2 \gamma_2, \gamma_2^2$	$b_1, 1, b_2, \gamma_2, \gamma_{21}$

summed to the anisotropic galaxy power spectrum *a posteriori*. Similarly, other analytical effects can be applied afterwards, such as the small-scale damping W_∞ , resulting in a reconstructed 2D spectrum given by

$$P_{\text{gg}}(k, \mu) = W_\infty(k, \mu, \sigma_v) \left[\sum_{\ell} \mathcal{L}_\ell(\mu) \sum_B \mathcal{B} P_{B,\ell}(k) \right] + P_{\text{gg}}^{\text{stoch}}(k, \mu), \quad (27)$$

where W_∞ corresponds to Eq. (16) for the VDG_∞ model or is equal to 1 for the EFT model.

The transversal and longitudinal components of the wave mode $\mathbf{k}$ are then distorted as

$$k'_\perp = q_\perp k_\perp, \quad (28)$$

$$k'_\parallel = q_\parallel k_\parallel, \quad (29)$$

to account for the impact of the fiducial cosmology assumed to convert redshifts into comoving distances. Here, $q_\perp$ and $q_\parallel$ are the Alcock-Paczynski (AP) parameters [85], defined as [86]

$$q_\perp(z) = \frac{D_M(z)}{D'_M(z)}, \quad (30)$$

$$q_\parallel(z) = \frac{D_H(z)}{D'_H(z)}, \quad (31)$$

where D_M and D_H are the comoving angular diameter and Hubble distances, respectively, and the prime indicates quantities in the fiducial cosmology.

The final step involves projecting the reconstructed anisotropic power spectrum $P(k, \mu)$ onto the Legendre polynomials $\mathcal{L}_\ell$ to derive the AP-distorted power spectrum multipoles,




FIG. 6. Flowchart illustrating the steps employed by COMET to derive the power spectrum multipoles $P_\ell(k, z)$ from the initial input parameters. These inputs include shape Θ_s , evolution Θ_e , and nuisance parameters Θ_{nuis} , as well as the fiducial cosmology Θ^{fid} used for computing AP corrections. The input and output elements are marked with blue and red boxes, respectively, while all the quantities that involve the evaluation of an emulator are denoted with an orange box. Gray boxes represent quantities that can be derived analytically, without the use of any emulator. The main difference between the emulators including massive neutrinos or not is that A_s is either a shape or evolution parameter. In addition, the evaluation of the linear power spectrum changes in the two configurations: in one case, the output of the P_{lin} emulator is rescaled *a posteriori* by the ratio of σ_{12} values as in Eqs. (36) and (37), while in the other, $\sigma_{12,*}$ is requested as input for the P_{lin} emulator itself.

$$P'_\ell(k') = \frac{2\ell + 1}{2q_\perp^2 q_\parallel} \int_{-1}^{+1} d\mu' \mathcal{L}_\ell(\mu') P_{\text{gg}}(k(k', \mu'), \mu(k', \mu')). \quad (32)$$

Given that the nonlinear models considered here receive contributions up to μ^8 from the coupling of the density and velocity fields, the reconstruction from Eq. (27) remains incomplete if the summation over ℓ only includes orders $\ell = 0, 2, 4$. This is further reinforced by the inclusion of IR resummation in the model, as described in Sec. V C, which leads to the appearance of nonzero even multipoles at all orders. To account for this effect, we include also the $\ell = 6$ multipole in the resummation. Since all the terms $\ell \geq 6$ display a small amplitude, we opt not to create a full-fledged emulator for it, as we have done for the $\ell = 0, 2, 4$ multipoles. Instead, we calculate a single realization of the corresponding tables using a *Planck* cosmology based on the TT, TE, EE + lowE + lensing constraints [39]. We then rescale each element $P_{B,6}^{\text{Planck}}$ according to the values of the amplitude σ_{12} and the growth rate f , as computed from the evolution parameters of the cosmology being considered. The rescaling is then realized according to the following formula:

$$P_{B,6} = \left(\frac{\sigma_{12}(\Theta_s, \Theta_e, z)}{\sigma_{12}(\Theta_s^{\text{Planck}}, \Theta_e^{\text{Planck}}, z = 1)} \right)^{2L} P_{B,6}^{\text{Planck}}, \quad (33)$$

where $L = 1$ for leading-order terms and $L = 2$ for next-to-leading-order terms, as indicated in Table II.

As demonstrated in Ref. [21], this approach effectively reduces systematic errors observed in the validation stage, particularly when applying AP corrections to the VDG_∞ model. Therefore, we continue to employ this method to ensure the accuracy and robustness of our model.

Figure 6 provides a schematic view of the workflow used by COMET to evaluate the galaxy power spectrum multipoles. The process begins with the input parameters (marked by blue circles) and concludes with the computation of the AP-corrected Legendre multipoles (red circles). The emulated quantities (shown in orange) are computed as described in Sec. IV B 2. All other quantities (gray circles) are computed analytically, as they are quick calculations and have negligible impact on the total evaluation time of the process.

B. Training and parameter priors

1. Gaussian processes

The latest version of COMET has been trained using a Gaussian process (GP) algorithm implemented via the public software GPy.¹¹ Gaussian processes (see, e.g., Ref. [87] for a review on the subject) are an unsupervised learning method designed to find an optimal interpolation scheme across a set of scattered n -dimensional points $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1,N}$. The accuracy of this process largely depends

¹¹<https://gpy.readthedocs.io/en/deploy/>

on the covariance function $K(\mathbf{x}_i, \mathbf{x}_j)$ between different input positions $\mathbf{x}_i$ and $\mathbf{x}_j$, which determines the relative weights assigned to the evaluation of the dependent variable $\mathbf{y}$ at a new position $\mathbf{x}$. Consistently with the training approach of COMET, we employ a covariance function that combines a radial basis function (RBF) kernel,

$$K_{\text{RBF}}(\mathbf{x}, \mathbf{x}') = \sigma_{\text{RBF}}^2 \exp \left[- \sum_{i=1}^d \frac{(\mathbf{x}_i - \mathbf{x}'_i)^2}{2l_{\text{RBF},i}^2} \right], \quad (34)$$

with a Matérn kernel of degree $\nu = 3/2$,

$$K_{\text{Mat}}(\mathbf{x}, \mathbf{x}') = \sigma_{\text{Mat}}^2 \left(1 + \sqrt{3} \sum_{i=1}^d \frac{(\mathbf{x}_i - \mathbf{x}'_i)^2}{l_{\text{Mat},i}^2} \right) \times \exp \left(-\sqrt{3} \sum_{i=1}^d \frac{(\mathbf{x}_i - \mathbf{x}'_i)^2}{l_{\text{Mat},i}^2} \right), \quad (35)$$

where d is the dimension of the input parameter space, while l and σ are hyperparameters characterizing the correlation length (of different input parameters) and amplitude, respectively, at different positions $\mathbf{x}$ and $\mathbf{x}'$. Due to incompatibilities with the latest available versions of NUMPY (≥ 1.24),¹² the optimal hyperparameters obtained from the training process have been reapplied to train new models using the SCIKIT-LEARN (SL) software,¹³ which maintains full compatibility with NUMPY, allowing for seamless integration. Due to the consistency in the internal structures of both packages, this procedure was implemented efficiently, yielding machine-level precision in the comparative outputs of the two emulators.

Future plans for the COMET software include transitioning the emulators from a GP-based architecture to a neural network. This change could reduce the evaluation time of the emulator without compromising its accuracy, leveraging the ability of neural networks to maintain performance regardless of the training set size.

2. List of emulators

To preserve the original accuracy level outlined in Ref. [21], the initial COMET emulators are retained as distinct entities. These emulators can be selected by users who require predictions that exclude the effects of massive neutrinos. However, considering new extra features (see Sec. IV C), such as an extended support in k , the extended input parameter space, and the introduction of a new IR resummation scheme (see Sec. V C) needed for the massive neutrino emulators, for consistency, we also decided to retrain the original configurations.

In COMET, we had two distinct groups of emulated quantities:

Shape: Emulating P_{lin} , σ_{12} , and σ_v as functions of Θ_s (with fixed evolution parameters Θ_e^{fixed} and $z = 1$).

Full: Emulating $\beta_{B,\ell}$ as a function of Θ_s , σ_{12} , and f .

