Astronomy & Astrophysics manuscript no. ThousandsEuclidSkies
July 17, 2025

Euclid preparation

Simulating thousands of *Euclid* spectroscopic skies

Euclid Collaboration: P. Monaco[⋆1,2,3,4], G. Parimbelli[5,6,7], M. Y. Elkhashab[2,3,1,8], J. Salvalaggio[2,8,1,3], T. Castro[2,3,8,4], M. D. Lepinzan[1,2], E. Sarpa[7,4,3], E. Sefusatti[2,8,3], L. Stanco[9], L. Tornatore[2], G. E. Addison[10], S. Bruton[11], C. Carbone[12], F. J. Castander[5,13], J. Carretero[14,15], S. de la Torre[16], P. Fosalba[13,5], G. Lavaux[17], S. Lee[18], K. Markovic[18], K. S. McCarthy[18,19], F. Passalacqua[20,9], W. J. Percival[21,22,23], I. Risso[24,25], C. Scarlata[26], P. Tallada-Crespí[14,15], M. Viel[8,2,7,3,4], Y. Wang[27], B. Altieri[28], S. Andreon[24], N. Auricchio[29], C. Baccigalupi[8,2,3,7], M. Baldi[30,29,31], S. Bardelli[29], P. Battaglia[29], F. Bernardeau[32,17], A. Biviano[2,8], E. Branchini[33,25,24], M. Brescia[34,35], J. Brinchmann[36,37,38], S. Camera[39,40,41], G. Cañas-Herrera[42,43,44], V. Capobianco[41], V. F. Cardone[45,46], S. Casas[47], M. Castellano[45], G. Castignani[29], S. Cavuoti[35,48], A. Cimatti[49], C. Colodro-Conde[50], G. Congedo[51], C. J. Conselice[52], L. Conversi[53,28], Y. Copin[54], F. Courbin[55,56], H. M. Courtois[57], A. Da Silva[58,59], H. Degaudenzi[60], G. De Lucia[2], A. M. Di Giorgio[61], F. Dubath[60], F. Ducret[16], C. A. J. Duncan[51,52], X. Dupac[28], S. Dusini[9], A. Ealet[54], S. Escoffier[62], M. Farina[61], R. Farinelli[29], S. Farrens[63], S. Ferriol[54], F. Finelli[29,64], N. Fourmanoit[62], M. Frailis[2], E. Franceschi[29], M. Fumana[12], S. Galeotta[2], K. George[65], B. Gillis[51], C. Giocoli[29,31], J. Gracia-Carpio[66], A. Grazian[67], F. Grupp[66,68], L. Guzzo[69,24,70], S. V. H. Haugan[71], W. Holmes[18], F. Hormuth[72], A. Hornstrup[73,74], K. Jahnke[75], M. Jhabvala[76], B. Joachimi[77], E. Keihänen[78], S. Kermiche[62], B. Kubik[54], M. Kümmel[68], M. Kunz[79], H. Kurki-Suonio[80,81], A. M. C. Le Brun[82], S. Ligori[41], P. B. Lilje[71], V. Lindholm[80,81], I. Lloro[83], D. Maino[69,12,70], E. Maiorano[29], O. Mansutti[2], O. Marggraf[84], M. Martinelli[45,46], N. Martinet[16], F. Marulli[85,29,31], R. Massey[86], E. Medinaceli[29], S. Mei[87,88], M. Melchior[89], Y. Mellier[90,17], M. Meneghetti[29,31], E. Merlin[45], G. Meylan[91], A. Mora[92], M. Moresco[85,29], L. Moscardini[85,29,31], E. Munari[2,8], R. Nakajima[84], C. Neissner[93,15], S.-M. Niemi[42], C. Padilla[93], S. Paltani[60], F. Pasian[2], K. Pedersen[94], V. Pettorino[42], S. Pires[63], G. Polenta[95], M. Poncet[96], L. A. Popa[97], L. Pozzetti[29], F. Raison[66], A. Renzi[20,9], J. Rhodes[18], G. Riccio[35], F. Rizzo[2], E. Romelli[2], M. Roncarelli[29], R. Saglia[68,66], Z. Sakr[98,99,100], A. G. Sánchez[66], D. Sapone[101], B. Sartoris[68,2], P. Schneider[84], T. Schrabback[102,84], M. Scodeggio[12], A. Secroun[62], G. Seidel[75], M. Seiffert[18], S. Serrano[13,103,5], C. Sirignano[20,9], G. Sirri[31], J. Steinwagner[66], D. Tavagnacco[2], A. N. Taylor[51], I. Tereno[58,104], N. Tessore[77], S. Toft[105,106], R. Toledo-Moreo[107], F. Torradeflot[15,14], I. Tutusaus[99], L. Valenziano[29,64], J. Valiviita[80,81], T. Vassallo[68,2], G. Verdoes Kleijn[108], A. Veropalumbo[24,25,33], J. Weller[68,66], G. Zamorani[29], E. Zucca[29], V. Allevato[35], M. Ballardini[109,110,29], C. Burigana[111,64], R. Cabanac[99], M. Calabrese[112,12], A. Cappi[29,113], D. Di Ferdinando[31], J. A. Escartin Vigo[66], G. Fabbian[114,115], L. Gabarra[116], J. Martín-Fleitas[117], S. Matthew[51], N. Mauri[49,31], R. B. Metcalf[85,29], A. Pezzotta[24], M. Pöntinen[80], C. Porciani[84], V. Scottez[90,118], M. Sereno[29,31], M. Tenti[31], M. Wiesmann[71], Y. Akrami[119,120], S. Alvi[109], I. T. Andika[121,122], S. Anselmi[9,20,123], M. Archidiacono[69,70], F. Atrio-Barandela[124], S. Avila[14], A. Balaguera-Antolinez[50], P. Bergamini[69,29], D. Bertacca[20,67,9], M. Bethermin[125], A. Blanchard[99], L. Blot[126,82], S. Borgani[1,8,2,3,4], M. L. Brown[52], A. Calabro[45], B. Camacho Quevedo[8,7,2], F. Caro[45], C. S. Carvalho[104], F. Cogato[85,29], S. Conseil[54], S. Contarini[66], A. R. Cooray[127], O. Cucciati[29], S. Davini[25], G. Desprez[108], A. Díaz-Sánchez[128], J. J. Diaz[50], S. Di Domizio[33,25], J. M. Diego[129], A. Enia[30,29], Y. Fang[68], A. G. Ferrari[31], A. Finoguenov[80], F. Fontanot[2,8], A. Franco[130,131,132], K. Ganga[87], J. García-Bellido[119], T. Gasparetto[45], V. Gautard[133], E. Gaztanaga[5,13,134], F. Giacomini[31], F. Gianotti[29], G. Gozaliasl[135,80], M. Guidi[30,29], C. M. Gutierrez[136], A. Hall[51], S. Hemmati[137], C. Hernández-Monteagudo[138,50], H. Hildebrandt[139], J. Hjorth[94], S. Joudaki[14], J. J. E. Kajava[140,141], Y. Kang[60], V. Kansal[142,143], D. Karagiannis[109,144], K. Kiiveri[78], C. C. Kirkpatrick[78], S. Kruk[28], V. Le Brun[16], J. Le Graet[62], L. Legrand[145,146], M. Lembo[17], F. Lepori[147], G. Leroy[148,86], G. F. Lesci[85,29], J. Lesgourgues[47], L. Leuzzi[29], T. I. Liaudat[149], J. Macias-Perez[150], G. Maggio[2], M. Magliocchetti[61], C. Mancini[12], F. Mannucci[151], R. Maoli[152,45], C. J. A. P. Martins[153,36], L. Maurin[114], M. Miluzio[28,154], A. Montoro[5,13], C. Moretti[2,8,7,3], G. Morgante[29], S. Nadathur[134], K. Naidoo[134], A. Navarro-Alsina[84], S. Nesseris[119], K. Paterson[75], A. Pisani[62], D. Potter[147], S. Quai[85,29], M. Radovich[67], G. Rodighiero[20,67], S. Sacquegna[155,131,130], M. Sahlén[156], D. B. Sanders[157], D. Sciotti[45,46], E. Sellentin[158,44], L. C. Smith[159], J. G. Sorce[160,114], K. Tanidis[116], C. Tao[62], G. Testera[25], R. Teyssier[161], S. Tosi[33,25,24], A. Troja[20,9], M. Tucci[60], C. Valieri[31], A. Venhola[162], F. Vernizzi[32], G. Verza[163], P. Vielzeuf[62], and N. A. Walton[159]

(Affiliations can be found after the references)

July 17, 2025

ABSTRACT

We present two extensive sets of 3500+1000 simulations of dark matter haloes on the past light cone, and two corresponding sets of simulated ('mock') galaxy catalogues that represent the *Euclid* spectroscopic sample. The simulations were produced with the latest version of the `Pinocchio` code, and provide the largest, public set of simulated skies. Mock galaxy catalogues were obtained by populating haloes with galaxies using an halo occupation distribution (HOD) model extracted from the Flagship galaxy catalogue provided by Euclid Collaboration. The Geppetto set of 3500 simulated skies was obtained by tiling a $1.2\,h^{-1}$ Gpc box to cover a light-cone whose sky footprint is a circle of 30° radius, for an area of 2763 deg^2 and a minimum halo mass of $1.5 \times 10^{11}\,h^{-1}\,M_\odot$. The relatively small box size makes this set unfit for measuring very large scales. The EuclidLargeBox set consists of 1000 simulations of $3.38\,h^{-1}$ Gpc, with the same mass resolution and a footprint that covers half of the sky, excluding the Milky Way zone of avoidance. From this we produced a set of 1000 EuclidLargeMocks on the 30° radius footprint, whose comoving volume is fully contained in the simulation box. We validated the two sets of catalogues by analysing number densities, power spectra, and 2-point correlation functions, showing that the Flagship spectroscopic catalogue is consistent with being one of the realisations of the simulated sets, although we noticed small deviations limited to the quadrupole at $k > 0.2\,h\,\mathrm{Mpc}^{-1}$. We show cosmological parameter inference from these catalogues and demonstrate that using one realisation of EuclidLargeMocks in place of the Flagship mock produces the same posteriors, to within the expected shift given by sample variance. These simulated skies will be used for the galaxy clustering analysis of *Euclid*'s Data Release 1 (DR1).

Key words. Methods: numerical; Surveys; Cosmology: observations; Large-scale structure of Universe; Cosmological parameters

1. Introduction

The *Euclid* satellite (Euclid Collaboration: Mellier et al. 2025), launched on July 1st, 2023, is mapping the visible Universe to redshifts of at least $z = 2$ with the aim of shedding light on its dark sector. Together with other Stage IV surveys, like DESI (DESI Collaboration: Adame et al. 2024), Rubin–LSST (LSST Science Collaboration: Abell et al. 2009), Roman (Spergel et al. 2015), SphereX (Doré et al. 2014), or the SKAO cosmological surveys (SKA Cosmology Group: Bacon et al. 2020), *Euclid* is starting to contribute a cosmological data flood, opening to the possibility of constraining cosmology at an unprecedented precision level, comparable to what the CMB provides at $z \sim 1100$.

In particular, *Euclid* is surveying the sky with the VIS visible imager (Euclid Collaboration: Cropper et al. 2025), designed for weak lensing studies, and the near-infrared imager and spectrometer (NISP, Euclid Collaboration: Jahnke et al. 2025) that uses slitless spectroscopy to measure redshifts for a sample of emission-line galaxies; these redshifts are the basis for galaxy clustering measurements, that is the focus of our research. The wavelength range of the 'red grism' used by NISP to disperse the light makes it possible to cover the H$\,\alpha$ line in emission in a redshift range of roughly $z \in [0.9, 1.8]$. We expect that most galaxies with reliable redshift have line flux $f_{\mathrm{H}\alpha}$ larger than a fiducial value of $f_0 = 2 \times 10^{-16}$ erg s^{-1} cm^{-2}. The Euclid wide survey (EWS, Euclid Collaboration: Scaramella et al. 2022) will cover $\sim 14\,000$ deg^2 of the sky, excluding the zone of avoidance of the Milky Way and the ecliptic plane where zodiacal light hampers deep observations. The EWS will be sided by the Euclid deep survey (EDS), consisting of 53 deg^2 in three disconnected regions; these fields will be spectroscopically observed 15 times with the red grism, plus 25 times using a 'blue grism' to extend the spectral coverage to shorter wavelengths. The EDS will provide a highly pure and complete version of the galaxy sample detected in the EWS, allowing a precise assessment of its purity and completeness, a crucial ingredient to control systematic effects.

The sheer amount of high-quality data is already challenging the traditional galaxy survey processing methods. Access to a large cosmic volume, sampled with a dense galaxy catalogue, will allow us to beat down the statistical uncertainty of the measurement, so the error budget will be dominated by systematic effects. Because galaxy clustering is based on counting galaxies down to the faintest accessible flux, control of the depth of the survey will be crucial to have an unbiased estimate of the galaxy density. Because most redshift measurements will be based on spectra with low signal-to-noise ratio, control of redshift errors (and the consequent emergence of interlopers where redshifts are catastrophically wrong) will be key to keeping the purity of the spectroscopic sample high. Eventually, in a standard likelihood approach, we will compress data into a summary statistics of galaxy clustering, like the 2-point correlation function in configuration or Fourier space (the power spectrum), possibly augmented with the corresponding 3-point functions, and we will compare these with a model that predicts the summary statistics as a function of cosmological and nuisance parameters. The difference between measurement and model must be related to a covariance matrix, which for its cosmological part is a higher-order moment of clustering. The covariance matrix should also contain contributions from systematics error, representing the uncertainty in the mitigation of known systematic effects. As shown by Colavincenzo et al. (2017), even with the simplest assumptions, where the correction to the galaxy density is multiplicative, the cosmological covariance cannot be simply separated as a sum of contributions from cosmology and systematic effects. A numerical characterisation of the covariance matrix, obtained by processing thousands of simulated skies, remains the most effective way to address this challenge.

In this paper we present the simulated spectroscopic skies that have been produced by the Euclid Collaboration to face these challenges. These are based on the largest set of simulated dark matter haloes on the past light cone ever produced. We discuss two sets of simulated *Euclid* spectroscopic skies (whose properties are reported in Table 1), performed with the `Pinocchio` code (Monaco et al. 2002) for a flat Λ cold dark matter (ΛCDM) model compatible with *Planck* results (Planck Collaboration: Aghanim et al. 2020). The first set of 3500 realisations, named Geppetto, is based on a relatively small box of $1.2\,h^{-1}$ Gpc, sampled with 2160^3 particles to achieve a minimum halo mass of $1.5{\times}10^{11}\,h^{-1}\,M_\odot$. Each light-cone covers a circle of 30 degrees of radius on the sky, covering 2763 deg^2, a bit larger

⋆ e-mail: `pierluigi.monaco@inaf.it`

than the area planned to be covered by *Euclid*'s DR1 (Euclid Collaboration: Scaramella et al. 2022), and extends to redshift $z = 2$. The second set of 1000 realisations (named EuclidLargeBox) is based on a $3.38\,h^{-1}$ Gpc box, sampled by 6144^3 particles and achieving the same mass resolution as Geppetto. The lightcones cover half of the sky, excluding a zone of avoidance of the Milky Way, for a semi-aperture angle of 70 deg, and extend to $z = 4$.

The paper is organised as follows. Section 2 expands on the context and on the motivation for the production of these simulations. Section 3 describes the improvements to the `Pinocchio` code that were required to scale up to this large number of particles. Section 4 describes in detail the simulations produced, while Sect. 5 describes the techniques adopted to populate haloes with galaxies and produce mock galaxy catalogues; details in the adopted matching of `Pinocchio` and N-body masses are reported in an Appendix. Section 6 presents a set of measurements performed to validate the sets of mock catalogues and demonstrate their utility to obtain numerical covariance matrices for *Euclid*. Finally, Sect. 7 gives the conclusions.

2. Context

Our starting point is the upcoming spectroscopic sample of *Euclid* (Euclid Collaboration: Mellier et al. 2025), consisting of ~ 30 million galaxies with reliable redshift measurement in the range $0.9 < z < 1.8$ over 14 000 deg^2. Systematic effects for the measurement of the galaxy density are represented by constructing a 'visibility mask', a function of sky position and redshift that quantifies, to the best of our knowledge, the probability that a galaxy gets a reliable redshift measurement given its properties and the local observing conditions. This visibility mask is applied to a dense set of unclustered galaxies to obtain a 'random catalogue', that is used as a reference to compute the galaxy density contrast, thus mitigating systematic effects induced by a non-uniform selection probability. All estimators of summary statistics that are based on the galaxy density rely on this random catalogue.

