

The effect of local Universe constraints on halo abundance and clustering

Maxwell L. Hutt¹,^{1,2} Harry Desmond¹,^{1,3,4} Julien Devriendt¹ and Adrienne Slyz¹

¹*Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*

²*The Blackett Laboratory, Imperial College London, Prince Consort Road, London SW7 2AZ, UK*

³*McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA*

⁴*Institute of Cosmology & Gravitation, University of Portsmouth, Dennis Sciana Building, Portsmouth PO1 3FX, UK*

Accepted 2022 August 22. Received 2022 August 8; in original form 2022 March 28

ABSTRACT

Cosmological N -body simulations of the dark matter component of the universe typically use initial conditions with a fixed power spectrum and random phases of the density field, leading to structure consistent with the local distribution of galaxies only in a statistical sense. It is, however, possible to infer the initial phases which lead to the configuration of galaxies and clusters that we see around us. We analyse the CSiBORG suite of 101 simulations, formed by constraining the density field within 155 $\text{Mpc } h^{-1}$ with dark matter particle mass $4.38 \times 10^9 M_\odot$, to quantify the degree to which constraints imposed on $2.65 \text{ Mpc } h^{-1}$ scales reduce variance in the halo mass function and halo–halo cross-correlation function on a range of scales. This is achieved by contrasting CSiBORG with a subset of the unconstrained Quijote simulations and expectations for the Λ CDM average. Using the FOF, PHEW, and HOP halofinders, we show that the CSiBORG suite beats cosmic variance at large mass scales ($\gtrsim 10^{14} M_\odot h^{-1}$), which are most strongly constrained by the initial conditions, and exhibits a significant halo–halo cross-correlation out to $\sim 30 \text{ Mpc } h^{-1}$. Moreover, the effect of the constraints percolates down to lower mass objects and to scales below those on which they are imposed. Finally, we develop an algorithm to ‘twin’ haloes between realizations and show that approximately 50 per cent of haloes with mass greater than $10^{15} M_\odot h^{-1}$ can be identified in all realizations of the CSiBORG suite. We make the CSiBORG halo catalogues publicly available for future applications requiring knowledge of the local halo field.

Key words: galaxies: haloes – galaxies: statistics – dark matter – large-scale structure of Universe – software: simulations.

1 INTRODUCTION

Undetectable except by its gravitational effect, dark matter (DM) is thought to be the dominant matter component of the Universe. Maps of large volumes of the Universe find galaxies to be arranged in a web-like structure of clusters, filaments and voids, tracing the distribution of the underlying dark matter density field. Understanding the evolution of this density field and its connection to visible tracers are currently key tasks in astrophysics.

Simulations of DM with initial conditions (ICs) based on a fixed power spectrum and random phases of the primordial matter field have been well studied for decades (Press & Schechter 1974; Davis et al. 1985; Lacey & Cole 1993; Eisenstein & Hut 1998; Tinker et al. 2008). Only more recently however have attempts been made to constrain the phases as well (i.e. the full 3D initial density field) so as to produce structures with positions matching those observed in the local Universe. This subfield originated with Hoffman & Ribak (1991), who developed a method for constructing constrained realizations of Gaussian random fields. Early applications of this algorithm developed galaxy count (Kolatt et al. 1996; Bistolas & Hoffman 1998) and peculiar velocity observations (van de Weygaert & Hoffman 2000; Klypin et al. 2003; Kravtsov, Klypin & Hoffman

2002) as means of deriving the constraints. Subsequent refinements increased the resolution of the simulations and the precision of the gravity model, improved the modelling used to generate the ICs (e.g. the galaxy bias), sampled the ICs rather than generating them in a maximum-likelihood sense, and implemented additional physics. Galaxy groups have also been used for deriving the constraints (Wang et al. 2014). A focus of constrained simulations has been the properties and formation history of the Local Group (Gottloeber, Hoffman & Yepes 2010; Sorce et al. 2015; Sawala et al. 2022; McAlpine et al. 2022), although they are finding use in an increasingly wide range of science applications.