Predicting P_{lin} solely as a function of Θ_s is feasible because, once the shape of the power spectrum is determined, the overall amplitude can be obtained using the predicted σ_{12} and then rescaling it by the ratio of the growth factors $D(z)$ and the scalar amplitudes A_s between the requested cosmology and the one with Θ_e^{fixed} and $z = 1$ as follows:

$$\begin{aligned} \sigma_{12}(\Theta_s, \Theta_e, z) &= \frac{D(\Theta_s, \Theta_e, z)}{D(\Theta_s, \Theta_e^{\text{fixed}}, z = 1)} \left(\frac{A_s}{A_s^{\text{fixed}}} \right)^{1/2} \sigma_{12}(\Theta_s, \Theta_e^{\text{fixed}}, z = 1), \end{aligned} \quad (36)$$

$$\begin{aligned} P_{\text{lin}}(k|\Theta_s, \Theta_e, z) &= \left[\frac{\sigma_{12}(\Theta_s, \Theta_e, z)}{\sigma_{12}(\Theta_s, \Theta_e^{\text{fixed}}, z = 1)} \right]^2 P_{\text{lin}}(k|\Theta_s, \Theta_e^{\text{fixed}}, z = 1). \end{aligned} \quad (37)$$

A similar argument can be used to obtain predictions for the bulk velocity dispersion σ_v .

With the inclusion of massive neutrinos, the previous approach becomes inadequate, as A_s can no longer be treated as an evolution parameter. In this scenario, we can distinguish three different sets of emulators:

Shape: Emulating $\sigma_{12,\star}$ as a function of Θ_s (with fixed evolution parameters Θ_e^{fixed} and $z = 1$).

Linear: Emulating P_{lin} and σ_v as a function of Θ_s , A_s , and $\sigma_{12,\star}$.

Full: Emulating $\beta_{B,\ell}$ as a function of Θ_s , A_s , $\sigma_{12,\star}$, and f .

The new mapping described in Sec. II relies on the $\sigma_{12,\star}$ computed from the equivalent cosmology with $\Theta_{s,\star}$. Thus, it remains feasible to predict $\sigma_{12,\star}$ using only the shape parameters, and then rescale its value as described in Eq. (36). On the contrary, both P_{lin} and σ_v [which directly depends on P_{lin} via Eq. (17)] must now be expressed in terms of the shape parameters Θ_s , which also include M_ν , the scalar amplitude A_s , and the large-scale amplitude $\sigma_{12,\star}$. This approach allows us to accurately predict both the neutrino-free power spectrum $P_\star$ (through $\sigma_{12,\star}$) and the suppression factor S (through the ratio $\sigma_{12,\star}/\sqrt{A_s}$). As a result, the total linear power spectrum can be directly written as in Eq. (10).

3. Parameter ranges

As detailed in the previous sections, the parameter space spanned by the emulators differs between the original version of COMET and the new one, which also incorporates massive neutrinos. The complete set of parameters and their

¹²<https://numpy.org/>

¹³<https://scikit-learn.org/>

TABLE III. Parameters and ranges covered by the emulators in COMET. Without massive neutrinos, they include the baryon (ω_b) and CDM (ω_c) density parameters, the scalar index (n_s), the amplitude $\sigma_{12,\star}$, and the growth rate (f). With massive neutrinos, it is required to specify also the total neutrino mass (M_ν) and the scalar amplitude (A_s).

Parameter	COMET		COMET v2	
	Min.	Max.	Min.	Max.
ω_b	0.02050	0.02415	0.01930	0.02535
ω_c	0.085	0.155	0.08	0.16
n_s	0.92	1.01	0.90	1.03
$\sigma_{12,\star}$	0.2	1.0	0.2	1.0
f	0.5	1.05	0.5	1.05
M_ν [eV]	...	...	0.0	1.0
$A_s[10^{-9}]$	...	...	1.0	3.5

ranges is listed in Table III, explicitly highlighting the differences between the two configurations.

In the previous version of COMET, the emulator parameter space included the physical baryon density (ω_b), the cold dark matter density (ω_c), the scalar spectral index (n_s), the amplitude $\sigma_{12,\star}$, and the growth rate f (the latter is not required in real space). In the updated version, the parameter ranges have been expanded: for ω_b and n_s , by almost a factor of 2 and 1.5, respectively, and slightly less so for ω_c . These ranges now correspond to 33, 20, and 15 times the 1σ constraints reported in Ref. [39] for ω_c , ω_b , and n_s , respectively.

With the inclusion of massive neutrinos, the emulator parameter space additionally includes the total neutrino mass (M_ν) and the scalar amplitude (A_s).¹⁴ We have adopted sufficiently broad ranges for these two parameters, particularly for M_ν , which spans the interval $[0, 1]$ eV. As highlighted by recent analyses of the combined BOSS galaxy and eBOSS quasar samples [37], these limits are consistent with the 2σ confidence intervals obtained when considering extensions to Λ CDM that include an evolving dark energy equation of state and nonzero curvature.

This extended and broadened parameter space necessitates a larger training dataset to achieve a similar level of accuracy as the previous version of COMET. However, instead of sampling the full parameter space as done before, we now exploit correlations between parameter combinations, such as between f and $\sigma_{12,\star}$, and between ω_c and A_s . This approach allows us to exclude parts of the

parameter space that are unlikely to be sampled in any realistic scenario, thereby minimizing the required increase of the training dataset size.

In practice, we proceed as follows: we begin by generating Latin hypercube (LHC) samples in the range from 0 to 1 for each dimension. We then establish upper and lower boundaries for the growth rate as a function of $\sigma_{12,\star}$ —i.e., $f_{\max}(\sigma_{12,\star})$ and $f_{\min}(\sigma_{12,\star})$ —and perform inverse transform sampling to ensure a uniform sample density in the $f - \sigma_{12,\star}$ parameter plane. This involves computing the cumulative distribution function

$$\text{CDF}(\sigma_{12,\star}) = \mathcal{N} \int_{0.2}^{\sigma_{12,\star}} dx [f_{\max}(x) - f_{\min}(x)] \quad (38)$$

with the normalization factor $\mathcal{N}$, and applying its inverse to the LHC samples. We obtain the ω_c samples analogously, using the upper and lower limits of A_s as a function of ω_c . Meanwhile, all other LHC samples are scaled according to the boundaries specified in Table III. The upper and lower boundaries for the growth rate are motivated by allowing a 10% variation around $f_{\text{Planck}}(z)$ —calculated using the cosmological parameters of Ref. [39]—and a 20% variation around $\sigma_{12,\text{Planck}}(z)$ within the redshift range of $z = 0$ to 3. Similarly, we establish the boundaries on A_s as a function of ω_c by translating a 20% variation in σ_{12} at different ω_c values, within the range specified in Table III, into A_s . For reference, the precise relationships for all these boundaries are given in the Appendix.

In Fig. 7, we present a section of the $f - \sigma_{12,\star}$ parameter space, indicated by the gray-shaded area. Additionally, we display posterior distributions, obtained when fitting synthetic data vectors at four different redshifts within the range (1, 1.6). For these tests, we used Gaussian covariance matrices assuming redshift shells with an angular footprint of 2500 deg^2 and a radial width of $\Delta z = 0.2$ (thus roughly corresponding to, at most, a quarter of the expected volume explored by Stage IV surveys such as DESI and *Euclid*). Two cosmological configurations are considered: a standard Λ CDM model and a $w_0 w_a$ CDM model, each one assuming a single massive neutrino with fixed mass $M_\nu = 0.06 \text{ eV}$. In both cases, the derived values for f and $\sigma_{12,\star}$ fully lie within the ranges covered by the training set, demonstrating that the layout of the training dataset can be confidently restricted to the region shown in Fig. 7.

Despite the more efficient distribution of the samples, we choose to increase the size of the training dataset to 2000 points for the “full” emulators without massive neutrinos and to 3000 otherwise, in order to account for the larger parameter space. Since the dimensionality of the “shape” and “linear” emulators is lower, smaller training sets are sufficient, and in these cases we use 1000 and 2000 samples, respectively.

¹⁴For this new release of COMET, we train the emulator assuming one massive and two massless neutrino species, corresponding to a total effective number of neutrino species fixed at $N_{\text{eff}} = 3.044$. Since changes to either the number of massive species or the value of N_{eff} alter the shape of the matter power spectrum, any deviation from this configuration would require a dedicated training set and emulator. Such extensions will be considered in future releases of COMET.




FIG. 7. Distribution in the $f - \sigma_{12,*}$ parameter space of the training set for the emulators that incorporate massive neutrinos (gray shaded area). For reference, we also show the 68% and 95% confidence intervals obtained from a set of points sampled during the posterior distribution exploration of four synthetic data vectors, corresponding to redshifts $z = (1.0, 1.2, 1.4, 1.6)$, as explained in the main text. The blue and red colors denote the posterior distributions for the Λ CDM and w_0w_a CDM parameter spaces, respectively. For clarity, we only show a fraction of the total $f - \sigma_{12,*}$ plane, which also extends beyond the top-left corner of the plot, as specified in Table III.