To compute the numerical covariance of a summary statistics we start from a set of mock catalogues that represent galaxies in the past light cone, on an angular footprint matching the *Euclid* survey. Selection biases can be imposed to mock catalogues by applying the same visibility mask used to construct the random catalogue, including catastrophic redshift errors. Then the same pipeline used to process the real data catalogue is applied to the mock catalogues. The resulting measurements are combined to produce a numerical covariance, which is then fed to a likelihood code to estimate cosmological parameters. This approach has been used in several surveys (e.g. Manera et al. 2013; Balaguera-Antolínez et al. 2023), and has the clear advantage of being realistic in representing the level of non-linearity of the density field, the complexity of galaxy bias and the treatment of systematic effects .

A numerical covariance matrix is affected by noise as it is sampled by a finite number of mocks; for *Euclid*, the requirement on its accuracy is that parameter uncertainties (for the final sample after the nominal six years of the mission) should vary by less than 10%, with respect to an infinite number of mocks. In Euclid Collaboration: Sanchez et al. (in prep.), we estimate that 3500 mocks are needed to have an accurate, brute-force numerical covariance matrix. This number is mostly determined by the total number of measurements for which the covariance is computed and by the accuracy requirement. That paper also presents tests of several methods to *de-noise* a numerical covariance matrix, demonstrating that ~ 100 mocks may be sufficient to achieve good accuracy (see also Fumagalli et al. 2022) when addressing the relatively limited data vector of 2-point correlation functions (in configuration or Fourier space). However, this figure may be optimistic: the inclusion of 3-point statistics will further increase the length of the data vector, moreover a precise characterisation of systematic effects requires to analyse a large number of mock catalogues (e.g. Euclid Collaboration: Risso et al., in prep; Euclid Collaboration: Lee et al., in prep.). Under these premises, we consider 1000 as a minimal requirement on the number of realisations for a set of simulated catalogues for *Euclid*.

A single simulation should ideally cover a volume that can include the largest connected patch of the surveyed volume, so as to properly represent super-sample covariance; this pushes requirements on the box size beyond $3\,h^{-1}$ Gpc. The simulation should also resolve the smallest halo that contains an observed galaxy, that (as argued in Sect. 5) is expected to be $\sim 2 \times 10^{11}\,h^{-1}\,M_\odot$ for the spectroscopic sample; this pushes requirements on particle mass towards $10^9\,h^{-1}\,M_\odot$. One such simulation was used to create the Flagship galaxy mock catalogue (Euclid Collaboration: Castander et al. 2025), a comprehensive catalogue representing galaxies in all aspects that are of interest for *Euclid*, including weak lensing, galaxy clusters, and galaxy evolution. The single Flagship numerical simulation used to create the galaxy mock catalogue costed $\sim 1\,000\,000$ node-hours (corresponding to 68 000 000 core hours) on 4000 nodes of Piz Daint supercomputer, and required storage of order of petabytes. Running thousands of such simulations is out of the question.

An alternative to the N-body approach is to run fast simulations using approximate methods, to obtain thousands of catalogues of dark matter haloes at the cost of one single equivalent simulation. Here speed is obtained by treating the evolution of perturbations in the mildly non-linear regime using some flavour of Lagrangian Perturbation Theory (LPT), avoiding to integrate the trajectories of particles within dark matter haloes. Approximate methods were reviewed by Monaco (2016). They can be broadly separated into two classes, 'predictive' and 'calibrated'. The first class is typically based on Lagrangian methods that range from semi-analytic, peak-based methods like `Peak-Patch` (Bond & Myers 1996; Stein et al. 2019) or `Pinocchio` to fast N-body codes, typically particle-mesh (PM) codes, like `FastPM` (Feng et al. 2016) or `COLA` (Tassev et al. 2013). The predictive nature of these codes implies that once they are properly calibrated on N-body simulations they can be applied to any box size, mass resolution, and cosmology; as a matter of fact their computational cost, at fixed box size and mass resolution, is broadly proportional to the accuracy they can achieve. The second class is typically based on a sophisticated bias scheme applied to a low-resolution density field, like `EZMOCKS` (Chuang et al. 2015a), `Patchy` (Kitaura et al. 2015), `BAM` (Balaguera-Antolínez et al. 2019), `COVMOS` (Baratta et al. 2023), `PineTree` (Ding et al. 2024), or `GOTHAM` (Pandey et al. 2024). These methods must be calibrated on an N-body simulation each time they have to be used, but the time needed to produce a single realisation is tiny, even when compared with `Pinocchio`, and so they are useful for a massive production of tens of thousands of mocks.

Approximate methods were thoroughly tested in Chuang et al. (2015b) and in a series of studies in preparation of *Euclid* (Lippich et al. 2019; Blot et al. 2019; Colavincenzo et al. 2019). In these three papers 300 N-body simulations were compared with several approximate methods run on the same initial conditions, addressing respectively the 2-point correlation

function, the power spectrum, and the bispectrum of dark matter haloes at various mass thresholds. Consistency was tested by using these covariances for parameter estimation and requiring that the relative increase of parameter uncertainties is less than 10%. As a result, while the average of the summary statistics may be biased, cosmological parameters inferred using the numerical covariance obtained from N-body were found consistent, for many methods, with those obtained using approximate simulations, with differences in parameter uncertainties within the requirement mentioned above.

To produce the simulations presented in this paper, we used the 5.0 version (V5) of `Pinocchio` (Monaco et al. 2002; Munari et al. 2017b).[1] The code is described in detail in Sect. 3, we just outline here the motivation for its use. As pointed out in Munari et al. (2017a), a predictive approximate method must solve two problems: first compute how particles move from their Lagrangian position, second determine how particles group into haloes. While the first problem is relatively easy to solve using LPT or a PM code, the second problem in principle requires accurate particle positions at scales smaller than the virial radius of the halo, making it challenging even for fast PM codes to reconstruct small haloes. `Pinocchio` can be seen as a (semi-analytic) halo finder in Lagrangian space, thus solving the difficult problem of associating particles to haloes even with rather approximate displacements. This makes it easier to LPT to represent the density field; for instance, Munari et al. (2017a) showed that, when multi-stream regions are correctly collapsed into haloes, 3LPT yields an increase in accuracy in halo positions with respect to 2LPT that is as large as the increase in accuracy going from the Zeldovich approximation to 2LPT. This is not true when straight LPT is applied to initial conditions, because higher LPT orders lead to stronger dispersion of particles after orbit crossing, thus losing the advantage of higher order. `Pinocchio` has been extensively used in the literature; for instance, Oddo et al. (2020, 2021) and Rizzo et al. (2023) used 10 000 realisations of the same setting of Lippich et al. (2019) to check that `Pinocchio` is correctly representing the covariance of the bispectrum of dark matter haloes.

3. Code

`Pinocchio`, a C code with `MPI` parallelization, has been designed to generate very good approximations of catalogues of dark matter haloes, both in periodic boxes and in the past light cone, with full information on their mass, position and merger history. Using excursion set theory and ellipsoidal collapse, the code computes for each particle an estimate of the time of orbit crossing (the collapse time), then groups the collapsed particles into haloes, following their merger history, and places haloes at the final position using 3LPT. The original code is described in Monaco et al. (2002), the 2LPT and 3LPT extensions in Monaco et al. (2013) and Munari et al. (2017b). This section describes the latest technical developments of the code.

The starting point is the realisation of a linear density field on a regular grid, as in the generation of the initial conditions of a cosmological N-body simulation. The code is made of two main parts, the computation of collapse times and LPT displacements for each particle, and the grouping of collapsed particles into haloes ('fragmentation'), with the construction of halo merger histories and a light-cone with continuous time sampling. Collapse times are computed by Gaussian-smoothing the linear density field on many smoothing radii, then computing the second derivatives of the potential with fast Fourier transforms (FFTs); these are used to compute the collapse redshift of each particle using ellipsoidal collapse. We define the inverse collapse time as $F = 1 + z_c$, and store its highest value F_{max} for all smoothing radii. At the final smoothing radius $R = 0$ (meaning that the variance of the linear density field is only limited by the Lagrangian grid), the LPT displacement fields are computed, amounting to four vectors for each particle (of the three 3LPT displacement fields we only compute the first two, the third rotational term being negligible). The second part of the code uses the collapse times and displacements to group particles into haloes, with an algorithm that mimics hierarchical clustering. Because collapse here is identified with orbit crossing, collapsed particles are not necessarily contained in haloes, they may be part of the filamentary network that joins haloes. So particles may be classified into uncollapsed (still in single-stream regime), filaments, and halo particles. Having recognised the haloes without really running the simulation (we just perform a single 3LPT time-step when needed), we can see `Pinocchio` as a halo finder that works on the Lagrangian space of initial conditions, plus a 3LPT engine to place the haloes at the right position.

The fourth version (V4) of `Pinocchio` is described in Munari et al. (2017b), where scaling properties are demonstrated to be nearly ideal for a range of conditions, culminating in boxes sampled with 2160^3 particles as the first Millennium simulation (Springel et al. 2005). Parallelisation is achieved with the MPI (message passing interface) protocol. However, scaling to larger sizes is hampered by two factors. Firstly, the FFT library `FFTW` (Frigo & Johnson 2012) distributes the 3D computational domain in planes, so each task must allocate memory for as many particles as a single plane. For a memory need of 350 bytes per particle, a plane of a 6144^3 box would require 12 GB, making it hard to run on all the available cores in a node. Secondly, fragmentation is an intrinsically scalar process, so parallelisation is achieved by dividing the computation domain in sub-boxes, chosen to present the smallest surface over volume; this implies a round of communications among tasks to move from the FFT domain to the sub-box domain, that has a small computational cost. However, haloes at the border of the domain will not be correctly reproduced, unless fragmentation is performed on an augmented domain that contains all the particles that are needed to properly produce haloes. These 'ghost regions', called boundary layer in the code jargon, create an overhead in memory and in computing time. They cannot be arbitrarily small as their size must be as large as the Lagrangian size of the largest halo that is predicted to be present in the box. At $z = 0$ this can amount to $\sim 35\,h^{-1}$ Mpc, placing a hard limit on the number of tasks the code can be distributed over.

3.1. FFT solver

The first limitation was overtaken by using the `pFFT` library.[2] This uses `FFTW` to perform 1D Fourier transforms, but wraps the 2D and 3D transforms differently, distributing memory in planes, pencils or sub-volumes. When memory constraints allow it, distribution in planes is the most convenient one, while distribution in pencils and in sub-volumes has a slightly increasing computational cost. While usage of `pFFT` simply required an adaptation of the code that performs FFTs, the code that creates the initial conditions required a deep re-design. The fourth version of `Pinocchio` implemented the same loop in k-space present in `N-GenIC` and `2lpt-IC` codes (Springel 2005; Crocce et al.

[1] https://github.com/pigimonaco/Pinocchio

[2] https://github.com/mpip/pfft

2006), making it possible to reproduce simulations started from initial conditions produced with those codes by just providing the same seed for random number generator. `N-GenIC` works as follows: each task creates a 2D table of random seeds, one for each row of cells along the box, then loops in k-space by drawing seeds consecutively from the random sequence relative to each pencil, assuring that the same result is achieved for any domain decomposition. Moreover, k-space is populated starting from the fundamental mode and progressing in concentric cubes in Fourier space. This way, doubling the number of grid points per side will result in a higher resolution realisation that has exactly the same larger-scale modes as the lower resolution one. This feature is maintained in the current version of the code; detailed technical aspects will be presented in a forthcoming technical publication.

3.2. Fragmentation code

The second limitation to code scaling concerns the construction of haloes from collapsed particle (fragmentation), and it was overtaken by changing the way the code stores and accesses memory. The first part of the code, that computes collapse times, produces for each particle its inverse collapse time $F_{\rm max}$ and the four LPT displacement fields. In the sub-box domain, the tasks do not really need these products for all particles, because particles that do not collapse by the final redshift will never be processed; moreover, not all the regions of the boundary layer are necessarily of interest. The code was then redesigned to have a more accurate selection of the particles that are loaded in the sub-box domain. To achieve this aim, each task gathers the properties of all collapsed particles in its 'well-resolved' region, that is the patch of Lagrangian space assigned to the task, augmented just by a 1-particle ghost region to properly compute the peaks of $F_{\rm max}$. This is done by creating a Boolean map of the whole computational domain (each point represented by a bit), setting to true the bits corresponding to the particles of interest. Of this domain, the task will receive (through a hypercubic communication scheme among tasks) only the products of the particles that are predicted to collapse by the end of the simulation. After that, the fragmentation code is run in a minimal configuration (no light-cone construction, no output) with the available information. At the end of this first fragmentation, haloes near the borders are not reconstructed properly but we know where they are. Another all-to-all communication round is then called: each task creates a second Boolean map that is set to `true` for all the particles that lie within a certain distance from the haloes and are beyond the well-resolved region. This distance is set as the Lagrangian radius of the halo (in grid units this is the cubic root of the number of its particles) times a boundary layer factor $f_{\rm bl} \sim 2$–3. After the products of these particles are communicated, the fragmentation is started again.

This implementation limits memory overhead by a large factor, improving the scaling of the fragmentation part of the code. Another large factor in memory requirement is obtained by storing the products of the collapse time part in single floating-point precision, while most of the code works in double precision; it has been checked that this creates an acceptable decrease in numerical accuracy, while yielding a drastic advantage in memory requirements. As a result, we pass from 350 to 150 bytes per particle as a minimum needed by 3LPT. However, while the previous version of the code had a predictable mapping between sub-box space domain and memory, accessing the memory location of a particle is less straightforward, and it is made possible by a series of pointers and some bisector searches. This increases the computing time in ideal cases, but the improvement in scaling overcomes this problem.

In this implementation, each task requires a variable amount of memory; for instance, if the largest cluster in the box falls at the border of a domain, its task will require a larger number of particles from the boundary layer, while a task with voids at its border will have a smaller overhead. This imbalance is handled by allowing the same overhead to all tasks; if a task fills all of its memory it issues a warning but does not halt the computation, and the largest overhead needed by the task is output at the end of the code, so it is easy to achieve an optimisation of memory usage after few test runs. This is considered as an acceptable strategy to prepare the production of thousands of massive runs.

4. Dark matter catalogues and light-cones

As discussed above, we produced two sets of simulations aimed at achieving a large number of realizations based on a small box (Geppetto) and a large box that can contain a survey, with a more limited number of realizations (EuclidLargeBox). The main properties of these sets are reported in Table 1. Both sets are based on a ΛCDM cosmology with parameters close to Planck Collaboration: Aghanim et al. (2020): $\Omega_{\rm m} = 0.32$, $\Omega_\Lambda = 0.68$, $\Omega_{\rm b} = 0.049$, $h = 0.67$, $\sigma_8 = 0.83$, $n_{\rm s} = 0.96$. These parameters are consistent with Flagship 2 simulation and mock galaxy catalogue; in the Flagship 2 simulation we have $\Omega_{\rm m} = 0.319$ and massive neutrinos, but these differences have no practical effect for the present purposes. For each simulation we produced a few dark matter halo catalogues in periodic boxes and a complete catalogue in the past light cone, with a sky footprint consisting of a circle with a given radius. Moreover, merger trees are available for all haloes. The minimum halo mass is set to 10 particles; this would be an unacceptably low value for an N-body simulation, but it is acceptable for a semi-analytic code where haloes are related to peaks of the inverse collapse time and each halo starts with one particle.

4.1. Geppetto simulations

The first set of $N_{\rm real} = 3500$ simulations, called Geppetto, was designed to be massive but relatively inexpensive, to produce a brute-force numerical covariance for the target sample of Hα emitters. It is based on a relatively small box of side $L_{\rm s} = 1200\,h^{-1}$ Mpc, sampled with $N_{\rm part} = 2160^3$ particles. With the given cosmology, the particle mass is $1.52 \times 10^{10}\,h^{-1}\,M_\odot$, and the smallest halo has a mass of 10 particles or $M_{\rm h} = 1.52 \times 10^{11}\,h^{-1}\,M_\odot$. The light-cones start at $z_{\rm start} = 2$ and cover a circle of radius of $\theta = 30°$, for a total area of $2\pi\,(1 - \cos\theta)$ steradians, equal to 2763 deg^2.[3] Outputs at fixed time are available at redshifts 2, 1.8, 1.5, 1.35, 1, 0.9, 0.5, and 0.