C. Additional features and improvements

In the following subsection, we describe additional new features available in the latest version of COMET, beyond the capability to explore massive neutrino cosmologies, as discussed in the previous sections.

1. Extended support in wave mode k

COMET v2 has been trained across a wider range of wave modes k , ranging from $k_{\min} = 7 \times 10^{-4} \text{ Mpc}^{-1}$ to $k_{\max} = 0.5 \text{ Mpc}^{-1}$ (compared to the original maximum wave mode of 0.35 Mpc^{-1}). This extension should not be mistaken for an attempt to model galaxy clustering and RSD up to these scales using the perturbative methods outlined in Sec. III. Instead, this broader range is intended to accommodate modeling scenarios that require a larger wave mode coverage, without having to resort to extrapolation. These scenarios include the convolution with the survey window function, as needed when analyzing real data samples, but also the conversion of the power spectrum multipoles into multipoles of the two-point correlation function via the Hankel transform.

2. Batch evaluation

The new version of COMET now allows for batch evaluation, enabling multiple parameter combinations to be processed simultaneously to obtain the corresponding

power spectrum multipoles in a single call. This is particularly advantageous when using COMET to evaluate the joint likelihood across multiple redshift bins. By leveraging SL's built-in capabilities for batch processing, computation time can be significantly reduced. For example, a single evaluation of three power spectrum multipoles over $100k$ -values using the EFT model without massive neutrinos takes around 11 ms on a single 2.6 GHz core machine. However, when evaluating ten samples simultaneously, the average time is only about 30 ms, roughly offering a speedup of a factor of 3 compared to running the emulator ten separate times. This effect is shown later in Fig. 12, in which we compare the total evaluation time for predicting a set of power spectrum multipoles with different emulators (EFT and VDG_{∞} , both with and without massive neutrinos).

3. Analytical marginalization

In many fields that involve parameter inference from data, it is common practice to address uncertainties in nuisance parameter values by marginalizing over them. This marginalization can be performed numerically, although it may become computationally expensive when dealing with a large number of nuisance parameters.

A more efficient approach relies on the analytical marginalization of nuisance parameters. This method is widely applied in cosmology, with examples spanning from CMB analyses [88] to LSS studies [7,20,89]. Analytical marginalization can be performed under two conditions. First, the parameter to be marginalized must appear linearly in the theory model, meaning that only some of the bias and RSD parameters in Sec. III can be treated this way. Specifically, these are the cubic nonlocal bias γ_{21} , the three leading-order and higher-order counterterms ($c_0, c_2, c_4, c_{\text{nlo}}$), and the three stochastic parameters ($N_0^P, N_{20}^P, N_{22}^P$). Second, the prior distributions for these parameters must have a form that can be integrated analytically, and for the implementation in COMET, we assume Gaussian priors.

Figure 8 illustrates the performance of this method by comparing two chains executed with identical setups (data vectors, scale cuts, and priors), but with one chain explicitly sampling the linear nuisance parameters, and the other employing analytical marginalization. The resulting contours show excellent agreement, while the chain with analytical marginalization achieves convergence approximately 10 times faster than the sampled case.¹⁵

¹⁵The exact speed gain of analytical marginalization is hard to quantify, since it ultimately depends on the number of parameters treated this way. Nevertheless, this approach is highly advantageous, particularly for tasks such as model selection or calibration, where multiple posterior distributions need to be computed efficiently in a short amount of time.




FIG. 8. Comparison of the marginalized posteriors for the parameters (ω_c, h, A_s) obtained with COMET when either explicitly sampling all linear bias, RSD, and shot-noise parameters or marginalizing over them analytically. The analysis is based on synthetic data vectors designed to replicate the clustering of H α galaxies, covering four redshift bins in the range $1 < z < 2$, and assuming a Gaussian covariance corresponding to a sky fraction of $1/3$.

4. New IR resummation scheme

As anticipated at the end of Sec. III, large-scale bulk flows induce a smearing of the BAO wiggles, an effect that is commonly addressed by the IR resummation procedure. Some implementations of this procedure (e.g., Refs. [76–78]) are based on an explicit decomposition of the linear power spectrum into its smooth and wiggly components, P_{nw} and P_{w} , such that $P_{\text{lin}} = P_{\text{nw}} + P_{\text{w}}$. However, there is no unique way to isolate the wiggly part, which makes the result of the IR resummation procedure somewhat dependent on the specific method adopted.

In the first release of COMET [21], we made use of the *Gaussian filter* approach, proposed by Ref. [90], which uses the featureless Eisenstein-Hu matter power spectrum [91], $P_{\text{EH}}(k)$, as a template for the nonwiggly component. The non-negligible differences between the broadband shape of P_{lin} and P_{EH} can be accounted for by rescaling P_{EH} with the smoothed ratio of the two power spectra [24,33,90,92]—i.e.,

$$P_{\text{nw}}(k) = P_{\text{EH}}(k) \mathcal{F} \left[\frac{P_{\text{lin}}(k)}{P_{\text{EH}}(k)} \right]. \quad (39)$$

The filter function $\mathcal{F}$ is chosen to be a one-dimensional Gaussian, specifically

$$\mathcal{F}[f(k)] = \frac{\log_{10}(e)}{\sqrt{2\pi}\lambda} \int dq \frac{f(q)}{q} \exp \left[-\frac{1}{2\lambda^2} \log_{10}^2 \left(\frac{k}{q} \right) \right], \quad (40)$$




FIG. 9. Ratio between the linear power spectrum P_{lin} and the smooth power spectrum P_{nw} according to different modeling recipes, the Gaussian filtering technique, and DST, as specified in the legend. The reference cosmology assumed to generate the power spectra is a flat Λ CDM, one without massive neutrinos. The gray shaded band represents a 1% difference between the linear power spectrum and the no-wiggle template.

where $\lambda = 0.25h^{-1} \text{ Mpc}$ is a parameter that determines the width of the filter. While this algorithm is well suited for a Λ CDM cosmology, the inclusion of massive neutrinos leads to further differences in the broadband shape that can no longer be corrected for in this approach, particularly for large values of the neutrino mass, $M_\nu \sim 1 \text{ eV}$.

For this reason, the resummation of IR modes in COMET v2 is based on an alternative approach, the *discrete sine transform* (DST) algorithm, as originally proposed in Ref. [93] (see, e.g., Refs. [19,32], for applications of the same algorithm to independent codes). The algorithm involves performing a discrete sine transform of P_{lin} , followed by the removal of the bump associated with the BAO peak. This process is done by analyzing the second derivative of the sine transform, which identifies the location of the peak. Afterward, the gap created by removing the peak is filled using cubic spline interpolation. Finally, the modified function is transformed back into Fourier space, resulting in a power spectrum that no longer contains BAO oscillations.

Figure 9 presents a comparison of the two methods for a single Λ CDM cosmology without the inclusion of massive neutrinos. Even in this case, the aforementioned mismatch in the broadband shape of P_{EH} and P_{lin} is evident, but it is significantly reduced after application of the Gaussian filter. However, a residual deficiency in power is observed in P_{nw} on scales between 0.2 Mpc^{-1} and 1 Mpc^{-1} , which can reach an amplitude of up to 1%. In contrast, the profile of the DST P_{nw} aligns more closely with the actual mean position between the BAO wiggles, resulting in symmetric fluctuations of the wiggles around unity. In Sec. V C, we will study how




FIG. 10. Cumulative distribution for the accuracy of the new set of emulators generated for COMET v2. The x axis represents the fraction of the validation sample, while the y axis indicates the maximum difference between the emulator output and the validation set. This difference is measured either as the fractional percent deviation [top row, see Eq. (41)], or in terms of the Gaussian statistical noise from a sample with 10 times the volume of a redshift bin $z = (0.9, 1.1)$ and an angular footprint of 15000 deg^2 [bottom row, see Eq. (42)]. The four columns correspond to four different emulators, for EFT and VDG_∞ , each without and with massive neutrinos. The three different colors correspond to the distinct power spectrum multipoles, while the varying line styles represent the maximum-mode k_{max} used for the validation.

these differences in the IR resummation schemes propagate to the outcomes of a likelihood analysis.

V. VALIDATION AND PERFORMANCES

In this section, we finally present the validation of the new set of emulators, both in terms of their accuracy and their computational efficiency. Furthermore, we quantify the impact of the different IR resummation scheme adopted for COMET v2 in terms of the posterior distribution inferred from the analysis of multiple redshift bins with a combined volume consistent with the expected one from Stage IV surveys, such as DESI and *Euclid*.

A. Accuracy of the emulators

To perform the validation of the new emulators, we constructed a dedicated validation set for each case. This validation set is generated from a distribution of ΛCDM cosmological parameters, specifically $(\omega_b, \omega_c, n_s, M_\nu, h, A_s)$. The shape parameters in this distribution have the same coverage as the training set, as detailed in Table III, while the evolution parameters (h on its own or in combination with A_s for the massless and massive neutrino cases, respectively) are chosen from sufficiently large ranges. We then compute the corresponding power spectrum multipoles at $z = 1$ to construct the final validation set. In this process, the bias, RSD, and noise parameters are fixed to partially reproduce the clustering properties of the BOSS CMASS sample, as described in Ref. [29].

The accuracy estimation is performed using two different metrics. The first metric evaluates the distance $d(E, V)$ between the emulation and validation sets by calculating the maximum relative percent difference, defined as

$$d_{\%}(E, V) = 100 \max_i \left\{ \frac{2|P_{\text{emu}}(k_i) - P_{\text{val}}(k_i)|}{[P_{\text{emu}}(k_i) + P_{\text{val}}(k_i)]} \right\}, \quad (41)$$

where $P_{\text{emu}}(k_i)$ and $P_{\text{val}}(k_i)$ are the emulated and validation power spectrum multipoles at each wave number k_i , respectively. On the other side, to better understand the impact of the emulator inaccuracies in a real likelihood analysis, we define a new distance metric tied to a specific power spectrum multipole uncertainty σ_p , as

$$d_{\sigma}(E, V) = \max_i \left\{ \frac{|P_{\text{emu}}(k_i) - P_{\text{val}}(k_i)|}{\sigma_p(k_i)} \right\}. \quad (42)$$

For this exercise, similarly to the original validation of COMET, we consider the theoretical Gaussian uncertainty [94] from a volume enclosed in an angular footprint of 15000 deg^2 and redshift bins of width $\Delta z = 0.2$. To conduct more stringent validation tests, the volume is further multiplied by a factor of 10, enhancing the sensitivity of the analysis.

The results of the accuracy tests are shown in Fig. 10, where we present both distance metrics for the full set of emulators that we have trained. To ensure completeness, two sets of curves are shown, one calculated over the entire




FIG. 11. Distribution of the relative percentage accuracy for the monopole of the EFT model, based on the corresponding validation set. The plot displays the parameters $(\omega_c, M_\nu, h, A_s)$, while we marginalize over the remaining parameters (ω_b, n_s) . Data points are color-coded according to the maximum relative difference.

range of wave numbers used to train the emulators ($k_{\max} = 0.5 \text{ Mpc}^{-1}$) and another limited to a more realistic scale cut, considering a more sensible range of validity of these perturbative models¹⁶ ($k_{\max} = 0.3 \text{ Mpc}^{-1}$).

In both cases, we observe that the accuracy remains consistently better than 0.3% for all configurations. The relative accuracy tends to be higher when considering lower-order multipoles and scenarios without massive neutrinos. This trend is reversed when using the second distance metric, where the tighter error bars on the monopoles cause the emulated quantities to appear more discrepant from the validation set compared to higher-order multipoles. Even when including massive neutrinos, the majority of the sample remains consistent with the validation set, achieving an accuracy that is better than 0.3σ .

In addition to assessing the overall accuracy of the new emulators, we are particularly interested in verifying that the extended recipe for evolution mapping in the presence of massive neutrinos performs equally well regardless of the total neutrino mass M_ν . Figure 11 illustrates the accuracy distribution estimated from the validation set as a function of the position in cosmological parameter space. The full parameter space is marginalized over (ω_b, n_s) , so we only display the accuracy distribution for the remaining parameters. Since the validation set is constrained to samples where $\sigma_{12,\star}$ and f deviate by no more than 20% from the corresponding *Planck* predictions at the

¹⁶As pointed out in Ref. [29], the VDG_∞ model can be used to robustly describe the shape of the galaxy power spectrum multipoles up to this value of $k_{\max}$, when considering HOD mocks based on the BOSS CMASS and LOWZ galaxy samples.




FIG. 12. Evaluation time required to compute the three power spectrum multipoles $[\ell = (0, 2, 4)]$ across different emulator models, indicated in the legend. The time per evaluation is plotted against the number of batch evaluations, illustrating the efficiency gain when generating multiple samples simultaneously compared to running individual emulator calls for each sample separately.

same redshift, the multidimensional distribution does not encompass the entire parameter space, as evident in, e.g., the $\omega_c - A_s$ panel. As the figure shows, we do not observe any significant dependence of the emulator accuracy on neutrino mass. We only note a slight increase in the maximum differences (on average at the level of 0.1%) for low M_ν values, likely reflecting the proximity to the edge of the parameter space for which the emulators have been trained.

B. Computational performance

In terms of computational efficiency, the new set of emulators aligns closely with the performance reported in Ref. [21], even with the increased demands of the COMET v2 setup, which includes an extended k range and a broader parameter space. Figure 12 illustrates the scaling of the total evaluation time for predicting the galaxy power spectrum multipoles (considering 100k-bins) with respect to the number of batch evaluations introduced in Sec. IV C 2. A single model evaluation on a single-core CPU takes approximately 11 ms for both the EFT and VDG_∞ models without massive neutrinos. On the other hand, emulators with massive neutrinos require more time per evaluation due to the expanded training set and parameter space, with a single evaluation taking around 19 ms. As discussed in Sec. IV C 2, performing N_{eval} evaluations simultaneously reduces the overall evaluation time compared to executing N_{eval} individual calls. When considering more than two simultaneous evaluations, the slope of the $t_{\text{eval}}(N_{\text{eval}})$ relation becomes shallower, leading to the gains in computation time already mentioned in Sec. IV C 2.

Overall, we find that replacing the direct computation of the real non-linear model with a COMET evaluation improves computational speed by approximately 3 orders of magnitude. This acceleration is crucial for enabling efficient statistical inference of cosmological parameters in a likelihood analysis. Additionally, using evolution mapping to account for different evolution parameters reduces the dimensionality of the parameter space. In turn, this improves the accuracy of the emulator at a fixed evaluation time.

C. Impact of different IR resummation schemes

In this section, we examine the impact of changing methods for isolating the no-wiggle power spectrum P_{nw} , which is required for the implementation of IR resummation. As noted in Sec. IV C 4, COMET v2 replaces the original Gaussian filtering technique applied to the featureless template of Ref. [91] with the DST algorithm. In Fig. 9, we already compared these methods at the level of P_{nw} , whereas here we are interested in determining how the residual differences propagate to the galaxy power spectrum multipoles and, ultimately, to the posterior distribution of cosmological parameters.

The comparison in terms of the final model for the $P_\ell(k)$ is illustrated in Fig. 13, where we use a fixed set of cosmological and nuisance parameters to plot the power spectrum Legendre multipoles, employing both the Gaussian smoothing technique and the DST framework. The relative differences in the bottom panel reveal that discrepancies emerge on mildly nonlinear scales ($k \gtrsim 0.1 \text{ Mpc}^{-1}$) and are more pronounced for higher-order multipoles. The residuals range from 1% for the monopole P_0 to over 5% for the hexadecapole P_4 . This level of discrepancy has also been reported in previous studies (e.g., Ref. [24]), raising the question of whether such differences could ultimately result in inconsistent cosmological constraints in a realistic likelihood analysis.

To address this issue, we generate a synthetic data vector consisting of the monopole, quadrupole, and hexadecapole at $z = 1$ obtained from the EFT model with the Gaussian filter approach. We then analyze this synthetic data vector using the EFT model combined with either the Gaussian filter or the DST approach. For this test, we adopt a Gaussian covariance matrix based on the predictions of Ref. [94] (a conservative choice for this test) along with a volume and sample number density that is representative for a Stage IV galaxy survey. Specifically, we use a volume corresponding to a single redshift bin of size $\Delta_z = 0.2$ over a sky area of 15000 deg^2 and a number density based on current forecasts for the spectroscopic H α sample [33,95,96] that will be targeted by *Euclid* [2,3]. In the likelihood analysis, we sample three cosmological parameters (ω_c, h, A_s) and six nuisance parameters ($b_1, b_2, c_0, c_2, c_4, N_0^P$), and their posteriors are shown in Fig. 14.