These simulations were mostly run on the Pleiadi system of INAF ('Istituto Nazionale di Astrofisica'), a Tier-2 facility with 66 computing nodes, each with 36 cores and 256 GB or RAM. Each run required 20 nodes, and lasted around 17 minutes, for a cost of 200 core-h per run. The total cost amounted to only 700 000 core-h. Each run required 23 GB of storage, for a total need of almost 80 TB. As a matter of fact, the first 600 realisations were run with V4.1.3 of the code, with minimal changes

[3] Noticeably, this specific sky footprint implies that the comoving volume of the light-cone is a 3D cone itself; however, the word 'cone' in the two cases refers to a spacetime and to a volume in 3D space, thus assuming a pretty different meaning.

Table 1. Main properties of the two sets of mock catalogues presented in this paper, Geppetto and EuclidLargeBox, together with the Minerva-like set used in Oddo et al. (2020) and NewClusterMocks used in Fumagalli et al. (2021). Columns give the simulation box side L_s, the number of particles in the simulation N_{part}, the number of realisations N_{real}, the minimum halo mass M_h, the simulation volume V_s and its total over the realisations, the semi-aperture of the light-cone θ, the survey area, the starting redshift of the light-cone z_{start}, the volume of a single light-cone V_{lc} and its total over the realisations.

Name	L_s h^{-1} Gpc	N_{part}	N_{real}	min. M_h $h^{-1}\,M_\odot$	V_s h^{-3} Gpc^3	tot. V_s h^{-3} Gpc^3	θ	area deg^2	z_{start}	V_{lc} h^{-3} Gpc^3	tot. V_{lc} h^{-3} Gpc^3
Geppetto	1.2	2160^3	3500	1.52×10^{11}	1.73	6047	30°	2763	2.0	12.72	44 529
EuclidLargeBox	3.38	6144^3	1000	1.48×10^{11}	38.61	38 614	70°	13 572	4.0	163.86	163 862
Minerva-like	1.5	1000^3	10 000	2.67×10^{12}	3.38	33 750	—	—	—	—	—
NewClusterMocks	3.87	2160^3	1000	4.90×10^{12}	57.96	57 960	60°	10 313	2.5	69.70	69 700

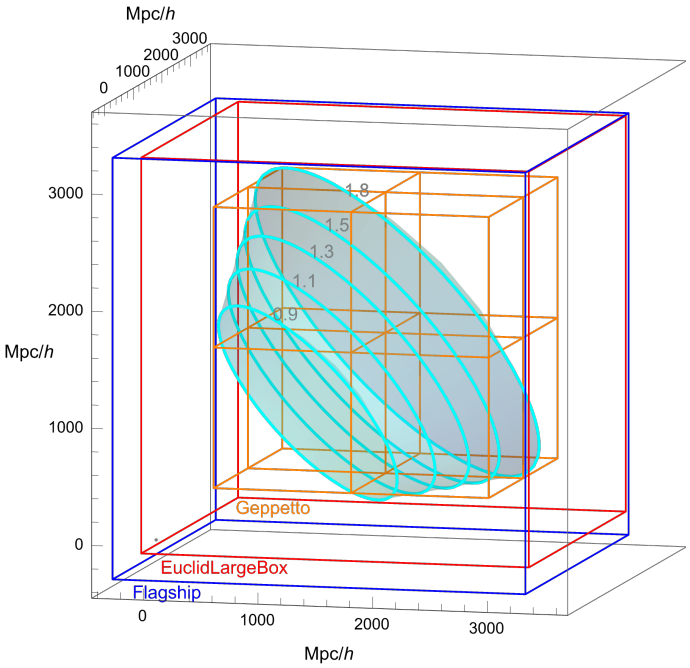

Fig. 1. Graphical representation of the comoving volume of the 30° radius survey footprint, highlighting the four redshift bins bounded by redshifts [0.9, 1.1, 1.3, 1.5, 1.8], immersed in the comoving volume of the Flagship simulation (that is our reference), of EuclidLargeBox, and of Geppetto simulation boxes. The Geppetto box is tiled a few times to cover the survey volume; for sake of clarity, the tiling is limited to the main replications. Here the central axis of the footprint lies on the main diagonal of the box, however all EuclidLargeBox and most Geppetto boxes have random orientations of this axis.

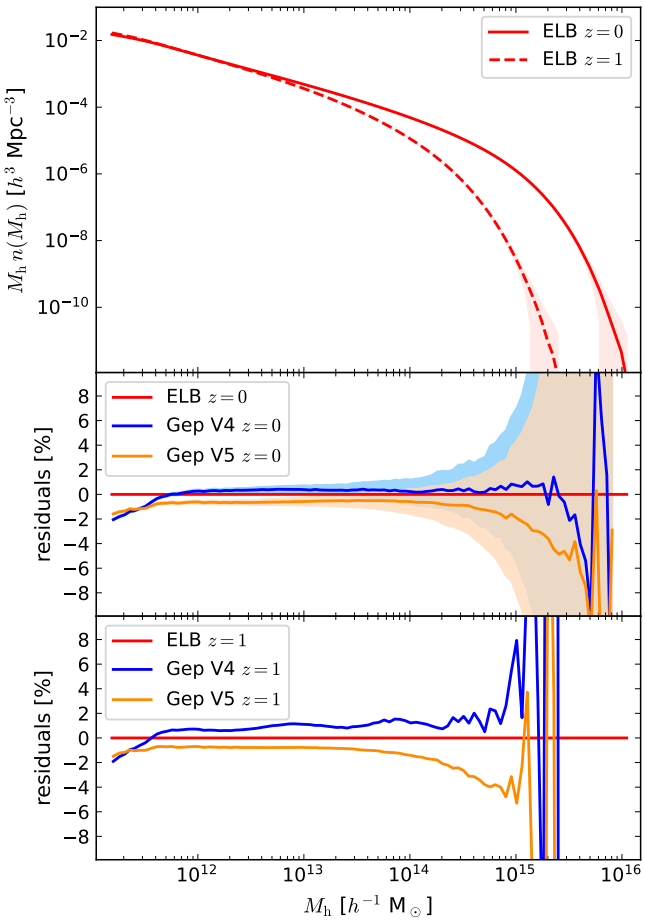

Fig. 2. Halo mass function $n(M_h)$ of the simulated sets. The upper panel only shows the mass function of EuclidLargeBox (ELB) for $z = 0$ and $z = 1$, where the line gives the average over the 1000 simulations and the shaded area gives its sample variance. The two lower panels show relative residuals, $n/n_{elb} - 1$ in percent, of the mass function of the first 600 Geppetto simulations, run with V4 of the code (Gep V4), and the other 2900, run with V5 (Gep V5), with respect to the EuclidLargeBox measurement, respectively at $z = 0$ (mid panel) and $z = 1$ (lower panel). Here the reported sample variance is relative only to the Geppetto sets.

(using single precision of the products of collapse time calculation and using a rather limited boundary layer to minimise memory overhead) to allow the code to scale to the needed mass resolution.

The comoving volume of the produced light-cones is 12.72 h^{-3} Gpc^3, that is 7.35 times larger than the simulated volume, 1.73 h^{-3} Gpc^3. Figure 1 reports the light-cone volume, limited to the four redshift bins that contain the spectroscopic sample (defined below). This is immersed in the Flagship simulation box (46.7 h^{-3} Gpc^3) that completely contains it. The figure also reports the Geppetto box, tiled to cover the volume. For sake of clarity we only report the main replications that cover most of the survey volume, in fact the tiling extends to all replications that have a non-null intersection with the cone.[4] Tiling is applied by using periodic boundary conditions to guarantee continuity of

[4] The code includes an algorithm dedicated to determining whether a given cone and cube intersect, a geometrically non-trivial problem despite involving only high-school-level maths.

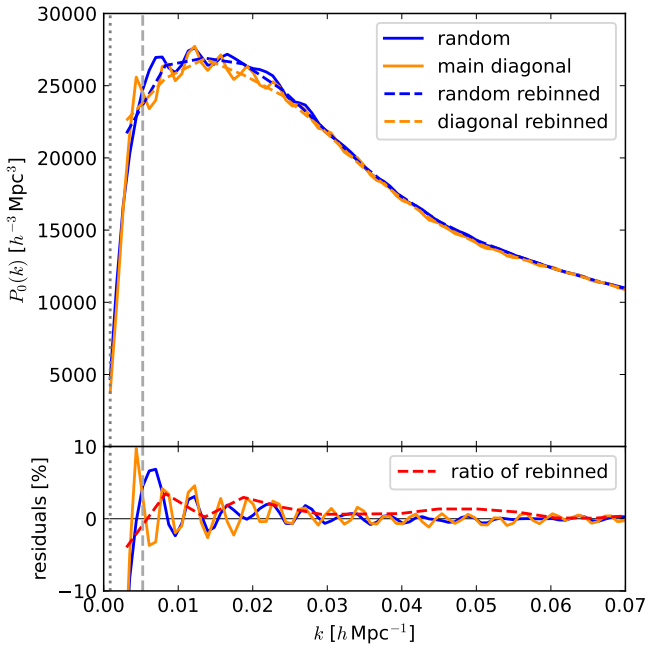

Fig. 3. Power spectrum monopole of `Pinocchio` dark matter haloes from the Geppetto light-cones, with $M_\mathrm{h} > 10^{12}\,h^{-1}\,M_\odot$ in the redshift range $z \in [0.9, 1.1]$. The continuous lines show measurements in a box of $L_\mathrm{m} = 7.2\,h^{-1}$ Gpc sampled on its fundamental frequency k_fm, the dashed lines show the same measurements with a sampling of $6k_\mathrm{fm}$, equal to the fundamental mode of the simulation box k_fs. Blue and orange lines show respectively averages over 100 mocks with random orientations of the cone axis and with axis aligned with the main diagonal of the simulation box. The lower panel shows the residuals of the measurements along a random orientation with respect to those with axis along the main diagonal, plus, in red, the residuals of the two rebinned measurements.

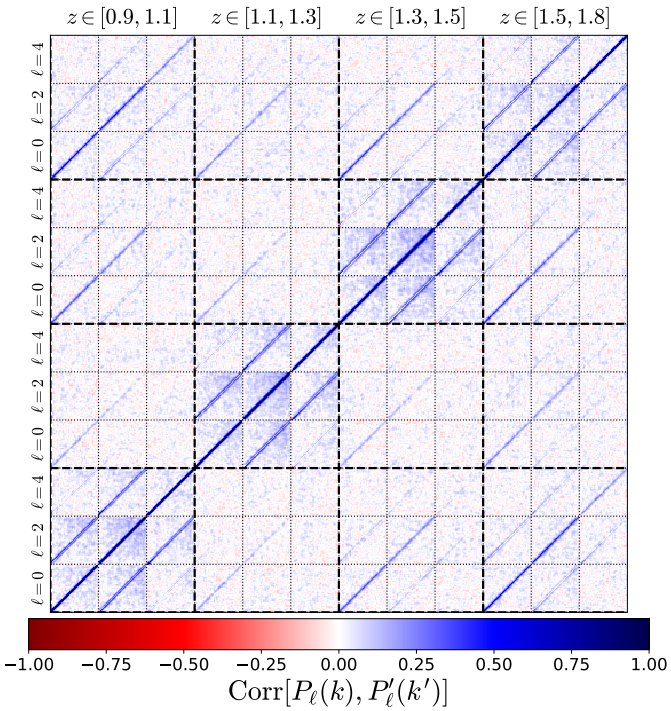

Fig. 4. Correlation matrix of the monopole of the power spectrum of dark matter haloes with $M_\mathrm{h} > 10^{12}\,h^{-1}$ Mpc, measured on the four redshift bins with edges in [0.9, 1.1, 1.3, 1.5, 1.8], for the Geppetto sets. The spurious cross-correlations due to replications are evident.

the density field. To minimise the statistical effect of these replications, the orientation of the light-cone volume is varied: while in the first 700 realisations the central axis of the cone, pointing from the observer to the centre of the sky footprint, is directed along the main diagonal of the box as in Fig. 1, in the other cases the axis is aligned along a random direction.

Averaging over so many realisations produces a very smooth halo mass function that beats down sample variance to negligible levels up to the mass of galaxy groups. The halo mass function is shown in Fig. 2; we separate the Geppetto mocks in two sets, the first 600 simulations, performed with V4 of the code, and the remaining 2900, run with V5. We also include results for the EuclidLargeBox set described below. The upper panel gives the halo mass function $M_\mathrm{h}\,n(M_\mathrm{h})$ at $z = 0$ and $z = 1$, and because the three curves are indistinguishable we only report results for the EuclidLargeBox set. The lower panels show the relative difference of the Geppetto halo mass functions with respect to the EuclidLargeBox one, $n/n_\mathrm{elb} - 1$ in percent, again at $z = 0$ (mid panel) and $z = 1$ (lower panel). In all cases the lines show the average mass functions while the shaded areas give their sample variance; in the residuals the variance refers only to the Geppetto sets, for the denominator n_elb we only use the average. The comparison of the results of V4 and V5 versions shows subtle percent-level differences. This discrepancy represents a residual of the calibration process, which is accounted for and reabsorbed during the subsequent HOD calibration (as discussed in Sect. 5).

4.2. The problem of replications

The starting point of any simulation is a realisation of the linear density field on a periodic box of length L_s, that samples the Fourier space modes on a regular grid with a cell size equal to the box fundamental mode $k_\mathrm{fs} = 2\pi/L_\mathrm{s}$. While, as illustrated in Fig. 1, the survey volume is fully contained in the EuclidLargeBox volume, in the Geppetto case the box is tiled several times to cover the light-cone. As pointed out in Howlett et al. (2015), when the measurement of the power spectrum is performed, a redshift bin is cut from the cone volume and immersed in a box of size L_m that is at least twice as large as the largest distance (in the three dimensions) among galaxies in the catalogue. Then the resulting density is Fourier-transformed to measure the power spectrum. The natural sampling of this measurement is on a Fourier-space grid of cell size $k_\mathrm{fm} = 2\pi/L_\mathrm{m}$, smaller than k_fs and not necessarily in some harmonic relation with it. This implies an oversampling of the original Fourier space, that becomes pretty evident when the survey volume covers several replications of the box.

To highlight this effect, we computed the power spectrum of the dark matter haloes of the Geppetto light-cones, with masses $> 10^{12}\,h^{-1}\,M_\odot$. In order to correct for the small difference in halo mass function between the first 600 light-cones and the remaining ones (Fig 2), we set the lower limits in halo mass at $10^{12}\,h^{-1}\,M_\odot$, corresponding to 66 particles, in the first 600 runs and to $0.99 \times 10^{12}\,h^{-1}\,M_\odot$, corresponding to 65 particles, in the other runs. Although the difference corresponds to only 1 particle, it is sufficient to ensure a consistent amplitude of the halo power spectrum. The power spectrum was computed using the *Euclid* ground-segment code (Euclid Collaboration: Salvalaggio et al., in prep.) that implements a Yamamoto–Bianchi estimator (Yamamoto et al. 2006; Bianchi et al. 2015). We produced catalogues of DM haloes in redshift space, in the red-

shift range $z \in [0.9, 1.1]$, and immersed them in boxes of side $L_m = 7.2\,h^{-1}$ Gpc. The resulting power spectrum is shown in Fig. 3, where we show averages for 100 mocks performed with V4 and with cone axis aligned with the main diagonal (orange lines) and 100 mocks performed with V5 with random direction of the cone axis (blue lines) to mitigate the replication problem. The figure also reports the fundamental modes of the simulation box k_{fs} and of the measurement box k_{fm}. The beating due to the oversampling of the Fourier space is very evident especially in the orange lines, and the randomisation of the cone direction mitigates but does not remove the oscillations (averaging over more realisations further dampens these oscillations without fully removing them).

One may argue that, more than being a problem of the data set, this is a problem of the measurement. By producing periodic boxes of $1.2\,h^{-1}$ Gpc we are sampling the density field with a hard limit on scales: scales larger than the box are simply not sampled. While replications are a way to produce larger volumes with a continuous density field, one cannot assume that they are adding information. So if the natural Fourier grid to measure the power spectrum of the periodic box is k_{fs}, one should use this same sampling for the power spectrum in the light-cone. Figure 3 shows the same measurements resampled over 6 times the fundamental mode, $k_{bin} = 6\,k_{fm} = k_{fs}$. It is evident that the beating disappears once the sampling of Fourier space respects the properties of the original box. However the two measurements show percent-level differences at large scales, that vanish toward the first baryonic acoustic oscillation at $k \sim 0.07\,h\,\mathrm{Mpc}^{-1}$, that may in part be due to the different code versions. These percent-level differences may artificially increase the covariance in a scale range where sample variance is already important, but they are unlikely to be a significant issue for a cosmological analysis.