FIG. 13. *Top*: comparison of the power spectrum multipoles within the EFT framework, using either the Gaussian filter or the DST methods to account for IR resummation effects. This analysis uses bias and RSD parameters derived from fits to synthetic mocks of H α galaxies, and it assumes a power spectrum amplitude of $\sigma_{12} = 0.7$. *Bottom*: percentage differences between the two approaches, with gray shaded areas highlighting the 1%, 3%, 5%, and 7% intervals.

Despite the residual differences between the two IR resummation methods, the figure demonstrates that the impact on cosmological parameters is minimal, with only a slight but negligible shift in the ω_c posterior. In contrast, more significant shifts are observed for some of the nuisance parameters, particularly the counterterms c_2 and c_4 . This is expected, given that they contribute almost exclusively to the quadrupole P_2 and the hexadecapole P_4 , respectively, which—as shown in Fig. 13—are the multipoles most affected by the different IR resummation methods.¹⁷ Other parameters, such as b_2 , also exhibit slight deviations from the fiducial values used to generate the data vectors, although these deviations remain within 1σ of the expected values. This indicates that the residual differences introduced by the different IR resummation schemes can be effectively absorbed by the nuisance parameters, without significantly impacting the cosmological parameters.

¹⁷We note that this statement may not hold when considering a parametrization of the RSD counterterms different from the one used in COMET. In our approach, each of the three leading-order terms primarily influences the amplitude of the corresponding power spectrum multipole. However, with an alternative parametrization—such as one using powers of μ instead of the Legendre polynomials $\mathcal{L}_\ell(\mu)$ —some of the observed discrepancies in c_2 and c_4 could potentially leak into c_0 .




FIG. 14. Marginalized posterior distribution of the cosmological parameters (ω_c, h, A_s) and the nuisance parameters ($b_1, b_2, c_0, c_2, c_4, N_{P,0}^P$) for the EFT modeling framework. The synthetic data vectors are generated at $z = 1$ using the Gaussian filter technique, applied to the smooth power spectrum from Ref. [91]. The two distinct marginalized contours represent cases where the theoretical model used in the likelihood is either the Gaussian filter or the DST method. In both scenarios, a Gaussian covariance matrix is assumed, derived from a redshift shell with width $\Delta z = 0.2$ and an angular footprint of 15000 deg^2 , with the fit extending up to $k_{\text{max}} = 0.3h \text{ Mpc}^{-1}$.

VI. CONCLUSIONS

This work extends the evolution mapping framework described in Ref. [26] to cosmological models that include massive neutrinos. In its original formulation, evolution mapping describes the evolution of the matter power spectrum only in terms of the shape parameters and a single amplitude parameter, σ_{12} . This can then be used as a new time variable to map different cosmologies with the same shape parameters into each other. This approach breaks down in the presence of massive neutrinos due to the suppression of power on scales below their typical free-streaming scale, since the strength of the damping is not uniquely determined in terms of the total neutrino mass M_ν and σ_{12} . To restore the mapping across different

cosmologies with massive neutrinos, we developed a more general formulation which also accounts for the mapping of the neutrino suppression factor $S(k)$. This is achieved in terms of the combination $\sigma_{12,\star}/\sqrt{A_s}$, where $\sigma_{12,\star}$ is evaluated for a cosmology with the same total matter density of the one under consideration, but composed solely of cold dark matter and baryons.

With these minor adjustments, it remains possible to achieve an almost exact mapping (to within $\sim 0.01\%$ accuracy) between cosmologies with different evolution parameters—such as expansion rate, curvature, dark energy, and redshift—and for a broad range of total neutrino masses (at least up to $M_\nu = 1 \text{ eV}$). We confirmed that this holds true when considering both the power spectrum of cold dark matter and baryons, as well as the

total matter density field, which also includes fluctuations in the neutrino density field. Furthermore, this approach is equally effective when varying the number of massive neutrino species, as we verified by explicitly testing configurations with $N_\nu = 1$ and $N_\nu = 3$.

As a follow-up to this result, we applied the extended evolution mapping framework to develop a new set of emulators of the galaxy power spectrum for cosmologies including massive neutrinos within the public COMET package [21]. This involved expanding the input parameter space to include both M_ν and A_s with sufficiently broad ranges, in order to encompass a wide variety of cosmological models. Specifically, we sampled M_ν in the range $[0, 1]$ eV and A_s in the range $[1.0, 3.5] \times 10^{-9}$. Additionally, we extended the prior ranges for the shape parameters already present in COMET, increasing them by factors of 1.65, 1.15, and 1.45 for ω_b , ω_c , and n_s , respectively. Given that the accuracy and computational performance of the emulators are affected by the two extra dimensions of the parameter space, we decided to retain a newly trained version of the emulators without massive neutrinos as separate tools. Users can select between the standard or massive neutrino versions by setting a dedicated flag, depending on the requirements of their analysis.

The newly created COMET emulator has been trained under the assumption of one massive and two massless neutrino species, with a total effective number of species fixed at $N_{\text{eff}} = 3.044$. We note that modifications to either the number of massive species or the value of N_{eff} would impact the shape of the matter power spectrum and thus require a new dedicated training set. Future releases of COMET will explore these directions to support a broader range of cosmological scenarios.

Consistently with the latest version of the software, COMET v2 allows users to obtain predictions using either the EFT or VDG $_\infty$ model (plus the real-space perturbative model), ensuring flexibility in the choice of theoretical framework. Similarly to the choice of including massive neutrinos or not, users can specify their preferred RSD model through a flag when initially creating an instance of the emulator. Beyond this selection, the code operates almost identically for both models.

Apart from the expanded parameter space, the new version of COMET introduces several additional features compared to its predecessor. For example, it offers an extended support for the wave mode k , as the new emulators have been trained up to $k_{\text{max}} = 0.5 \text{ Mpc}^{-1}$, an increase from the original limit of $k_{\text{max}} = 0.35 \text{ Mpc}^{-1}$. Most notably, the new version allows for simultaneous predictions of multiple Legendre multipoles with a single call to the emulator. This improvement, coupled with the transition to a different Gaussian process software, has reduced the total evaluation time by approximately 50%. Such efficiency gains can offer significant advantages, particularly when combining galaxy clustering constraints

from multiple redshift bins or from different datasets into a single posterior distribution. The process of sampling the posterior distribution over a given parameter space can be further accelerated through the implementation of analytical marginalization. This approach reduces the number of sampled bias, RSD, and noise parameters to only 3–4 per spectroscopic bin, depending on the chosen RSD model. This feature can be activated by the user when utilizing the built-in routine to calculate the χ^2 statistic between the theoretical model and user-provided data vectors.

Finally, the new sets of emulators include a different IR resummation scheme, compared to the original version of COMET. This transition was motivated by the poor performances of the Gaussian filtering technique applied to the featureless transfer function of Ref. [91] with cosmologies including massive neutrinos. Specifically, using this approach, residuals in the wiggle/no-wiggle split are observed at mildly nonlinear scales, potentially affecting the construction of the final galaxy power spectrum model. The new technique introduced in COMET v2 is based on a discrete sine transform, as also used in other public software packages (i.e., Ref. [19]). This approach successfully eliminates residuals, even for cosmologies that include massive neutrinos. We quantify the impact of different IR resummation schemes on the final galaxy power spectrum multipoles for a neutrino-free cosmology and find that higher-order multipoles are more sensitive to the chosen method, with residuals ranging from $\mathcal{O}(1\%)$ in the monopole up to $\mathcal{O}(10\%)$ in the hexadecapole. Despite these differences, we show that the inference of cosmological parameters is only minimally affected, with noticeable deviations confined to specific RSD parameters, such as the EFT counterterms. Since these parameters are typically marginalized in Bayesian analyses, we conclude that switching between IR resummation schemes can be considered safe, depending on user preference.

This paper is accompanied by the public release of COMET v2, which incorporates all the new features outlined in the main body of this article and summarized in the previous sections. The updated code can be installed via PyPI¹⁸ or by cloning the corresponding GIT repository.¹⁹

ACKNOWLEDGMENTS

We acknowledge useful discussions with Matteo Esposito, Daniel Farrow, Martha Lippich, and Agne Semenaite. A.P. acknowledges support from the INAF AstroFit fellowship (Iniziativa di Ricerca Fondamentale 2021–2025). A.E. is supported at the Argelander Institut für Astronomie by an Argelander Fellowship. G.G., B.C., M.C., and G.P. acknowledge support from the Spanish Ministerio de Ciencia, Innovación y Universidades, Project No. PID2021-128989NB. A.G.S. is supported by the

¹⁸<https://pypi.org/project/comet-emu/>

¹⁹<https://gitlab.com/aegge/comet-emu>

Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2094–390783311. N.L. was supported by the James Arthur Graduate Associate Fellowship from the Center for Cosmology and Particle Physics at New York University and the Horizon Fellowship from Johns Hopkins University. Our research made use of MATPLOTLIB, a Python library for publication-quality graphics [97].