Another effect of replications is to create cross-correlations of different redshift bins. We show in Fig. 4 the correlation matrix (the covariance matrix normalised to its diagonal) of the power spectrum monopole of dark matter haloes over four redshift bins with bin edges in [0.9, 1.1, 1.3, 1.5, 1.8], obtained from the whole set of 3500 measurement. The cross-correlation of different redshift bins is expected to be non-zero due to lensing (e.g. Euclid Collaboration: Lepori et al. 2022) or to systematic effects (Monaco et al. 2019), but none of these effects is present in the mocks, so these cross-correlations arise due to the replications, that make the density field of different redshift bins not independent. While this is a limitation, non-null cross-covariances may be needed only for very specific cases, and setting them to zero is a convenient way to solve the problem in practice.

4.2.1. EuclidLargeBox simulations

The second set of $N_{real} = 1000$ simulations, called EuclidLargeBox, was designed to have a more limited set of much larger and more expensive simulations, not affected by the replications problem. They are based on a $L_s = 3.38\,h^{-1}$ Gpc box sampled with $N_{part} = 6144^3$ particles. With these choices, the particle mass is $1.48 \times 10^{10}\,h^{-1}\,M_\odot$, very similar to that of Geppetto simulations, and the smallest output halo (with at least 10 particles) has a mass of $M_h = 1.48 \times 10^{11}\,h^{-1}\,M_\odot$. This box is marginally smaller than that ($3.6\,h^{-1}$ Gpc) of the Flagship simulation, that is however sampled with $16\,000^3$ particles, reaching a particle mass of $10^9\,h^{-1}\,M_\odot$. The halo mass functions of the EuclidLargeBox set, shown in Fig. 2, are consistent to within 1% with the ones from the Geppetto mocks. These simulations produced much larger light-cones, covering half of the sky and starting at redshift $z = 4$. Because the *Euclid* survey will never go deep in the zone of avoidance of the Milky Way, sampling Galactic latitudes always with $|b| > 20$, light-cones were produced with a sky footprint of a circle with radius of $\theta = 70°$, thus reducing storage needs by 34% with respect to a really half-sky output. With the same aim of reducing storage, output in periodic boxes was limited to redshifts 0 and 1. This produced 215 GB per run, 210 TB in total, of which 80 GB per run were taken by the light-cones.

To handle the memory overhead of the fragmentation and its variations over many realisations, it was necessary to allocate 250 bytes per particle, for a total memory requirement of 52 TB. After a few experiments, we decided to run the code on 24 fat nodes of the machine Galileo100@CINECA, equipped with 3 TB each and 48 cores per node. Each run took on average of 3 hours 45 minutes, for a computational cost of 4276 core-h per run, amounting to a total of about 4.3 million core-h.

Due to both larger redshift range and wider sky area, the volume sampled by the light-cone grows from $12.72\,h^{-3}\,\mathrm{Gpc}^3$ to $163.86\,h^{-3}\,\mathrm{Gpc}^3$, that is 4.24 times larger than the simulated box, whose volume is $38.61\,h^{-3}\,\mathrm{Gpc}^3$. Indeed, the light-cone is covered by tiling the box 23 times, including tiles that have very little overlap in volume. However, much of the volume increase is due to the higher redshift reached by the light-cone, that is relevant for adding interlopers to the spectroscopic sample due to catastrophic redshift errors. As shown in Fig. 1, the comoving volume of a survey from $z = 0.9$ to 1.8 and with a 30° radius sky footprint is contained in one simulated box. We will use the 30° radius footprint for most of the preparation papers, so the measurements of the EuclidLargeMocks galaxy catalogues produced from the EuclidLargeBox simulations, that are described below, are expected to be free of the problem of replications.

The volume sampled by the Geppetto and EuclidLargeBox sets is respectively $6047\,h^{-3}\,\mathrm{Gpc}^3$ and $38\,614\,h^{-3}\,\mathrm{Gpc}^3$, for a total of $44\,661\,h^{-3}\,\mathrm{Gpc}^3$, while the light-cones have volumes of $44\,529\,h^{-3}\,\mathrm{Gpc}^3$ and $163\,862\,h^{-3}\,\mathrm{Gpc}^3$. These can be compared to the $3400\,h^{-3}\,\mathrm{Gpc}^3$ of our visible Universe. All these numbers are reported in Table 1, together with the same numbers for two other sets of simulations produced with `Pinocchio` and presented in previous papers, namely the 10 000 Minerva-like simulations used in Oddo et al. (2020) and the 1000 NewClusterMock simulations aimed at describing galaxy clusters and presented in Fumagalli et al. (2021). The simulation sets presented here go down by a factor of ∼30 in halo mass, necessary to describe the spectroscopic sample that *Euclid* will observe.

This archive of data sets can be compared with other simulation sets available in the literatures, as the remarkable Quijote[5] project (Villaescusa-Navarro et al. 2020), that samples a volume of $82\,000\,h^{-3}\,\mathrm{Gpc}^3$ in boxes of $4\,h^{-1}$ Gpc of side. While the volume covered by Geppetto and EuclidLargeBox together is only half of it, the total volume produced with `Pinocchio` for *Euclid* studies, including the low-resolution boxes, is even larger than this ground-breaking figure; of course, our sets do not (yet) cover that rich variety of cosmologies. Our data set is unique in the box size, while still achieving a suitable mass resolution to sample galaxies, and in the number of light-cones, that presently provide the largest data set of dark matter haloes in the light-cone ever produced and shared with the cosmological community.

5. From haloes to galaxies

In this section we describe the procedure that we implemented to obtain *Euclid* spectroscopic skies from our collection of halo

[5] https://quijote-simulations.readthedocs.io

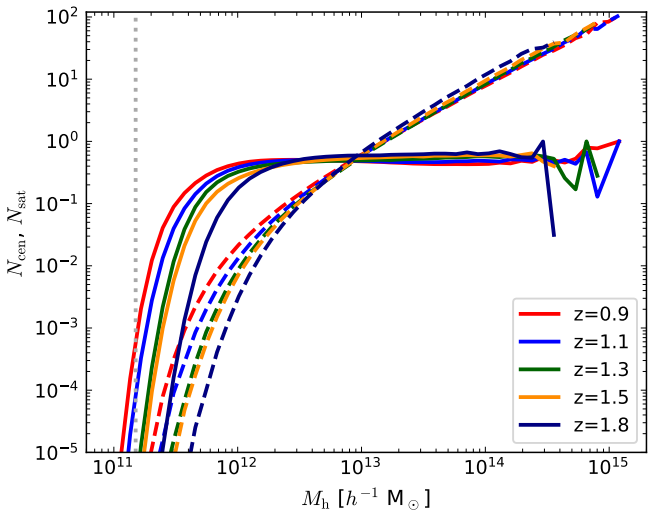

Fig. 5. HOD curves $N_{\rm cen}(M_{\rm h}, z)$ and $N_{\rm sat}(M_{\rm h}, z)$, respectively as continuous and dashed lines, measured from the Flagship mock catalogue as described in the main text. The five colours give the curves at five redshifts, as specified in the legend. The vertical dotted line gives the minimum halo mass in the `Pinocchio` mocks.

light-cones, thus bridging the gap from unobservable dark matter haloes to visible galaxies.

The reference data set for *Euclid* preparation papers is the Flagship simulation and mock catalogue described in the Flagship paper. We aim at creating a set of galaxy catalogues that have the same mean number density and 2-point clustering as a spectroscopic sample extracted from the Flagship mock, for a broad range of selections. All the galaxy catalogues that we produce are contained in a 30° radius footprint, as the Geppetto light-cones; the EuclidLargeBox footprint is much larger, but to have a consistent approach in the test of the galaxy power spectrum and 2-point correlation function we decided to use, across several *Euclid* preparation papers, this footprint; we thus call EuclidLargeMocks the galaxy catalogues extracted from the EuclidLargeBox set of simulations using a 30° radius footprint. This footprint can be fully immersed in the Flagship octant, all comparisons with Flagship will be done consistently using this footprint.

The steps that lead from the Flagship N-body simulation to a set of galaxies with measurable H$\,\alpha$ flux are thoroughly described in the Flagship paper. The procedure is based on an HOD conceived and calibrated to produce a very large range of galaxy properties, with the flux of galaxy emission lines coming at the end of the process. This makes it impossible to analytically infer HOD curves for galaxies selected in H$\,\alpha$ line flux $f_{\rm H\alpha}$. The number density of galaxies with $f_{\rm H\alpha}$ larger than a threshold (of order of f_0 defined in Sect. 1) is close to model 3 of Pozzetti et al. (2016). When plotted as a function of redshift, this number density shows little blips due to an issue with the interpolation of internal galaxy extinction; these blips have no appreciable practical consequences but will be sometimes visible in the figures that we will show below.

It is not convenient to directly implement such an assignment scheme to populate thousands of Flagship-like simulations; we then bypass this whole procedure by measuring the HOD functions directly from the Flagship catalogue. The algorithm, very similar to that used in the perturbation theory challenge (Euclid Collaboration: Pezzotta et al. 2024, where it was applied to Flagship haloes in the periodic box), runs as follows. We first set a selection criterion for the galaxy catalogue, based on a measured galaxy property, in this case the H$\,\alpha$ line flux: $f_{\rm H\alpha} > f_{\rm lim}$. We repeat the selection using a very fine grid of limiting flux values $f_{\rm lim}$, starting from $0.5 f_0 = 10^{-16}$ erg s^{-1} cm^{-2}, below which we assume that a reliable measure of redshift is very unlikely. For each flux cut we measure the two HOD curves relative to central and satellite galaxies in a grid of halo mass and redshift (without redshift-space distortions):

$$N_{\rm cen}(M_{\rm h}, z\,|\,f_{\rm H\alpha} \geq f_{\rm lim}) = \frac{\#\text{ of central galaxies}}{\#\text{ of haloes}}, \quad (1)$$

$$N_{\rm sat}(M_{\rm h}, z\,|\,f_{\rm H\alpha} \geq f_{\rm lim}) = \frac{\#\text{ of satellite galaxies}}{\#\text{ of haloes}}, \quad (2)$$

where the numbers are relative to that specific bin and flux cut. Here $N_{\rm cen}$ should be interpreted as the probability that a halo of given mass and at a given redshift hosts a central galaxy that satisfies our selection criterion, $N_{\rm sat}$ as the average number of satellites in that same halo. We choose a redshift bin $\delta z = 0.01$ from 0 to 3 (where the Flagship mock ends) and a mass bin of $\delta \log_{10}(M_{\rm h}/M_\odot) = 0.083$; to limit noise from sample variance we smooth these curves in redshift with a Gaussian kernel of width of five bins. The resulting curves are presented in Fig. 5, as a function of $M_{\rm h}$, for $f_{\rm H\alpha} > f_0$ and for several redshifts from 0.9 to 1.8. While $N_{\rm sat}$ can overshoot unity, $N_{\rm cen}$ is constrained to be ≤ 1, and its values flatten at large masses at ~ 0.4. The plots also report the mass of the smallest halo in our Geppetto and EuclidLargeBox simulations as a vertical dotted line, to show that the number of H$\,\alpha$-emitting galaxies hosted in dark matter haloes below our mass resolution is expected to be negligible, at least according to this HOD.

The same algorithm can be used to work out HOD curves for any other observational selection, like a limit in the *Euclid* magnitude $H_{\rm E} < 24$ that defines the photometric sample for which a measurement of the spectrum is attempted. However, the mass resolution of our simulation sets has been chosen to resolve the haloes that contain galaxies of the *Euclid* spectroscopic sample, but galaxies of the photometric sample are hosted by smaller haloes and so a photometric mock catalogue would be incomplete, especially at low redshift.

To each halo we associate a single central galaxy, with a probability $N_{\rm cen}$, and a number of satellite galaxies drawn from a Poissonian distribution with mean $N_{\rm sat}$. Central galaxies inherit the halo position and velocity, while satellite galaxies are distributed following a Navarro–Frenk–White profile (NFW, Navarro et al. 1996), with concentration given by the Diemer & Joyce (2019) relation and velocities around the halo centre of mass with magnitude

$$V_{\rm sat}(r) = f_{\rm v} \sqrt{\frac{G M_{\rm h}(<r)}{r}}, \quad (3)$$

where $f_{\rm v}$ is a calibration constant and $M_{\rm h}(<r)$ is the halo mass within distance from halo centre r, given by the NFW profile. Orientations of satellite velocities are randomly drawn from a sphere. The calibration constant was added to recover the velocity dispersion of Flagship satellites, that is computed with a more sophisticated scheme aimed at reproducing in detail the velocity dispersion of galaxies in cluster. Its value after calibration is $f_{\rm v} = 0.7$.

`Pinocchio` and N-body halo masses are not expected to be equivalent, both because `Pinocchio`'s claimed accuracy is $\sim 5\%$

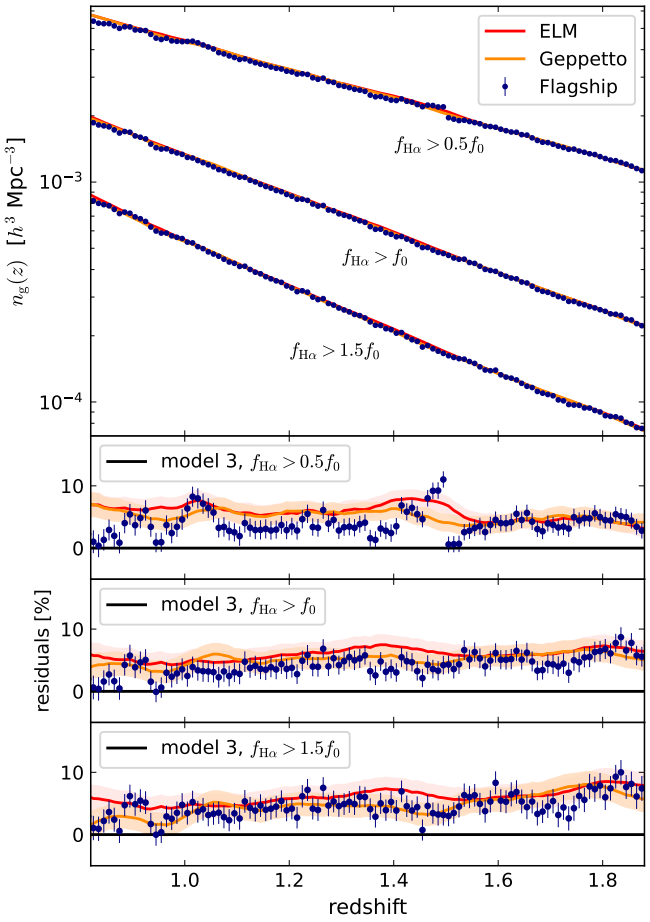

Fig. 6. Comoving number density n_g, in h^3 Mpc^{-3}, of galaxies as a function of redshift for Geppetto (orange), and EuclidLargeMocks (red) catalogues, compared with the Flagship catalogue (dark blue points). For the sets of `Pinocchio` catalogues we give the average number density as a continuous line and its standard deviation as a shaded area, while Flagship values, reported as circles, are assigned the EuclidLargeMocks variance, reported as an errorbar. We show in the top panel results for three limiting fluxes, $f_{H\alpha} > 0.5 f_0$, f_0, and $1.5 f_0$. The other panels give, for each limiting flux, the residuals $(n_g/n_{\text{model 3}} - 1)$ in percent of simulated number densities with respect to model 3 from Pozzetti et al. (2016). The black lines denote the zero value of the residuals.