DATA AVAILABILITY

The data that support the findings of this article are openly available [98].

APPENDIX: TRAINING SET BOUNDARIES FOR CORRELATED PARAMETERS

To reduce the number of samples needed for training the emulator, we exploit the fact that some of the parameters exhibit stronger correlations, specifically between f and $\sigma_{12,\star}$, as well as A_s and ω_c . Based on these correlations, we can define boundaries, $f_{\max}(\sigma_{12,\star})$ and $f_{\min}(\sigma_{12,\star})$, or $A_{s,\max}(\tilde{\omega}_c)$ and $A_{s,\min}(\tilde{\omega}_c)$, which indicate regions of the parameter space that are unlikely to be sampled in any realistic application of the emulator, allowing us to safely exclude them from the construction of our training dataset. Following the procedure outlined in Sec. IV B 3 and combining with the strict limits $0.48 \leq f \leq 1.05$ and $A_s \leq 3.5 \times 10^{-9}$, we determine the boundaries:

$$f_{\max}(\sigma_{12,\star}) \approx \begin{cases} 1.05, & \text{for } \sigma_{12,\star} \leq 0.44 \\ -0.73\sigma_{12,\star}^2 + 0.19\sigma_{12,\star} + 1.11, & \text{otherwise} \end{cases}, \quad (\text{A1})$$

$$f_{\min}(\sigma_{12,\star}) \approx \begin{cases} -1.69\sigma_{12,\star}^2 + 0.59\sigma_{12,\star} + 0.82, & \text{for } \sigma_{12,\star} \leq 0.65 \\ 0.48, & \text{otherwise} \end{cases}, \quad (\text{A2})$$

and

$$A_{s,\max}(\tilde{\omega}_c) \approx 10^{-9} \begin{cases} 3.5, & \text{for } \tilde{\omega}_c \leq 0.85 \\ 2.1\tilde{\omega}_c^2 - 5.5\tilde{\omega}_c + 4.9, & \text{otherwise} \end{cases}, \quad (\text{A3})$$

$$A_{s,\min}(\tilde{\omega}_c) \approx 10^{-9}(0.45\tilde{\omega}_c^2 - 1.6\tilde{\omega}_c + 1.8), \quad (\text{A4})$$

where $\tilde{\omega}_c \equiv \omega_c/0.1198$.

-
- [1] A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland *et al.* (DESI Collaboration), The DESI experiment part I: Science, targeting, and survey design, [arXiv:1611.00036](#).
 - [2] R. Laureijs, J. Amiaux, S. Arduini, J. L. Auguères, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet *et al.*, Euclid definition study report, [arXiv:1110.3193](#).
 - [3] Y. Mellier, Abdurro'uf, J. A. A. Barroso, A. Achúcarro, J. Adamek, R. Adam, G. E. Addison, N. Aghanim, M. Agüena *et al.* (Euclid Collaboration), Euclid. I. Overview of the Euclid mission, [Astron. Astrophys. **697**, A1 \(2025\)](#).
 - [4] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon *et al.*, Wide-field infrared survey telescope-astronomy focused telescope assets WFIRST-AFTA 2015 report, [arXiv:1503.03757](#).
 - [5] Z. Ivezic, T. Axelrod, W. N. Brandt, D. L. Burke, C. F. Claver, A. Connolly, K. H. Cook, P. Gee, D. K. Gilmore, S. H. Jacoby *et al.* (LSST Collaboration), Large synoptic survey telescope: From science drivers to reference design, [Serb. Astron. J. **176**, 1 \(2008\)](#).
 - [6] B. Hadzhiyska, K. Wolz, S. Azzoni, D. Alonso, C. García-García, J. Ruiz-Zapatero, and A. Slosar, Cosmology with 6 parameters in the Stage-IV era: Efficient marginalisation over nuisance parameters, [Open J. Astrophys. **6**, 23 \(2023\)](#).
 - [7] P. Carrilho, C. Moretti, and A. Pourtsidou, Cosmology with the EFTofLSS and BOSS: Dark energy constraints and a note on priors, [J. Cosmol. Astropart. Phys. **01** \(2023\) 028](#).
 - [8] T. Simon, P. Zhang, V. Poulin, and T. L. Smith, Consistency of effective field theory analyses of the BOSS power spectrum, [Phys. Rev. D **107**, 123530 \(2023\)](#).
 - [9] E. Lawrence, K. Heitmann, M. White, D. Higdon, C. Wagner, S. Habib, and B. Williams, The Coyote Universe. III. Simulation suite and precision emulator for the non-linear matter power spectrum, [Astrophys. J. **713**, 1322 \(2010\)](#).

- [10] Z. Zhai, J. L. Tinker, M. R. Becker, J. DeRose, Y.-Y. Mao, T. McClintock, S. McLaughlin, E. Rozo, and R. H. Wechsler, The aemulus project. III. Emulation of the galaxy correlation function, *Astrophys. J.* **874**, 95 (2019).
- [11] T. Nishimichi, M. Takada, R. Takahashi, K. Osato, M. Shirasaki, T. Oogi, H. Miyatake, M. Oguri, R. Murata, Y. Kobayashi, and N. Yoshida, Dark quest. I. Fast and accurate emulation of halo clustering statistics and its application to galaxy clustering, *Astrophys. J.* **884**, 29 (2019).
- [12] M. Knabenhans, J. Stadel, D. Potter, J. Dakin, S. Hannestad, T. Tram, S. Marelli, A. Schneider, R. Teyssier, P. Fosalba *et al.* (Euclid Collaboration), Euclid preparation: IX. EuclidEmulator2—power spectrum emulation with massive neutrinos and self-consistent dark energy perturbations, *Mon. Not. R. Astron. Soc.* **505**, 2840 (2021).
- [13] R. E. Angulo, M. Zennaro, S. Contreras, G. Aricò, M. Pellejero-Ibañez, and J. Stücker, The BACCO simulation project: Exploiting the full power of large-scale structure for cosmology, *Mon. Not. R. Astron. Soc.* **507**, 5869 (2021).
- [14] G. Aricò, R. E. Angulo, S. Contreras, L. Ondaro-Mallea, M. Pellejero-Ibañez, and M. Zennaro, The BACCO simulation project: A baryonification emulator with neural networks, *Mon. Not. R. Astron. Soc.* **506**, 4070 (2021).
- [15] M. Zennaro, R. E. Angulo, M. Pellejero-Ibañez, J. Stücker, S. Contreras, and G. Aricò, The BACCO simulation project: Biased tracers in real space, *Mon. Not. R. Astron. Soc.* **524**, 2407 (2023).
- [16] M. Pellejero Ibañez, R. E. Angulo, M. Zennaro, J. Stücker, S. Contreras, G. Aricò, and F. Maion, The bacco simulation project: Bacco hybrid Lagrangian bias expansion model in redshift space, *Mon. Not. R. Astron. Soc.* **520**, 3725 (2023).
- [17] J. Kwan, K. Heitmann, S. Habib, N. Padmanabhan, E. Lawrence, H. Finkel, N. Frontiere, and A. Pope, Cosmic emulation: Fast predictions for the galaxy power spectrum, *Astrophys. J.* **810**, 35 (2015).
- [18] Y. Kobayashi, T. Nishimichi, M. Takada, R. Takahashi, and K. Osato, Accurate emulator for the redshift-space power spectrum of dark matter halos and its application to galaxy power spectrum, *Phys. Rev. D* **102**, 063504 (2020).
- [19] A. Chudaykin, M. M. Ivanov, O. H. E. Philcox, and M. Simonović, Nonlinear perturbation theory extension of the Boltzmann code CLASS, *Phys. Rev. D* **102**, 063533 (2020).
- [20] G. D’Amico, L. Senatore, and P. Zhang, Limits on Λ CDM from the EFTofLSS with the PyBird code, *J. Cosmol. Astropart. Phys.* **01** (2021) 006.
- [21] A. Eggemeier, B. Camacho-Quevedo, A. Pezzotta, M. Crocce, R. Scoccimarro, and A. G. Sánchez, COMET: Clustering observables modelled by emulated perturbation theory, *Mon. Not. R. Astron. Soc.* **519**, 2962 (2022).
- [22] J. Donald-McCann, K. Koyama, and F. Beutler, Matryoshka II: Accelerating effective field theory analyses of the galaxy power spectrum, *Mon. Not. R. Astron. Soc.* **518**, 3106 (2023).
- [23] M. Maus, S. Chen, M. White, J. Aguilar, S. Ahlen, A. Aviles, S. Brieden, D. Brooks, T. Claybaugh, S. Cole *et al.*, An analysis of parameter compression and full-modeling techniques with Velocileptors for DESI 2024 and beyond, *J. Cosmol. Astropart. Phys.* **01** (2025) 138.
- [24] D. Linde, A. M. Dizgah, C. Radermacher, S. Casas, and J. Lesgourgues, CLASS-OneLoop: Accurate and unbiased inference from spectroscopic galaxy surveys, *J. Cosmol. Astropart. Phys.* **07** (2024) 068.
- [25] M. Bonici, G. D’Amico, J. Bel, and C. Carbone, Effort: A fast and differentiable emulator for the effective field theory of the large scale structure of the universe, *arXiv:2501.04639*.
- [26] A. G. Sánchez, A. N. Ruiz, J. G. Jara, and N. D. Padilla, Evolution mapping: A new approach to describe matter clustering in the non-linear regime, *Mon. Not. R. Astron. Soc.* **514**, 5673 (2022).
- [27] A. G. Sánchez, Arguments against using h^{-1} Mpc units in observational cosmology, *Phys. Rev. D* **102**, 123511 (2020).
- [28] M. Esposito, A. G. Sánchez, J. Bel, and A. N. Ruiz, Evolution mapping II: Describing statistics of the non-linear cosmic velocity field, *Mon. Not. R. Astron. Soc.* **534**, 3906 (2024).
- [29] A. Eggemeier, N. Lee, R. Scoccimarro, B. Camacho-Quevedo, A. Pezzotta, M. Crocce, and A. G. Sánchez, Boosting galaxy clustering analyses with non-perturbative modelling of redshift-space distortions, *arXiv:2501.18597*.
- [30] J. J. M. Carrasco, M. P. Hertzberg, and L. Senatore, The effective field theory of cosmological large scale structures, *J. High Energy Phys.* **09** (2012) 082.
- [31] A. Perko, L. Senatore, E. Jennings, and R. H. Wechsler, Biased tracers in redshift space in the EFT of large-scale structure, *arXiv:1610.09321*.
- [32] M. M. Ivanov, M. Simonović, and M. Zaldarriaga, Cosmological parameters from the BOSS galaxy power spectrum, *J. Cosmol. Astropart. Phys.* **05** (2020) 042.
- [33] A. Pezzotta, C. Moretti, M. Zennaro, A. Moradinezhad Dizgah, M. Crocce, E. Sefusatti, I. Ferrero, K. Pardede, A. Eggemeier, A. Barreira *et al.* (Euclid Collaboration), Euclid preparation—XLI. Galaxy power spectrum modelling in real space, *Astron. Astrophys.* **687**, A216 (2024).
- [34] J. Lesgourgues and S. Pastor, Massive neutrinos and cosmology, *Phys. Rep.* **429**, 307 (2006).
- [35] M. Chevalier and D. Polarski, Accelerating universes with scaling dark matter, *Int. J. Mod. Phys. D* **10**, 213 (2001).
- [36] E. V. Linder, Exploring the expansion history of the Universe, *Phys. Rev. Lett.* **90**, 091301 (2003).
- [37] A. Semenaite, A. G. Sánchez, A. Pezzotta, J. Hou, A. Eggemeier, M. Crocce, C. Zhao, J. R. Brownstein, G. Rossi, and D. P. Schneider, Beyond Λ CDM constraints from the full shape clustering measurements from BOSS and eBOSS, *Mon. Not. R. Astron. Soc.* **521**, 5013 (2023).
- [38] C. Moretti, M. Tsedrik, P. Carrilho, and A. Pourtsidou, Modified gravity and massive neutrinos: Constraints from the full shape analysis of boss galaxies and forecasts for stage iv surveys, *J. Cosmol. Astropart. Phys.* **12** (2023) 025.
- [39] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak *et al.* (Planck Collaboration), Planck 2018 results—VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [40] D. Blas, M. Garny, T. Konstandin, and J. Lesgourgues, Structure formation with massive neutrinos: going beyond linear theory, *J. Cosmol. Astropart. Phys.* **11** (2014) 039.
- [41] M. Zennaro, J. Bel, F. Villaescusa-Navarro, C. Carbone, E. Sefusatti, and L. Guzzo, Initial conditions for accurate

- N-body simulations of massive neutrino cosmologies, *Mon. Not. R. Astron. Soc.* **466**, 3244 (2017).
- [42] M. Zennaro, J. Bel, J. Dossett, C. Carbone, and L. Guzzo, Cosmological constraints from galaxy clustering in the presence of massive neutrinos, *Mon. Not. R. Astron. Soc.* **477**, 491 (2018).
- [43] F. Villaescusa-Navarro, F. Marulli, M. Viel, E. Branchini, E. Castorina, E. Sefusatti, and S. Saito, Cosmology with massive neutrinos I: Towards a realistic modeling of the relation between matter, haloes and galaxies, *J. Cosmol. Astropart. Phys.* **03** (2014) 011.
- [44] E. Castorina, E. Sefusatti, R. K. Sheth, F. Villaescusa-Navarro, and M. Viel, Cosmology with massive neutrinos II: On the universality of the halo mass function and bias, *J. Cosmol. Astropart. Phys.* **02** (2014) 049.
- [45] S. Vagnozzi, T. Brinckmann, M. Archidiacono, K. Freese, M. Gerbino, J. Lesgourgues, and T. Sprenger, Bias due to neutrinos must not uncorrect'd go, *J. Cosmol. Astropart. Phys.* **09** (2018) 001.
- [46] H. E. Noriega, A. Aviles, S. Fromenteau, and M. Vargas-Magaña, Fast computation of non-linear power spectrum in cosmologies with massive neutrinos, *J. Cosmol. Astropart. Phys.* **11** (2022) 038.
- [47] P. McDonald and A. Roy, Clustering of dark matter tracers: Generalizing bias for the coming era of precision LSS, *J. Cosmol. Astropart. Phys.* **08** (2009) 020.
- [48] V. Assassi, D. Baumann, D. Green, and M. Zaldarriaga, Renormalized halo bias, *J. Cosmol. Astropart. Phys.* **08** (2014) 056.
- [49] V. Desjacques, D. Jeong, and F. Schmidt, Large-scale galaxy bias, *Phys. Rep.* **733**, 1 (2018).
- [50] A. Eggemeier, R. Scoccimarro, and R. E. Smith, Bias loop corrections to the galaxy bispectrum, *Phys. Rev. D* **99**, 123514 (2019).
- [51] K. C. Chan, R. Scoccimarro, and R. K. Sheth, Gravity and large-scale nonlocal bias, *Phys. Rev. D* **85**, 083509 (2012).
- [52] E. Castorina, C. Carbone, J. Bel, E. Sefusatti, and K. Dolag, DEMNUni: The clustering of large-scale structures in the presence of massive neutrinos, *J. Cosmol. Astropart. Phys.* **07** (2015) 043.
- [53] L. Hui and K. P. Parfrey, Evolution of galaxy bias, generalized, *Phys. Rev. D* **77**, 043527 (2008).
- [54] K. Parfrey, L. Hui, and R. K. Sheth, Scale-dependent halo bias from scale-dependent growth, *Phys. Rev. D* **83**, 063511 (2011).
- [55] R. Scoccimarro, H. M. P. Couchman, and J. A. Frieman, The bispectrum as a signature of gravitational instability in redshift space, *Astrophys. J.* **517**, 531 (1999).
- [56] F. Villaescusa-Navarro, A. Banerjee, N. Dalal, E. Castorina, R. Scoccimarro, R. Angulo, and D. N. Spergel, The imprint of neutrinos on clustering in redshift space, *Astrophys. J.* **861**, 53 (2018).
- [57] F. Verdiani, E. Bellini, C. Moretti, E. Sefusatti, C. Carbone, and M. Viel, Redshift-space distortions in massive neutrinos cosmologies, [arXiv:2503.06655](https://arxiv.org/abs/2503.06655).
- [58] A. Chudaykin and M. M. Ivanov, Measuring neutrino masses with large-scale structure: Euclid forecast with controlled theoretical error, *J. Cosmol. Astropart. Phys.* **11** (2019) 034.
- [59] A. Aviles, A. Banerjee, G. Niz, and Z. Slepian, Clustering in massive neutrino cosmologies via Eulerian Perturbation Theory, *J. Cosmol. Astropart. Phys.* **11** (2021) 028.
- [60] H. E. Noriega, A. Aviles, H. Gil-Marín, S. Ramirez-Solano, S. Fromenteau, M. Vargas-Magaña, J. Aguilar, S. Ahlen, O. Alves, S. Brieden *et al.*, Comparing compressed and full-modeling analyses with FOLPS: Implications for DESI 2024 and beyond, *J. Cosmol. Astropart. Phys.* **01** (2025) 136.
- [61] N. Kaiser, Clustering in real space and in redshift space, *Mon. Not. R. Astron. Soc.* **227**, 1 (1987).
- [62] A. J. S. Hamilton, Measuring omega and the real correlation function from the redshift correlation function, *Astrophys. J. Lett.* **385**, L5 (1992).
- [63] J. C. Jackson, A critique of Rees's theory of primordial gravitational radiation, *Mon. Not. R. Astron. Soc.* **156**, 1P (1972).
- [64] R. Scoccimarro, Redshift-space distortions, pairwise velocities, and nonlinearities, *Phys. Rev. D* **70**, 083007 (2004).
- [65] A. Taruya, T. Nishimichi, and S. Saito, Baryon acoustic oscillations in 2D: Modeling redshift-space power spectrum from perturbation theory, *Phys. Rev. D* **82**, 063522 (2010).
- [66] A. G. Sánchez, R. Scoccimarro, M. Crocce, J. N. Grieb, S. Salazar-Albornoz, C. Dalla Vecchia, M. Lippich, F. Beutler, J. R. Brownstein, C.-H. Chuang *et al.*, The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: Cosmological implications of the configuration-space clustering wedges, *Mon. Not. R. Astron. Soc.* **464**, 1640 (2017).
- [67] D. Bianchi, M. Chiesa, and L. Guzzo, Improving the modelling of redshift-space distortions—I. A bivariate Gaussian description for the galaxy pairwise velocity distributions, *Mon. Not. R. Astron. Soc.* **446**, 75 (2015).
- [68] J. N. Grieb, A. G. Sánchez, S. Salazar-Albornoz, R. Scoccimarro, M. Crocce, C. Dalla Vecchia, F. Montesano, H. Gil-Marín, A. J. Ross, F. Beutler *et al.*, The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: Cosmological implications of the Fourier space wedges of the final sample, *Mon. Not. R. Astron. Soc.* **467**, 2085 (2017).
- [69] A. Semenaite, A. G. Sánchez, A. Pezzotta, J. Hou, R. Scoccimarro, A. Eggemeier, M. Crocce, C.-H. Chuang, A. Smith, C. Zhao, J. R. Brownstein, G. Rossi, and D. P. Schneider, Cosmological implications of the full shape of anisotropic clustering measurements in BOSS and eBOSS, *Mon. Not. R. Astron. Soc.* **512**, 5657 (2022).
- [70] L. Senatore and M. Zaldarriaga, Redshift space distortions in the effective field theory of large scale structures, [arXiv:1409.1225](https://arxiv.org/abs/1409.1225).
- [71] M. Crocce and R. Scoccimarro, Memory of initial conditions in gravitational clustering, *Phys. Rev. D* **73**, 063520 (2006).
- [72] D. J. Eisenstein, H.-J. Seo, and M. White, On the robustness of the acoustic scale in the low-redshift clustering of matter, *Astrophys. J.* **664**, 660 (2007).
- [73] D. J. Eisenstein, H.-J. Seo, E. Sirko, and D. N. Spergel, Improving cosmological distance measurements by reconstruction of the baryon acoustic peak, *Astrophys. J.* **664**, 675 (2007).