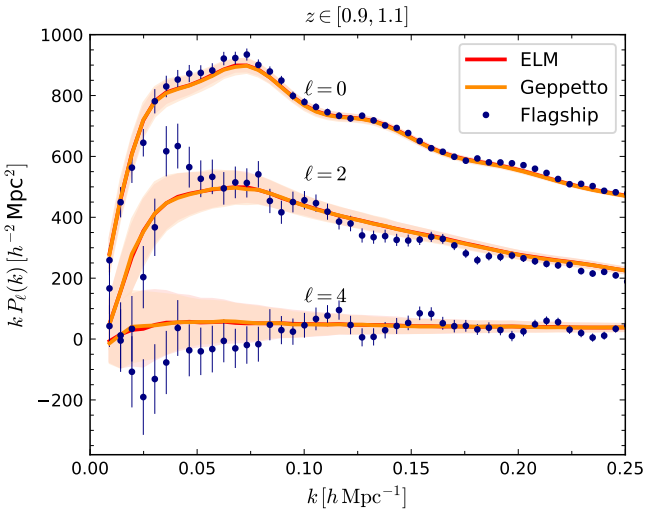

Fig. 7. Power spectrum multipoles of galaxies with $f_{H\alpha} > f_0$, for Flagship (blue points), the first 1000 Geppetto mocks (orange lines and shaded areas), and the 1000 EuclidLargeMocks (ELM, red lines and shaded areas). We show here the first redshift bin, $z \in [0.9, 1.1]$. For the `Pinocchio` sets we report the average measurement as a line and its variance as a shaded area; Geppetto and ELM lines are very similar and hard to distinguish. The Flagship points are assigned an errorbar equal to the standard deviation of the EuclidLargeMocks.

in the halo mass function and because it has been calibrated to reproduce haloes found with the friends-of-friends algorithm, not those found with the `Rockstar` halo finder (Behroozi et al. 2013) used in the Flagship simulation (see Euclid Collaboration: Castro et al. 2023 for a discussion of the differences among halo finders). Moreover, `Pinocchio` is able to reproduce the linear halo bias to within a few percent (Paranjape et al. 2013; Munari et al. 2017b), and this difference in clustering amplitude can be absorbed by halo calibration. This was done in Oddo et al. (2020), where an optimal reproduction of the halo bispectrum and its covariance was achieved by calibrating the mass cuts of simulated and `Pinocchio` haloes so as to have the same amplitude of the power spectrum. To apply this HOD to our halo lightcones, we calibrate halo masses with a clustering matching (CM) procedure, described in detail in the Appendix, and different from the abundance matching (AM) procedure that would guarantee consistency in the galaxy number density. Once the clustering level is matched, matching of number density is achieved by multiplying the HOD curves N_{cen} and N_{sat} by the ratio of a target number density and the measured one. Indeed, as long as the largest value of N_{cen} remains below unity, these curves can be shifted up or down with no effect on the clustering amplitude (apart the obviously different shot noise level), as the changing amplitude leads to a different sampling of the parent halo distribution, keeping the ratios of haloes of different mass constant.

The number density of the Flagship catalogue was calibrated to reproduce model 3 of Pozzetti et al. (2016), achieving an accuracy of roughly 5%, larger than the expected sample variance though not larger than the observational uncertainty. As a target number density we then used that of the Flagship catalogue. We fitted the ratio of the number density obtained with CM, averaged over 100 mocks, and the Flagship density with a 4th-order polynomial, and used this fitted ratio to modulate the two HOD curves as a function of redshift. We give the resulting number density n_g in Fig. 6, compared with that of the Flagship mock on the same 30° radius footprint and with model 3, for three flux limits, $f_{H\alpha} > 0.5 f_0$, f_0, and $1.5 f_0$, amounting to 1, 2 and 3×10^{-16} erg s^{-1} cm^{-2}. For the Geppetto and EuclidLargeMocks sets we report the standard deviation of the number density as a shaded area; the same standard deviation (from the EuclidLargeMocks) is assigned to the Flagship and shown as an errorbar. The three lower panels report the difference of the number densities with respect to model 3 predictions (not reported in the above panel) for the three flux limits. The Flagship mock fluctuates around the Geppetto and EuclidLargeMocks average values in a way that is compatible with the expected standard deviation, with some difference at the lowest redshift; at low flux the little blips in the Flagship number density, at $z = 1$ and 1.5, is due to the issue connected to the interpolation of internal extinction mentioned in Sect. 5.

We conclude this section by arguing that the small differences in number density and halo clustering between N-body and `Pinocchio` haloes found in Sect. 4 are not relevant in practice as they can be calibrated away. The aim of these simulation sets is to provide the covariance of clustering measurements of

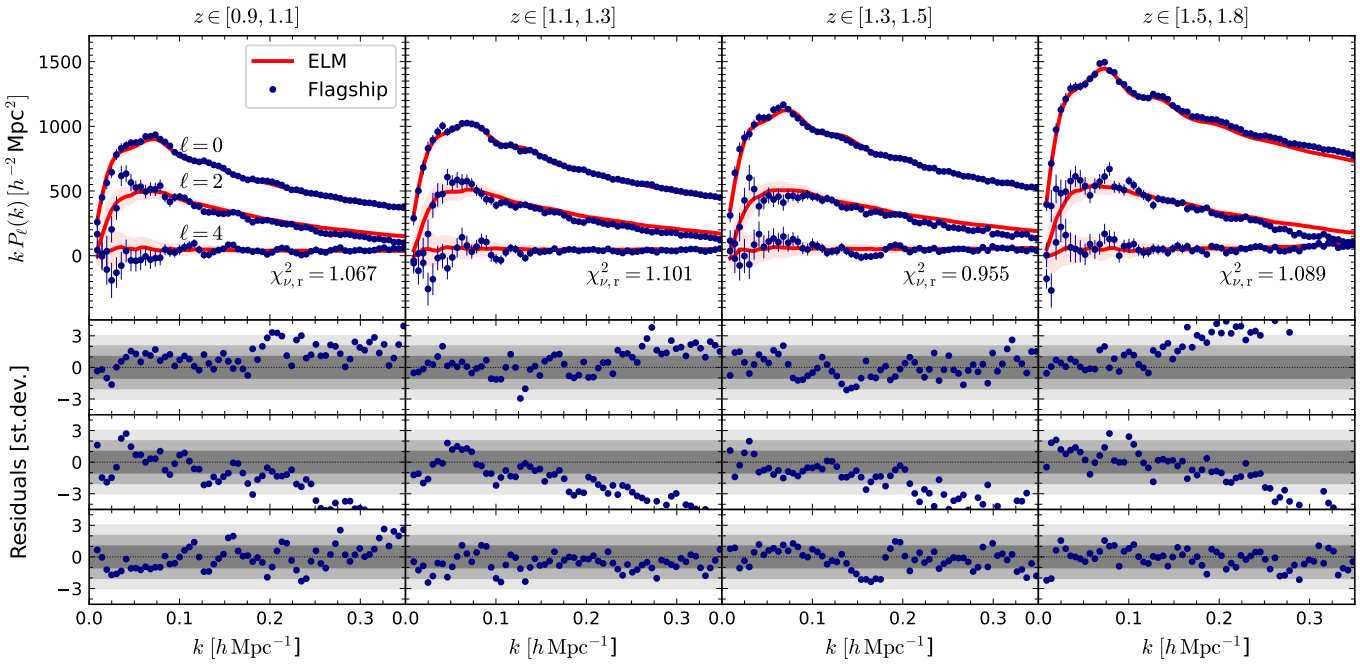

Fig. 8. Power spectrum multipoles of galaxies with $f_{\mathrm{H}\alpha} > f_0$, for Flagship (blue points), and EuclidLargeMocks (red lines), in the four redshift bins, as indicated at the top of each column of panels. Lower panels show the residuals of the Flagship measurements (from top: monopole, quadrupole, hexadecapole) with respect to the average EuclidLargeMocks in units of the standard deviation of the latter.

an observed galaxy sample, and this aim is achieved by calibrating the HOD (or in general the algorithm used to populate haloes with galaxies) to reproduce its properties. Suppose that one calibrates an HOD on the galaxy number density and 2-point clustering amplitude, and performs this calibration independently using an N-body and a `Pinocchio` halo light-cone. The small differences in halo properties will result in slightly different HOD parameters, with virtually identical results on the calibrated quantities. Such small differences will have negligible impact on the clustering covariance, whose leading terms are related to the amplitude of the power spectrum, that is matched to the data. Moreover, the accuracy requirement on the covariance is weaker than the one on the average measurement, so percent-level effects will likely be negligible. At the same time, given that we are not predictive in how a specific galaxy sample populates DM haloes, we can consider the HOD as a nuisance for cosmological parameter inference; as long as we are able to effectively marginalise over nuisance parameters, the details of the HOD will be immaterial to the final results.

6. Validation

In this section we present the validation tests that we have performed to demonstrate that the mocks can faithfully represent the covariance of the Flagship galaxy mock. Most effort is devoted to test the EuclidLargeMocks set, that is not affected by the issue of replications (Sect. 4.2), but given the very consistent results of the code in the two configurations one can safely extend the validity of the tests to the Geppetto set, at least on the scales well sampled by the simulation box of $1.2\,h^{-1}$ Gpc. In all cases the Flagship is projected on a 30° radius circle to have the same footprint as the other mock catalogues.

We first test the results for the galaxy power spectrum at the nominal flux limit f_0, for both the Geppetto and the EuclidLargeMocks sets. We use all the 1000 EuclidLargeMocks and the first 1000 Geppetto mocks, the binning of the power spectrum is at 6 fundamental modes to suppress the beating of Geppetto measurements due to replications. We show in Fig. 7 the first three even multipoles of the first redshift bin, $z \in [0.9, 1.1]$. The power spectra of Geppetto and EuclidLargeMocks are so similar that it is hard to distinguish them in the figure, both in the mean and in the standard deviation; in particular, a (very) close analysis of the figure shows that the standard deviation of the Geppetto mocks is some $\sim 5\%$ smaller than that of the EuclidLargeMocks, and this is broadly consistent with the missing super-sample covariance in the smaller simulated volume, while adding together simulations performed with different code versions does not seem to contribute. Both average power spectra reproduce very closely the measurement of the Flagship mock in the same footprint; for this summary statistics Flagship is consistent with being a realisation drawn from the set, the largest discrepancy being a flare in the Flagship quadrupole on large scales that is however consistent with being a statistical fluctuation.

To give a more comprehensive view of the results, Figs. 8 and 9 show the power spectrum and 2-point correlation function multipoles of the galaxy catalogue cut at the fiducial flux limit, $f_{\mathrm{H}\alpha} > f_0$, for the four redshift bins bounded by $[0.9, 1.1, 1.3, 1.5, 1.8]$, for Flagship and for the 1000 EuclidLargeMocks. Each column of panels reports the value of the reduced chi-squared, χ^2_{r} relative to that redshift bin, computed for $k < 0.2\,h\,\mathrm{Mpc}^{-1}$ (power spectrum) and r between 50 and 150 h^{-1} Mpc (two-point correlation function), using the three multipoles of the Flagship mock as data vector and the EuclidLargeMocks as model and covariance (see Fig 11 below). In all cases the agreement is excellent: for the power spectrum, and in the wavelength range $k < 0.2\,h\,\mathrm{Mpc}^{-1}$ typically used in standard inference of cosmological parameters, the Flagship measurement is consistent with being drawn from a realisation of the EuclidLargeMocks, while at higher wave-numbers the EuclidLargeMocks tend to overestimate its quadrupole, and underestimate the monopole in

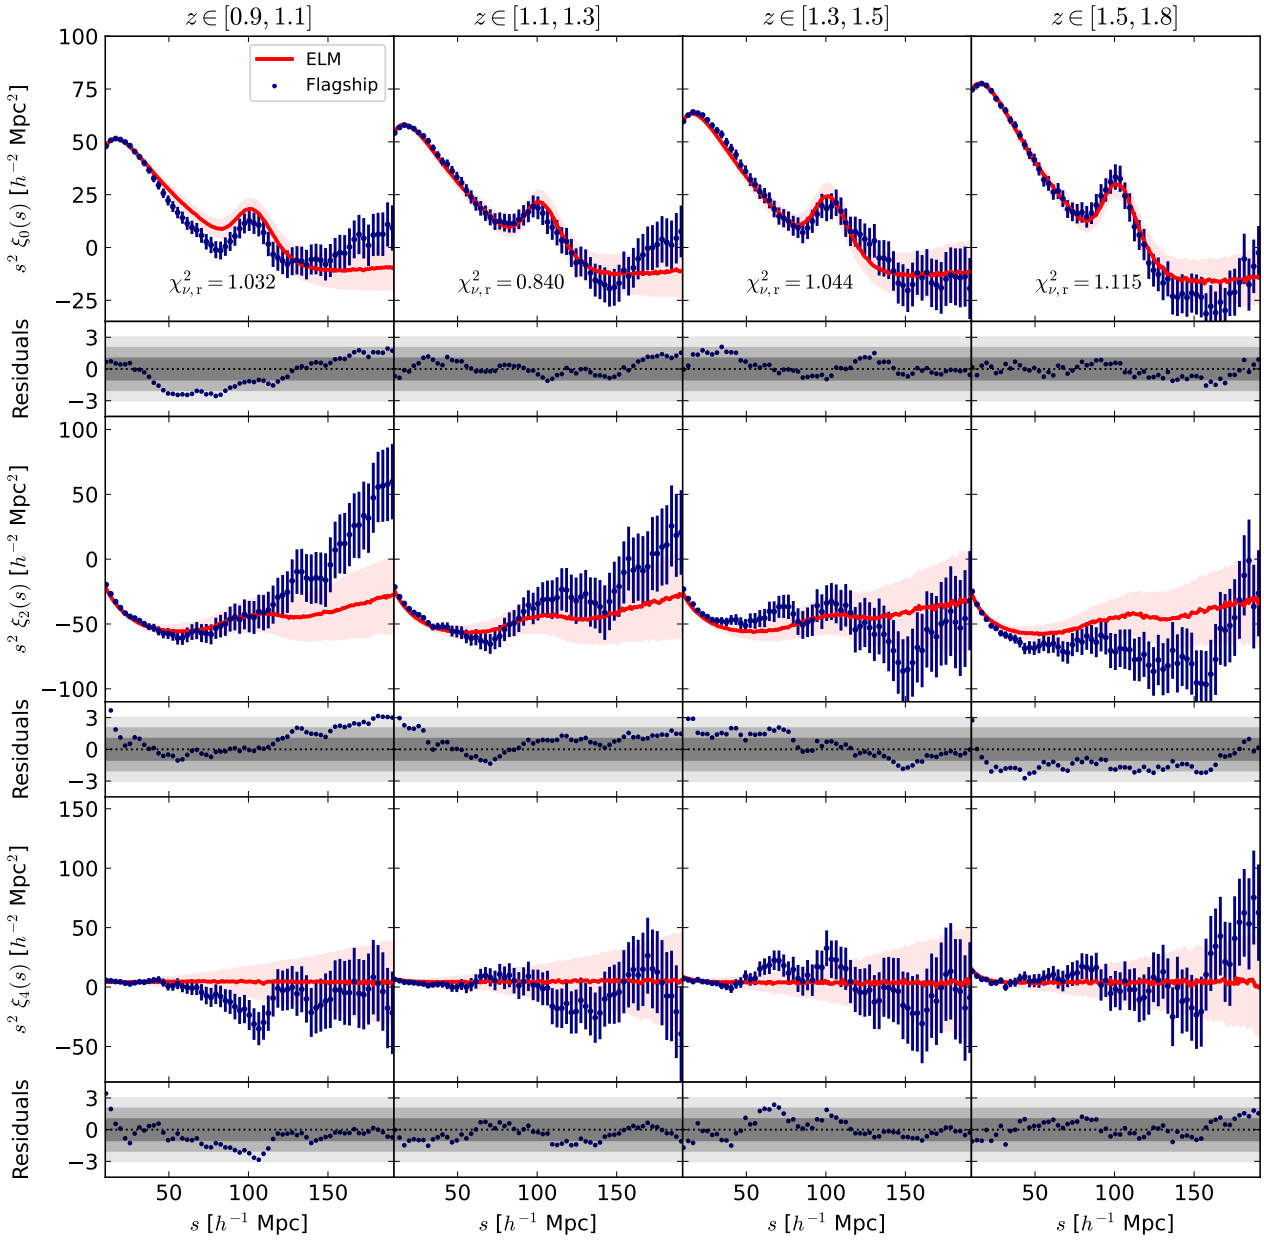

Fig. 9. Two-point correlation function multipoles of galaxies with $f_{\mathrm{H}\alpha} > f_0$, for Flagship (blue points), and EuclidLargeMocks (red lines), in the four redshift bins as indicated at the top of each column of panels. Lower panels show the residuals of the Flagship measurements with respect to the average EuclidLargeMocks in units of the standard deviation of the latter.

the last redshift bin by 7%. As mentioned in the Appendix, we calibrated the CM relation by requiring an accuracy better than 10% on the scales dominated by the one-halo term, so this difference can be decreased by improving the HOD calibration. The overestimation of the quadrupole was already noticed by Munari et al. (2017b) and, not only for the `Pinocchio` code, by Blot et al. (2019), and will need a specific mitigation strategy if the scale cut is going to be more aggressive than $k < 0.3\,h\,\mathrm{Mpc}^{-1}$. These differences are less visible in the 2-point correlation function (Fig. 9), where the (much more correlated) measurements of the Flagship mock are always consistent with being drawn from the distribution of the EuclidLargeMocks. Here the small apparent shift in the baryonic-acoustic-oscillation (BAO) peak visible in the monopole is consistent with sample variance; this is evident from the low χ^2_{r} values, but we also checked that fitting the BAO peak for the average measurement and for the Flagship mock provides very consistent results (Euclid Collaboration: Sarpa et al., in prep.). As noticed above, in both estimators the Flagship measurements show large and sometimes significant ($\sim 3\sigma$) flares of the quadrupole on large scales.

The accuracy with which luminosity-dependent galaxy bias is reproduced is shown in Fig. 10. We report here the monopole

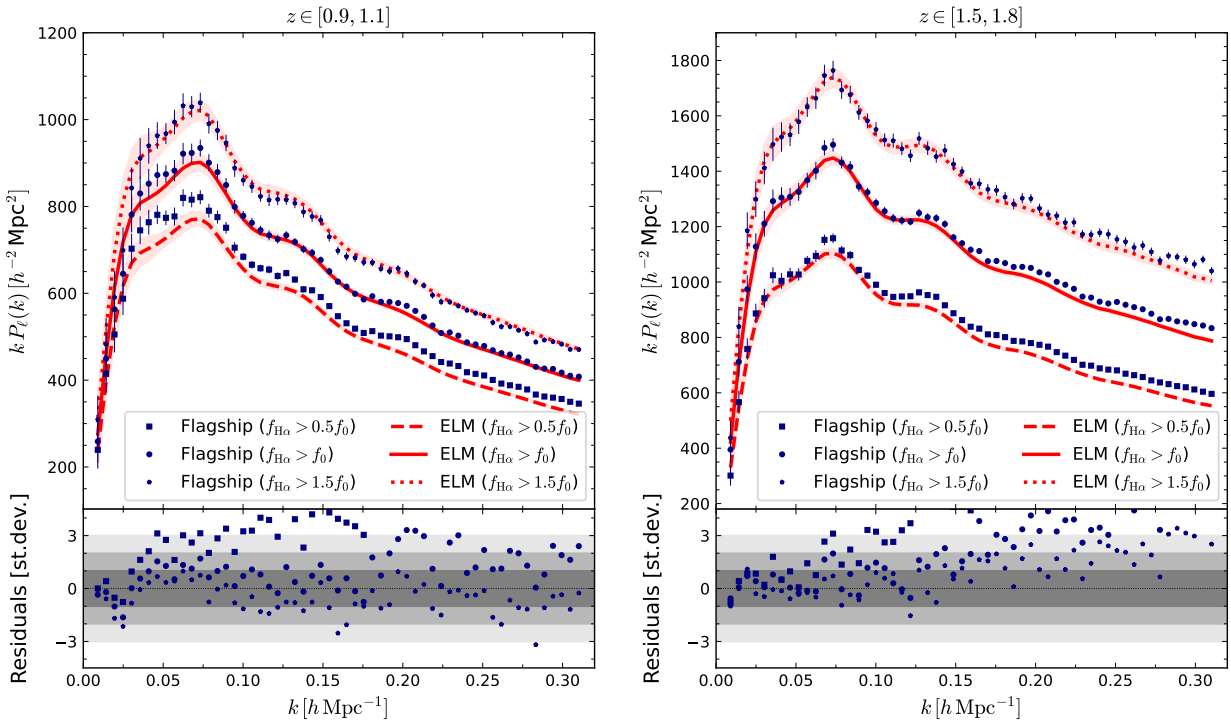

Fig. 10. Power spectrum monopole of galaxies in the Flagship mock and in the EuclidLargeMocks, for $f_{\mathrm{H}\alpha} > 0.5 f_0$, f_0, and $1.5 f_0$, and for $z \in [0.9, 1.1]$ (left panel) and $z \in [1.5, 1.8]$. Lower panels show the residuals of the Flagship measurements for the three flux limits with respect to the average EuclidLargeMocks in units of the standard deviation of the latter.

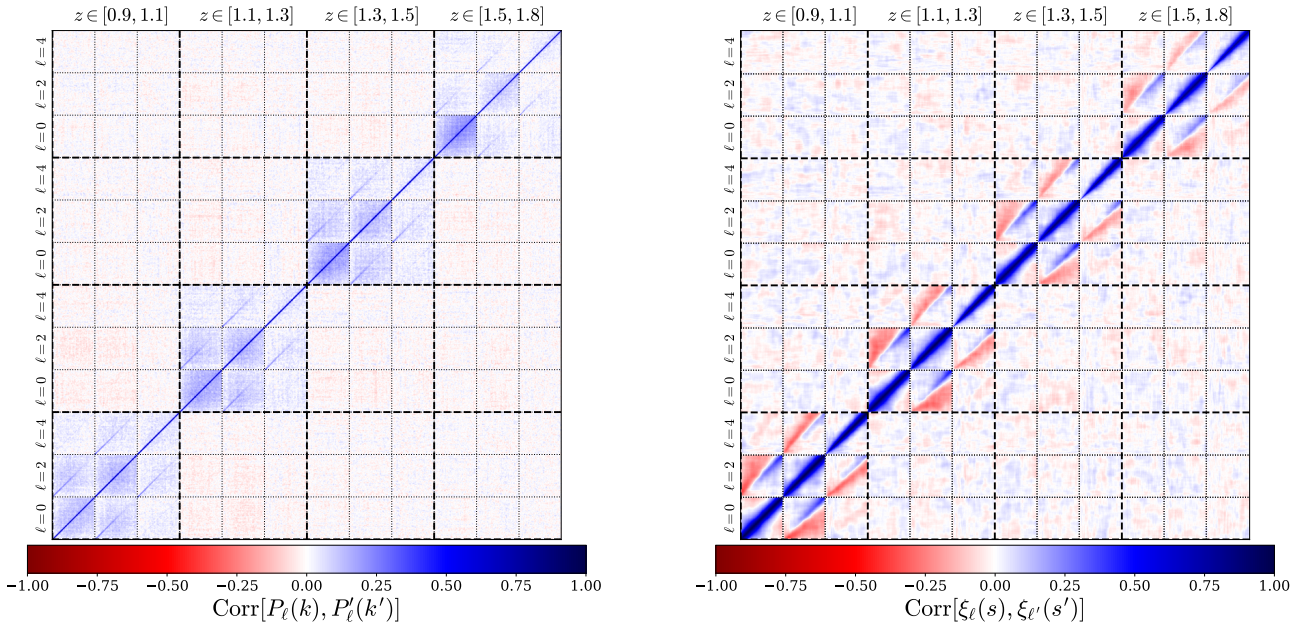

Fig. 11. Correlation matrices of the first three even multipoles of the power spectrum (left) and of the 2-point correlation function (right), for all the redshift bins.

of the power spectrum for flux cuts $f_{\mathrm{H}\alpha} > 0.5 f_0$, f_0, and $1.5 f_0$, in the first and last redshift bin ($z \in [0.9, 1.1]$ and $z \in [1.5, 1.8]$). Luminosity-dependent bias is well reproduced, although at the lowest flux we notice an underestimation of the monopole that would require a more sophisticated implementation of CM to be removed; the number of galaxies detected in the actual survey at such low fluxes is expected to be small, so this discrepancy is unlikely to be a significant problem.

These measurements allow us to construct a numerical covariance matrix that takes into account in full detail the geometry of the survey. Figure 11 shows the correlation matrices of all the multipoles of the power spectrum for all the redshift bins, on

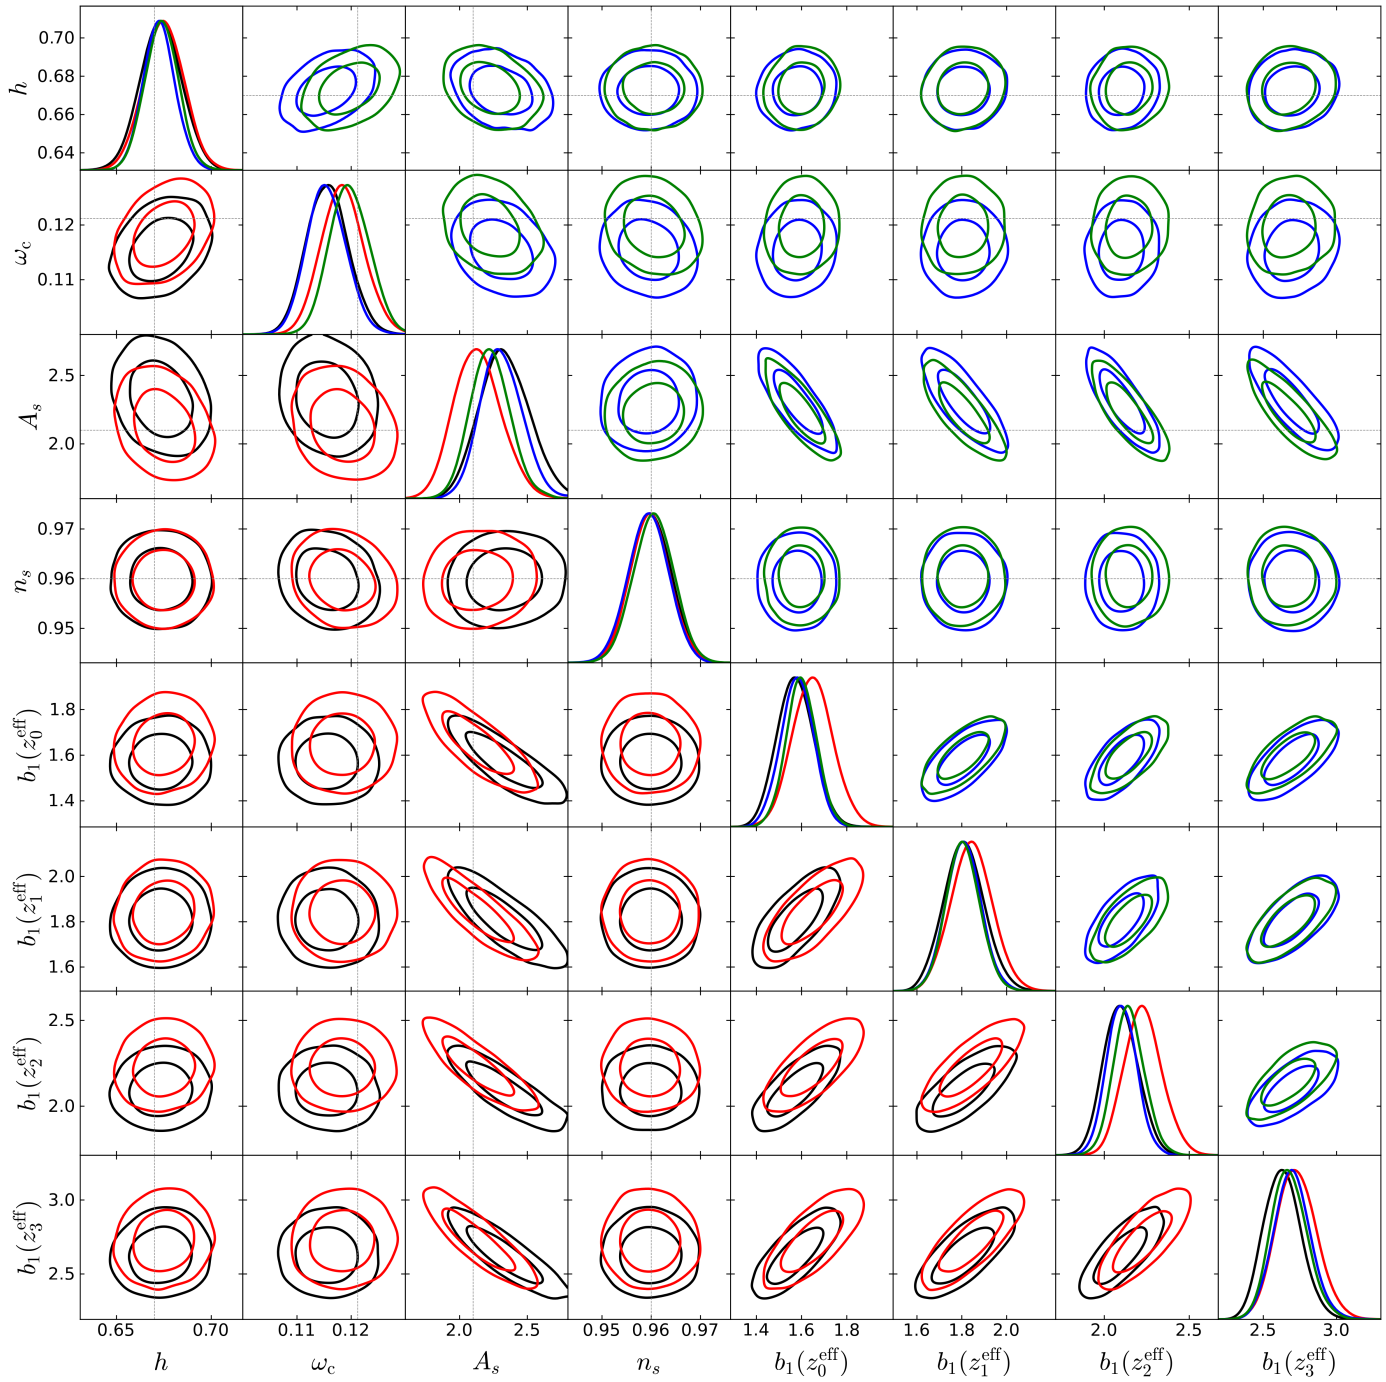

Fig. 12. Posteriors of parameter inference obtained by fitting the Flagship (black and blue lines) and one of the EuclidLargeMocks (red and green lines) with cuts at $k < 0.2\,h\,\mathrm{Mpc}^{-1}$ (lower-left corner with red and black lines) and $k < 0.3\,h\,\mathrm{Mpc}^{-1}$ (upper-right corner with blue and green lines).

the left for the power spectrum and on the right for the 2-point correlation function. In this case the correlation of different redshift bins is very low, as expected, showing that the problem of replications is absent in the EuclidLargeMocks.

We use these covariances in one example of inference, similar to what is done in Euclid Collaboration: P. Monaco et al. (in prep.). We fit the Flagship power spectrum multipoles across four redshift bins using a model based on the effective field theory of large-scale structure (see Cabass et al. 2023, for a recent review). Model and likelihood evaluations are performed using the Gaussian process emulator `Comet` (Eggemeier et al. 2022), while posterior estimation is conducted with `NAUTILUS` (Lange 2023). We fit the monopole and quadrupole of the power spectrum, binned in k with an initial bin $k_\mathrm{i} = 6\,k_\mathrm{F} \sim 5.4 \times 10^{-3}\,h\,\mathrm{Mpc}^{-1}$ and bin size $\Delta k = 6\,k_\mathrm{F}$. The maximum wavenumber k_max is varied across different fits. The model includes five cosmological parameters: $\{h, \omega_\mathrm{c}, A_\mathrm{s}, n_\mathrm{s}, \omega_\mathrm{b}\}$, where $\omega = \Omega\,h^2$ and c stands for cold dark matter, b for baryons, while A_s and n_s are the normalisation and the spectral slope of primordial perturbations. We adopt wide uniform priors for all parameters except for the spectral index n_s, which follows a Gaussian prior $\mathcal{N}(0.96, 0.041)$, and the baryon density parameter, which follows $\mathcal{N}(2.218 \times 10^{-2}, 0.055 \times 10^{-2})$. Additionally, we vary three

bias parameters $\{b_1, b_2, \gamma_{21}\}$ for each redshift bin and analytically marginalise over the counter- and shot noise terms.

We also use one of the measurements from the EuclidLargeMocks as a data vector and rerun the chains to get another set of parameter values. Figure 12 shows the resulting parameter posteriors (with the exception of ω_b, that is prior-dominated), where the black and red lines in the lower left triangle give results obtained respectively with Flagship and EuclidLargeMocks using a conservative cut of $k < 0.2\,h\,\mathrm{Mpc}^{-1}$, while the blue and green lines in the upper-right triangle give results obtained using a more aggressive $k < 0.3\,h\,\mathrm{Mpc}^{-1}$ cut. In all cases the contours are consistent among themselves to within 1σ, while inferred values of cosmological parameters are consistent with the true ones. This test demonstrates that the EuclidLargeMocks can be safely used to do parameter inference for *Euclid* galaxy clustering.