- [74] S. Matarrese and M. Pietroni, Baryonic acoustic oscillations via the renormalization group, *Mod. Phys. Lett. A* **23**, 25 (2008).
- [75] M. Crocce and R. Scoccimarro, Nonlinear evolution of baryon acoustic oscillations, *Phys. Rev. D* **77**, 023533 (2008).
- [76] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, Equivalence principle and the baryon acoustic peak, *Phys. Rev. D* **92**, 043514 (2015).
- [77] D. Blas, M. Garny, M. M. Ivanov, and S. Sibiryakov, Time-sliced perturbation theory II: Baryon acoustic oscillations and infrared resummation, *J. Cosmol. Astropart. Phys.* **07** (2016) 028.
- [78] M. M. Ivanov and S. Sibiryakov, Infrared resummation for biased tracers in redshift space, *J. Cosmol. Astropart. Phys.* **07** (2018) 053.
- [79] A. Taruya, Constructing perturbation theory kernels for large-scale structure in generalized cosmologies, *Phys. Rev. D* **94**, 023504 (2016).
- [80] G. D’Amico, M. Marinucci, M. Pietroni, and F. Vernizzi, The large scale structure bootstrap: Perturbation theory and bias expansion from symmetries, *J. Cosmol. Astropart. Phys.* **10** (2021) 069.
- [81] M. Garny, D. Laxhuber, and R. Scoccimarro, Perturbation theory with dispersion and higher cumulants: Framework and linear theory, *Phys. Rev. D* **107**, 063539 (2023).
- [82] M. Hartmeier and M. Garny, Minimal basis for exact time dependent kernels in cosmological perturbation theory and application to Λ CDM and w_0w_a CDM, *J. Cosmol. Astropart. Phys.* **12** (2023) 027.
- [83] S. Saito, M. Takada, and A. Taruya, Nonlinear power spectrum in the presence of massive neutrinos: Perturbation theory approach, galaxy bias, and parameter forecasts, *Phys. Rev. D* **80**, 083528 (2009).
- [84] G. Parimbelli, C. Carbone, J. Bel, B. Bose, M. Calabrese, E. Carella, and M. Zennaro, DEMNUni: Comparing nonlinear power spectra prescriptions in the presence of massive neutrinos and dynamical dark energy, *J. Cosmol. Astropart. Phys.* **11** (2022) 041.
- [85] C. Alcock and B. Paczynski, An evolution free test for non-zero cosmological constant, *Nature (London)* **281**, 358 (1979).
- [86] W. E. Ballinger, J. A. Peacock, and A. F. Heavens, Measuring the cosmological constant with redshift surveys, *Mon. Not. R. Astron. Soc.* **282**, 877 (1996).
- [87] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, MA, 2005).
- [88] S. L. Bridle, R. Crittenden, A. Melchiorri, M. P. Hobson, R. Kneissl, and A. N. Lasenby, Analytic marginalization over CMB calibration and beam uncertainty, *Mon. Not. R. Astron. Soc.* **335**, 1193 (2002).
- [89] J. Ruiz-Zapatero, B. Hadzhiyska, D. Alonso, P. G. Ferreira, C. García-García, and A. Mootooyaloo, Analytical marginalization over photometric redshift uncertainties in cosmic shear analyses, *Mon. Not. R. Astron. Soc.* **522**, 5037 (2023).
- [90] Z. Vlah, U. Seljak, M. Yat Chu, and Y. Feng, Perturbation theory, effective field theory, and oscillations in the power spectrum, *J. Cosmol. Astropart. Phys.* **03** (2016) 057.
- [91] D. J. Eisenstein and W. Hu, Power spectra for cold dark matter and its variants, *Astrophys. J.* **511**, 5 (1999).
- [92] K. Osato, T. Nishimichi, F. Bernardeau, and A. Taruya, Perturbation theory challenge for cosmological parameters estimation: Matter power spectrum in real space, *Phys. Rev. D* **99**, 063530 (2019).
- [93] J. Hamann, S. Hannestad, J. Lesgourgues, C. Rampf, and Y. Y. Y. Wong, Cosmological parameters from large scale structure—geometric versus shape information, *J. Cosmol. Astropart. Phys.* **07** (2010) 022.
- [94] J. N. Grieb, A. G. Sánchez, S. Salazar-Albornoz, and C. Dalla Vecchia, Gaussian covariance matrices for anisotropic galaxy clustering measurements, *Mon. Not. R. Astron. Soc.* **457**, 1577 (2016).
- [95] L. Pozzetti, C. M. Hirata, J. E. Geach, A. Cimatti, C. Baugh, O. Cucciati, A. Merson, P. Norberg, and D. Shi, Modelling the number density of H α emitters for future spectroscopic near-IR space missions, *Astron. Astrophys.* **590**, A3 (2016).
- [96] F. J. Castander, P. Fosalba, J. Stadel, D. Potter, J. Carretero, P. Tallada-Crespí, L. Pozzetti, M. Bolzonella, G. A. Mamon, and L. Blot (Euclid Collaboration), Euclid. V. The Flagship galaxy mock catalogue: A comprehensive simulation for the Euclid mission, *Astron. Astrophys.* **697**, A5 (2025).
- [97] J. D. Hunter, Matplotlib: A 2d graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [98] <https://pypi.org/project/comet-emu/>