7. Prospects and conclusions

We presented the largest collection of simulated dark matter haloes in the light-cone ever produced, and a corresponding set of 3500+1000 galaxy mock catalogues of the *Euclid* spectroscopic sample that cover each an area of 2763 deg^2, larger than the planned area of DR1 in Euclid Collaboration: Scaramella et al. (2022) and roughly one fifth of the final EWS. The simulations were produced with the `Pinocchio` code for the generation of approximate catalogues of dark matter haloes, in a numerical effort of $\sim 5\,000\,000$ core-h that produced ~ 300 TB of halo catalogues. The galaxy mock catalogues were obtained by extracting an HOD from the Flagship mock catalogue described in the Flagship paper, and calibrating it to absorb the little differences in the clustering and number density of `Pinocchio` and simulated haloes.

We validated these sets of galaxy mock catalogues by comparing number densities, power spectra and 2-point correlation functions, of haloes and galaxies with different halo mass or line flux cuts and in four redshift bins from $z = 0.9$ to 1.8, with the same measurements of the Flagship spectroscopic catalogue. We found that the spectroscopic sample drawn from the Flagship mock catalogue is consistent with being a realisation drawn from one of our sets, and this is testified by the reasonable values of χ^2_r obtained using the Flagship multipoles as a data vector and the EuclidLargeMocks for the model and the covariance. We identified two main limitations of our mocks, in particular the Geppetto set is affected by the relatively small box size, replicated several times to fill the light-cone volume. This results in spurious oscillations of the galaxy power spectrum that can be mitigated by performing the measurements in units of the fundamental mode of the simulation box. This problem does not affect the EuclidLargeMocks, where the survey volume is contained in the simulation box. Also, For both sets we noticed an overestimation of the power spectrum quadrupole at $k > 0.2\,h\,\mathrm{Mpc}^{-1}$. Conversely, we noticed a flare of the quadrupole of the Flagship galaxies on the largest scale that is consistent with being a statistical fluctuations, while some blips in the Flagship galaxy number density are due to a known issue in the interpolation of internal extinction needed to compute the H$\,\alpha$ line flux.

As a test of this consistency, we verified that using the power spectrum of a single EuclidLargeMocks or the Flagship mock for inferring parameters in a maximum likelihood fit, where the numerical covariance is obtained from the EuclidLargeMocks, gives consistent posteriors for all the cosmological and nuisance parameters. This remains true even when pushing the scale cut to $k < 0.3\,h\,\mathrm{Mpc}^{-1}$, showing that the overestimation of the quadrupole does not bias cosmological inference. This fully demonstrates the usability of these mock catalogues for analysing *Euclid*'s DR1.

The EuclidLargeMocks are being extensively used in many *Euclid* preparation papers, especially in the papers of the Observational Systematics Key Project presented by Euclid Collaboration: P. Monaco et al. (in prep.), and in the papers of the Organisation Unit for Level-3 products, in particular for testing the estimator of the power spectrum (Euclid Collaboration: Salvalaggio et al., in prep.). Simulated catalogues of dark matter haloes on the light-cone are going to me made available at `https://adlibitum.oats.inaf.it/pierluigi.monaco/euclid_mocks.html` by August 2026, Euclid mock catalogs will be released together with the Flagship mock catalogue.

Acknowledgements. The Euclid Consortium acknowledges the European Space Agency and a number of agencies and institutes that have supported the development of *Euclid*, in particular the Agenzia Spaziale Italiana, the Austrian Forschungsförderungsgesellschaft funded through BMK, the Belgian Science Policy, the Canadian Euclid Consortium, the Deutsches Zentrum für Luft- und Raumfahrt, the DTU Space and the Niels Bohr Institute in Denmark, the French Centre National d'Etudes Spatiales, the Fundação para a Ciência e a Tecnologia, the Hungarian Academy of Sciences, the Ministerio de Ciencia, Innovación y Universidades, the National Aeronautics and Space Administration, the National Astronomical Observatory of Japan, the Netherlandse Onderzoekschool Voor Astronomie, the Norwegian Space Agency, the Research Council of Finland, the Romanian Space Agency, the State Secretariat for Education, Research, and Innovation (SERI) at the Swiss Space Office (SSO), and the United Kingdom Space Agency. A complete and detailed list is available on the *Euclid* web site (`www.euclid-ec.org`). This work has made use of CosmoHub, developed by PIC (maintained by IFAE and CIEMAT) in collaboration with ICE-CSIC. CosmoHub received funding from the Spanish government (MCIN/AEI/10.13039/501100011033), the EU NextGeneration/PRTR (PRTR-C17.I1), and the Generalitat de Catalunya. This paper has been supported by: the Fondazione ICSC, Spoke 3 Astrophysics and Cosmos Observations. National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) Project ID CN_00000013 'Italian Research Center on High-Performance Computing, Big Data and Quantum Computing' funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di 'campioni nazionali di R&S (M4C2-19)' - Next Generation EU (NGEU); by the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, Call for tender No. 1409 published on 14.9.2022 by the Italian Ministry of University and Research (MUR), funded by the European Union – NextGenerationEU– Project Title 'Space-based cosmology with Euclid: the role of High-Performance Computing' – CUP J53D23019100001 - Grant Assignment Decree No. 962 adopted on 30/06/2023 by the Italian Ministry of Ministry of University and Research (MUR). Computing time was obtained from CINECA ISCRA-B grant 'EuMocks', from INFN and from the Pleiadi system of INAF. (Taffoni et al. 2020; Bertocco et al. 2020)

References

Balaguera-Antolínez, A., Kitaura, F.-S., Alam, S., et al. 2023, A&A, 673, A130
Balaguera-Antolínez, A., Kitaura, F.-S., Pellejero-Ibáñez, M., et al. 2019, MNRAS, 483, L58
Baratta, P., Bel, J., Gouyou Beauchamps, S., et al. 2023, A&A, 673, A1
Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, ApJ, 762, 109
Bertocco, S., Goz, D., Tornatore, L., et al. 2020, in Astronomical Society of the Pacific Conference Series, Vol. 527, Astronomical Data Analysis Software and Systems XXIX, ed. R. Pizzo, E. R. Deul, J. D. Mol, J. de Plaa, & H. Verkouter, 303
Bianchi, D., Gil-Marín, H., Ruggeri, R., et al. 2015, MNRAS, 453, L11
Blot, L., Crocce, M., Sefusatti, E., et al. 2019, MNRAS, 485, 2806
Bond, J. R. & Myers, S. T. 1996, ApJS, 103, 1
Cabass, G., Ivanov, M. M., Lewandowski, M., Mirbabayi, M., & Simonović, M. 2023, Phys. Dark Univ., 40, 101193
Chuang, C.-H., Kitaura, F.-S., Prada, F., et al. 2015a, MNRAS, 446, 2621
Chuang, C.-H., Zhao, C., Prada, F., et al. 2015b, MNRAS, 452, 686
Colavincenzo, M., Monaco, P., Sefusatti, E., et al. 2017, JCAP, 03, 052
Colavincenzo, M., Sefusatti, E., Monaco, P., et al. 2019, MNRAS, 482, 4883
Crocce, M., Pueblas, S., & Scoccimarro, R. 2006, MNRAS, 373, 369
DESI Collaboration: Adame, A. G., Aguilar, J., et al. 2024, arXiv:2404.03002
Diemer, B. & Joyce, M. 2019, ApJ, 871, 168

Ding, S., Lavaux, G., & Jasche, J. 2024, A&A, 690, A236
Doré, O., Bock, J., Ashby, M., et al. 2014, arXiv:1412.4872
Eggemeier, A., Camacho-Quevedo, B., Pezzotta, A., et al. 2022, MNRAS, 519, 2962
Euclid Collaboration: Castander, F., Fosalba, P., Stadel, J., et al. 2025, A&A, 697, A5
Euclid Collaboration: Castro, T., Fumagalli, A., Angulo, R. E., et al. 2023, A&A, 671, A100
Euclid Collaboration: Cropper, M., Al-Bahlawan, A., Amiaux, J., et al. 2025, A&A, 697, A2
Euclid Collaboration: Jahnke, K., Gillard, W., Schirmer, M., et al. 2025, A&A, 697, A3
Euclid Collaboration: Lepori, F., Tutusaus, I., Viglione, C., et al. 2022, A&A, 662, A93
Euclid Collaboration: Mellier, Y., Abdurro'uf, Acevedo Barroso, J., et al. 2025, A&A, 697, A1
Euclid Collaboration: Pezzotta, A., Moretti, C., Zennaro, M., et al. 2024, A&A, 687, A216
Euclid Collaboration: Scaramella, R., Amiaux, J., Mellier, Y., et al. 2022, A&A, 662, A112
Feng, Y., Chu, M.-Y., Seljak, U., & McDonald, P. 2016, MNRAS, 463, 2273
Frigo, M. & Johnson, S. G. 2012, FFTW: Fastest Fourier Transform in the West, Astrophysics Source Code Library, record ascl:1201.015
Fumagalli, A., Biagetti, M., Saro, A., et al. 2022, JCAP, 12, 022
Fumagalli, A., Saro, A., Borgani, S., et al. 2021, A&A, 652, A21
Howlett, C., Manera, M., & Percival, W. J. 2015, Astronomy and Computing, 12, 109
Kitaura, F.-S., Gil-Marín, H., Scóccola, C. G., et al. 2015, MNRAS, 450, 1836
Lange, J. U. 2023, MNRAS, 525, 3181
Lippich, M., Sánchez, A. G., Colavincenzo, M., et al. 2019, MNRAS, 482, 1786
LSST Science Collaboration: Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
Manera, M., Scoccimarro, R., Percival, W. J., et al. 2013, MNRAS, 428, 1036
Monaco, P. 2016, Galaxies, 4, 53
Monaco, P., Di Dio, E., & Sefusatti, E. 2019, JCAP, 04, 023
Monaco, P., Sefusatti, E., Borgani, S., et al. 2013, MNRAS, 433, 2389
Monaco, P., Theuns, T., & Taffoni, G. 2002, MNRAS, 331, 587
Munari, E., Monaco, P., Koda, J., et al. 2017a, JCAP, 07, 050
Munari, E., Monaco, P., Sefusatti, E., et al. 2017b, MNRAS, 465, 4658
Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, ApJ, 462, 563
Oddo, A., Rizzo, F., Sefusatti, E., et al. 2021, JCAP, 11, 038
Oddo, A., Sefusatti, E., Porciani, C., et al. 2020, JCAP, 03, 056
Pandey, S., Lanusse, F., Modi, C., et al. 2024, arXiv:2409.11401
Paranjape, A., Sefusatti, E., Chan, K. C., et al. 2013, MNRAS, 436, 449
Planck Collaboration: Aghanim, N., Akrami, Y., et al. 2020, A&A, 641, A6
Pozzetti, L., Hirata, C. M., Geach, J. E., et al. 2016, A&A, 590, A3
Rizzo, F., Moretti, C., Pardede, K., et al. 2023, JCAP, 01, 031
SKA Cosmology Group: Bacon, D. J., Battye, R. A., et al. 2020, PASA, 37, e007
Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757
Springel, V. 2005, MNRAS, 364, 1105
Springel, V., White, S. D. M., Jenkins, A., et al. 2005, Nature, 435, 629
Stein, G., Alvarez, M. A., & Bond, J. R. 2019, MNRAS, 483, 2236
Taffoni, G., Becciani, U., Garilli, B., et al. 2020, in Astronomical Society of the Pacific Conference Series, Vol. 527, Astronomical Data Analysis Software and Systems XXIX, ed. R. Pizzo, E. R. Deul, J. D. Mol, J. de Plaa, & H. Verkouter, 307
Tassev, S., Zaldarriaga, M., & Eisenstein, D. J. 2013, JCAP, 06, 036
Villaescusa-Navarro, F., Hahn, C., Massara, E., et al. 2020, ApJS, 250, 2
Yamamoto, K., Nakamichi, M., Kamino, A., et al. 2006, PASJ, 58, 93

[1] Dipartimento di Fisica - Sezione di Astronomia, Università di Trieste, Via Tiepolo 11, 34131 Trieste, Italy
[2] INAF-Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, 34143 Trieste, Italy
[3] INFN, Sezione di Trieste, Via Valerio 2, 34127 Trieste TS, Italy
[4] ICSC - Centro Nazionale di Ricerca in High Performance Computing, Big Data e Quantum Computing, Via Magnanelli 2, Bologna, Italy
[5] Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain
[6] Dipartimento di Fisica, Università degli studi di Genova, and INFN-Sezione di Genova, via Dodecaneso 33, 16146, Genova, Italy
[7] SISSA, International School for Advanced Studies, Via Bonomea 265, 34136 Trieste TS, Italy
[8] IFPU, Institute for Fundamental Physics of the Universe, via Beirut 2, 34151 Trieste, Italy
[9] INFN-Padova, Via Marzolo 8, 35131 Padova, Italy
[10] Johns Hopkins University 3400 North Charles Street Baltimore, MD 21218, USA
[11] California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, USA
[12] INAF-IASF Milano, Via Alfonso Corti 12, 20133 Milano, Italy
[13] Institut d'Estudis Espacials de Catalunya (IEEC), Edifici RDIT, Campus UPC, 08860 Castelldefels, Barcelona, Spain
[14] Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Avenida Complutense 40, 28040 Madrid, Spain
[15] Port d'Informació Científica, Campus UAB, C. Albareda s/n, 08193 Bellaterra (Barcelona), Spain
[16] Aix-Marseille Université, CNRS, CNES, LAM, Marseille, France
[17] Institut d'Astrophysique de Paris, UMR 7095, CNRS, and Sorbonne Université, 98 bis boulevard Arago, 75014 Paris, France
[18] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109, USA
[19] Kavli Institute for the Physics and Mathematics of the Universe (WPI), University of Tokyo, Kashiwa, Chiba 277-8583, Japan
[20] Dipartimento di Fisica e Astronomia "G. Galilei", Università di Padova, Via Marzolo 8, 35131 Padova, Italy
[21] Waterloo Centre for Astrophysics, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada
[22] Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada
[23] Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada
[24] INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy
[25] INFN-Sezione di Genova, Via Dodecaneso 33, 16146, Genova, Italy
[26] Minnesota Institute for Astrophysics, University of Minnesota, 116 Church St SE, Minneapolis, MN 55455, USA
[27] Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA
[28] ESAC/ESA, Camino Bajo del Castillo, s/n., Urb. Villafranca del Castillo, 28692 Villanueva de la Cañada, Madrid, Spain
[29] INAF-Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Via Piero Gobetti 93/3, 40129 Bologna, Italy
[30] Dipartimento di Fisica e Astronomia, Università di Bologna, Via Gobetti 93/2, 40129 Bologna, Italy
[31] INFN-Sezione di Bologna, Viale Berti Pichat 6/2, 40127 Bologna, Italy
[32] Institut de Physique Théorique, CEA, CNRS, Université Paris-Saclay 91191 Gif-sur-Yvette Cedex, France
[33] Dipartimento di Fisica, Università di Genova, Via Dodecaneso 33, 16146, Genova, Italy
[34] Department of Physics "E. Pancini", University Federico II, Via Cinthia 6, 80126, Napoli, Italy
[35] INAF-Osservatorio Astronomico di Capodimonte, Via Moiariello 16, 80131 Napoli, Italy
[36] Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto, CAUP, Rua das Estrelas, PT4150-762 Porto, Portugal
[37] Faculdade de Ciências da Universidade do Porto, Rua do Campo de Alegre, 4150-007 Porto, Portugal
[38] European Southern Observatory, Karl-Schwarzschild-Str. 2, 85748 Garching, Germany
[39] Dipartimento di Fisica, Università degli Studi di Torino, Via P. Giuria 1, 10125 Torino, Italy
[40] INFN-Sezione di Torino, Via P. Giuria 1, 10125 Torino, Italy
[41] INAF-Osservatorio Astrofisico di Torino, Via Osservatorio 20, 10025 Pino Torinese (TO), Italy
[42] European Space Agency/ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
[43] Institute Lorentz, Leiden University, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands

[44] Leiden Observatory, Leiden University, Einsteinweg 55, 2333 CC Leiden, The Netherlands
[45] INAF-Osservatorio Astronomico di Roma, Via Frascati 33, 00078 Monteporzio Catone, Italy
[46] INFN-Sezione di Roma, Piazzale Aldo Moro, 2 - c/o Dipartimento di Fisica, Edificio G. Marconi, 00185 Roma, Italy
[47] Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, 52056 Aachen, Germany
[48] INFN section of Naples, Via Cinthia 6, 80126, Napoli, Italy
[49] Dipartimento di Fisica e Astronomia "Augusto Righi" - Alma Mater Studiorum Università di Bologna, Viale Berti Pichat 6/2, 40127 Bologna, Italy
[50] Instituto de Astrofísica de Canarias, Vía Láctea, 38205 La Laguna, Tenerife, Spain
[51] Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK
[52] Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL, UK
[53] European Space Agency/ESRIN, Largo Galileo Galilei 1, 00044 Frascati, Roma, Italy
[54] Université Claude Bernard Lyon 1, CNRS/IN2P3, IP2I Lyon, UMR 5822, Villeurbanne, F-69100, France
[55] Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona (IEEC-UB), Martí i Franquès 1, 08028 Barcelona, Spain
[56] Institució Catalana de Recerca i Estudis Avançats (ICREA), Passeig de Lluís Companys 23, 08010 Barcelona, Spain
[57] UCB Lyon 1, CNRS/IN2P3, IUF, IP2I Lyon, 4 rue Enrico Fermi, 69622 Villeurbanne, France
[58] Departamento de Física, Faculdade de Ciências, Universidade de Lisboa, Edifício C8, Campo Grande, PT1749-016 Lisboa, Portugal
[59] Instituto de Astrofísica e Ciências do Espaço, Faculdade de Ciências, Universidade de Lisboa, Campo Grande, 1749-016 Lisboa, Portugal
[60] Department of Astronomy, University of Geneva, ch. d'Ecogia 16, 1290 Versoix, Switzerland
[61] INAF-Istituto di Astrofisica e Planetologia Spaziali, via del Fosso del Cavaliere, 100, 00100 Roma, Italy
[62] Aix-Marseille Université, CNRS/IN2P3, CPPM, Marseille, France
[63] Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191, Gif-sur-Yvette, France
[64] INFN-Bologna, Via Irnerio 46, 40126 Bologna, Italy
[65] University Observatory, LMU Faculty of Physics, Scheinerstrasse 1, 81679 Munich, Germany
[66] Max Planck Institute for Extraterrestrial Physics, Giessenbachstr. 1, 85748 Garching, Germany
[67] INAF-Osservatorio Astronomico di Padova, Via dell'Osservatorio 5, 35122 Padova, Italy
[68] Universitäts-Sternwarte München, Fakultät für Physik, Ludwig-Maximilians-Universität München, Scheinerstrasse 1, 81679 München, Germany
[69] Dipartimento di Fisica "Aldo Pontremoli", Università degli Studi di Milano, Via Celoria 16, 20133 Milano, Italy
[70] INFN-Sezione di Milano, Via Celoria 16, 20133 Milano, Italy
[71] Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, 0315 Oslo, Norway
[72] Felix Hormuth Engineering, Goethestr. 17, 69181 Leimen, Germany
[73] Technical University of Denmark, Elektrovej 327, 2800 Kgs. Lyngby, Denmark
[74] Cosmic Dawn Center (DAWN), Denmark
[75] Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany
[76] NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
[77] Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK
[78] Department of Physics and Helsinki Institute of Physics, Gustaf Hällströmin katu 2, 00014 University of Helsinki, Finland
[79] Université de Genève, Département de Physique Théorique and Centre for Astroparticle Physics, 24 quai Ernest-Ansermet, CH-1211 Genève 4, Switzerland
[80] Department of Physics, P.O. Box 64, 00014 University of Helsinki, Finland
[81] Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, Helsinki, Finland
[82] Laboratoire d'etude de l'Univers et des phenomenes eXtremes, Observatoire de Paris, Université PSL, Sorbonne Université, CNRS, 92190 Meudon, France
[83] SKA Observatory, Jodrell Bank, Lower Withington, Macclesfield, Cheshire SK11 9FT, UK
[84] Universität Bonn, Argelander-Institut für Astronomie, Auf dem Hügel 71, 53121 Bonn, Germany
[85] Dipartimento di Fisica e Astronomia "Augusto Righi" - Alma Mater Studiorum Università di Bologna, via Piero Gobetti 93/2, 40129 Bologna, Italy
[86] Department of Physics, Institute for Computational Cosmology, Durham University, South Road, Durham, DH1 3LE, UK
[87] Université Paris Cité, CNRS, Astroparticule et Cosmologie, 75013 Paris, France
[88] CNRS-UCB International Research Laboratory, Centre Pierre Binétruy, IRL2007, CPB-IN2P3, Berkeley, USA
[89] University of Applied Sciences and Arts of Northwestern Switzerland, School of Engineering, 5210 Windisch, Switzerland
[90] Institut d'Astrophysique de Paris, 98bis Boulevard Arago, 75014, Paris, France
[91] Institute of Physics, Laboratory of Astrophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, 1290 Versoix, Switzerland
[92] Telespazio UK S.L. for European Space Agency (ESA), Camino bajo del Castillo, s/n, Urbanizacion Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain
[93] Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona), Spain
[94] DARK, Niels Bohr Institute, University of Copenhagen, Jagtvej 155, 2200 Copenhagen, Denmark
[95] Space Science Data Center, Italian Space Agency, via del Politecnico snc, 00133 Roma, Italy
[96] Centre National d'Etudes Spatiales – Centre spatial de Toulouse, 18 avenue Edouard Belin, 31401 Toulouse Cedex 9, France
[97] Institute of Space Science, Str. Atomistilor, nr. 409 Măgurele, Ilfov, 077125, Romania
[98] Institut für Theoretische Physik, University of Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany
[99] Institut de Recherche en Astrophysique et Planétologie (IRAP), Université de Toulouse, CNRS, UPS, CNES, 14 Av. Edouard Belin, 31400 Toulouse, France
[100] Université St Joseph; Faculty of Sciences, Beirut, Lebanon
[101] Departamento de Física, FCFM, Universidad de Chile, Blanco Encalada 2008, Santiago, Chile
[102] Universität Innsbruck, Institut für Astro- und Teilchenphysik, Technikerstr. 25/8, 6020 Innsbruck, Austria
[103] Satlantis, University Science Park, Sede Bld 48940, Leioa-Bilbao, Spain
[104] Instituto de Astrofísica e Ciências do Espaço, Faculdade de Ciências, Universidade de Lisboa, Tapada da Ajuda, 1349-018 Lisboa, Portugal
[105] Cosmic Dawn Center (DAWN)
[106] Niels Bohr Institute, University of Copenhagen, Jagtvej 128, 2200 Copenhagen, Denmark
[107] Universidad Politécnica de Cartagena, Departamento de Electrónica y Tecnología de Computadoras, Plaza del Hospital 1, 30202 Cartagena, Spain
[108] Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands
[109] Dipartimento di Fisica e Scienze della Terra, Università degli Studi di Ferrara, Via Giuseppe Saragat 1, 44122 Ferrara, Italy

[110] Istituto Nazionale di Fisica Nucleare, Sezione di Ferrara, Via Giuseppe Saragat 1, 44122 Ferrara, Italy
[111] INAF, Istituto di Radioastronomia, Via Piero Gobetti 101, 40129 Bologna, Italy
[112] Astronomical Observatory of the Autonomous Region of the Aosta Valley (OAVdA), Loc. Lignan 39, I-11020, Nus (Aosta Valley), Italy
[113] Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS, Laboratoire Lagrange, Bd de l'Observatoire, CS 34229, 06304 Nice cedex 4, France
[114] Université Paris-Saclay, CNRS, Institut d'astrophysique spatiale, 91405, Orsay, France
[115] School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, CF24 3AA, UK
[116] Department of Physics, Oxford University, Keble Road, Oxford OX1 3RH, UK
[117] Aurora Technology for European Space Agency (ESA), Camino bajo del Castillo, s/n, Urbanizacion Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain
[118] ICL, Junia, Université Catholique de Lille, LITL, 59000 Lille, France
[119] Instituto de Física Teórica UAM-CSIC, Campus de Cantoblanco, 28049 Madrid, Spain
[120] CERCA/ISO, Department of Physics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, USA
[121] Technical University of Munich, TUM School of Natural Sciences, Physics Department, James-Franck-Str. 1, 85748 Garching, Germany
[122] Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany
[123] Laboratoire Univers et Théorie, Observatoire de Paris, Université PSL, Université Paris Cité, CNRS, 92190 Meudon, France
[124] Departamento de Física Fundamental. Universidad de Salamanca. Plaza de la Merced s/n. 37008 Salamanca, Spain
[125] Université de Strasbourg, CNRS, Observatoire astronomique de Strasbourg, UMR 7550, 67000 Strasbourg, France
[126] Center for Data-Driven Discovery, Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan
[127] Department of Physics & Astronomy, University of California Irvine, Irvine CA 92697, USA
[128] Departamento Física Aplicada, Universidad Politécnica de Cartagena, Campus Muralla del Mar, 30202 Cartagena, Murcia, Spain
[129] Instituto de Física de Cantabria, Edificio Juan Jordá, Avenida de los Castros, 39005 Santander, Spain
[130] INFN, Sezione di Lecce, Via per Arnesano, CP-193, 73100, Lecce, Italy
[131] Department of Mathematics and Physics E. De Giorgi, University of Salento, Via per Arnesano, CP-I93, 73100, Lecce, Italy
[132] INAF-Sezione di Lecce, c/o Dipartimento Matematica e Fisica, Via per Arnesano, 73100, Lecce, Italy
[133] CEA Saclay, DFR/IRFU, Service d'Astrophysique, Bat. 709, 91191 Gif-sur-Yvette, France
[134] Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK
[135] Department of Computer Science, Aalto University, PO Box 15400, Espoo, FI-00 076, Finland
[136] Instituto de Astrofísica de Canarias, c/ Via Lactea s/n, La Laguna 38200, Spain. Departamento de Astrofísica de la Universidad de La Laguna, Avda. Francisco Sanchez, La Laguna, 38200, Spain
[137] Caltech/IPAC, 1200 E. California Blvd., Pasadena, CA 91125, USA
[138] Universidad de La Laguna, Departamento de Astrofísica, 38206 La Laguna, Tenerife, Spain
[139] Ruhr University Bochum, Faculty of Physics and Astronomy, Astronomical Institute (AIRUB), German Centre for Cosmological Lensing (GCCL), 44780 Bochum, Germany
[140] Department of Physics and Astronomy, Vesilinnantie 5, 20014 University of Turku, Finland
[141] Serco for European Space Agency (ESA), Camino bajo del Castillo, s/n, Urbanizacion Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain
[142] ARC Centre of Excellence for Dark Matter Particle Physics, Melbourne, Australia
[143] Centre for Astrophysics & Supercomputing, Swinburne University of Technology, Hawthorn, Victoria 3122, Australia
[144] Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town, 7535, South Africa
[145] DAMTP, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK
[146] Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, UK
[147] Department of Astrophysics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland
[148] Department of Physics, Centre for Extragalactic Astronomy, Durham University, South Road, Durham, DH1 3LE, UK
[149] IRFU, CEA, Université Paris-Saclay 91191 Gif-sur-Yvette Cedex, France
[150] Univ. Grenoble Alpes, CNRS, Grenoble INP, LPSC-IN2P3, 53, Avenue des Martyrs, 38000, Grenoble, France
[151] INAF-Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, 50125, Firenze, Italy
[152] Dipartimento di Fisica, Sapienza Università di Roma, Piazzale Aldo Moro 2, 00185 Roma, Italy
[153] Centro de Astrofísica da Universidade do Porto, Rua das Estrelas, 4150-762 Porto, Portugal
[154] HE Space for European Space Agency (ESA), Camino bajo del Castillo, s/n, Urbanizacion Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain
[155] INAF - Osservatorio Astronomico d'Abruzzo, Via Maggini, 64100, Teramo, Italy
[156] Theoretical astrophysics, Department of Physics and Astronomy, Uppsala University, Box 516, 751 37 Uppsala, Sweden
[157] Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA
[158] Mathematical Institute, University of Leiden, Einsteinweg 55, 2333 CA Leiden, The Netherlands
[159] Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK
[160] Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRIStAL, 59000 Lille, France
[161] Department of Astrophysical Sciences, Peyton Hall, Princeton University, Princeton, NJ 08544, USA
[162] Space physics and astronomy research unit, University of Oulu, Pentti Kaiteran katu 1, FI-90014 Oulu, Finland
[163] Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, 10010, New York, NY, USA

Table A.1. Coefficients of the relations that map `Pinocchio` halo masses to Flagship AM masses (Eq. A.1), for the two code versions used in this paper.

	V4	V5
a_1	0.0243	0.0302
a_2	−0.0463	−0.0458
b_1	0.923	0.923
b_2	0.0220	0.0219
c_1	0.007 29	0.006 73
c_2	0.002 98	0.002 99

Appendix A: Calibration of `Pinocchio` and Flagship halo masses

This Appendix describes the adopted procedure to match `Pinocchio` and Flagship haloes in AM and CM (see Sect. 5). Binning haloes in redshift with bin size $\delta z = 0.01$, we measure the cumulative halo mass function of the Flagship simulation (over its octant footprint) and the stacked one from all the `Pinocchio` mocks of a set. AM of halo masses is then obtained by computing the Flagship and `Pinocchio` halo masses that give the same halo number densities. Because we have only one Flagship simulation, a direct match in number density would propagate its noise, so we fit the relation between the logarithms of the two masses with a second-order polynomial; this was found to be very accurate in the halo mass range of interest, namely from $10^{11}\,h^{-1}\,M_\odot$ to $10^{13}\,h^{-1}\,M_\odot$. We then linearly interpolate in redshift the three fitting coefficients of the power laws, thus creating an AM model for mapping `Pinocchio` halo masses to Flagship halo masses, $M_{\rm AM}$

$$M_{\rm AM,12} = (a_1+a_2 z)+(b_1+b_2 z)\,M_{\rm pin,12}+(c_1+c_2 z)\,M_{\rm pin,12}^2\,, \quad \text{(A.1)}$$

where $M_{\rm AM,12} := \log_{10}(M_{\rm AM}/10^{12}\,h^{-1}\,M_\odot)$ and $M_{\rm pin,12} := \log_{10}(M_{\rm pin}/10^{12}\,h^{-1}\,M_\odot)$. The coefficients are given in Table A.1 for V4 and V5 of the code.

A brute-force determination of CM would imply to measure power spectra for hundreds of mock catalogues for many halo mass thresholds and redshifts; we resort to a more convenient procedure by assuming that the relation between AM and CM masses is well represented by a power law with redshift-dependent coefficients. This leaves us with two parameters, that are calibrated using the galaxy catalogues. We require that, at $k < 0.2\,h\,{\rm Mpc}^{-1}$, the average clustering amplitude of 100 `Pinocchio` galaxy catalogues with $f_{{\rm H}\alpha} > f_0$, measured in the four redshift bins with edges in [0.9, 1.1, 1.3, 1.5, 1.8], is consistent with the Flagship one projected onto the same 30° radius footprint, to within the predicted variance. On smaller scales, that are hardly used by the standard inference methods but are necessary to constrain the one-halo term, we require that the clustering amplitude is always reproduced to better than 10%. After a trial-and-error procedure based on 100 catalogues, we find that the CM masses of the Flagship mock is well reproduced with the present scaling

$$\log_{10}(M_{\rm CM}/M_\odot) = \log_{10}(M_{\rm AM}/M_\odot) - 0.125 - 0.175\,(1-z)\,, \quad \text{(A.2)}$$

with no dependence on the code version. Figure A.1 shows the resulting relation between `Pinocchio` and Flagship halo masses valid for V5 of the code, in the case of AM and CM, at redshifts $z = 0.9$ and $z = 1.8$.

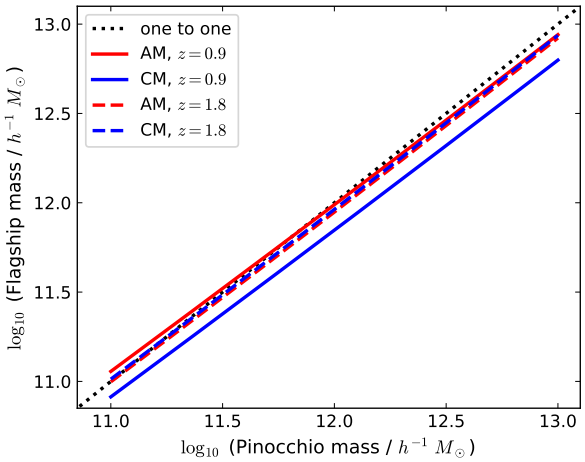

Fig. A.1. Relation between Flagship and `Pinocchio` masses (valid for V5 of the code) in the case of AM and CM, at the two redshifts $z = 0.9$ and $z = 1.8$.