It was realized already a decade ago (Jasche et al. 2010; Jasche & Wandelt 2012, 2013) that deriving the ICs of the local universe is best viewed as a Bayesian inference problem and solved by forward modelling. The method for doing this – which expanded to include full models for galaxy bias, redshift space distortions and observational selection effects – became known as the *Bayesian Origin Reconstruction from Galaxies* (BORG) algorithm, which has been applied to the 2M++ (Lavaux & Hudson 2011; Lavaux & Jasche 2016; Jasche & Lavaux 2019) and SDSS-III/BOSS (Lavaux, Jasche & Leclercq 2019) galaxy fields. This uses an efficient Hamiltonian Markov Chain Monte Carlo algorithm to sample the ICs at $z = 1000$ as well as the nuisance parameters. Although other constrained simulations have been run using ICs derived with similar techniques (e.g. Wang et al. 2014; Tweed et al. 2017; Wang et al. 2016), there has not before now been a statistical ensemble of such

* E-mail: m.hutt22@imperial.ac.uk (MLH); harry.desmond@port.ac.uk (HD)

simulations that fully samples the IC posteriors. This is needed to propagate the uncertainties in the IC inference into the low-redshift density field at high resolution, and hence quantify the extent to which our present knowledge of the local Universe’s ICs permits us to localize structures of various masses on various scales and reduce cosmic variance in their N -point statistics. Such an ensemble is also necessary to marginalize over the IC uncertainties in post-processing applications.

To fill this gap, we analyse CSiBORG (*Constrained Simulations in BORG*), a suite of constrained simulations with ICs derived from the 2M++ particle mesh BORG reconstruction (Jasche & Lavaux 2019). The 101 realizations of the suite span the converged part of the BORG posterior, allowing us to investigate the effect of the uncertainty it introduces. CSiBORG has been used previously to test galileon modified gravity in Bartlett, Desmond & Ferreira (2021) and to identify local voids-as-antihalos in Desmond et al. (2022).

Our focus here is on quantifying the effects of the constraints on the variance in halo abundance and clustering (i.e. their uncertainty), as a function of mass and scale, while accounting fully for the BORG IC uncertainties. This is important to understand the exact manner in which constraints on one scale affect structure on others, an important issue for all constrained simulations with ICs specified on a certain scale or set of scales. While the constraints will not affect the finite-volume uncertainty in halo statistics (for fixed box size), they will reduce shot noise on all scales because pinning down larger-scale modes will induce correlations on smaller scales. We assess this by comparing the CSiBORG statistics with those of Quijote (Villaescusa-Navarro et al. 2020), an unconstrained suite with otherwise similar properties. In particular, we address the following questions:

- (i) What is the impact of the BORG constraints on the variance in the halo mass function (HMF) between realizations of the local universe?
- (ii) How do the constraints propagate into the halo–halo cross-correlation function on a range of scales?
- (iii) Which haloes should be considered ‘the same’ between the different local-universe realizations, and how does the fraction of matched haloes depend on the properties of the haloes?
- (iv) How do these results depend on the halofinder used [Parallel HiEarchical Watershed (PHEW; Bleuler et al. 2015), friends-of-friends (FOF; Davis et al. 1985), or HOP (Eisenstein & Hut 1998)]?

This paper is structured as follows. In Section 2, we describe the CSiBORG and Quijote simulation suites. In Section 3, we present the halo statistics that we investigate, describe the halofinders used, and document our method for matching haloes between realizations. In Section 4, we compare the variance in the HMF and halo–halo cross-correlation function of CSiBORG to those of Quijote and analytic expectations, and show the results of our matching procedure. In Section 5, we discuss systematic uncertainties and summarize our results.

2 SIMULATED DATA: CSiBORG & Quijote

The CSiBORG suite comprises 101 realizations of a $677.7 \text{ Mpc } h^{-1}$ box centred on the MW with ICs inferred from the 2M++ survey (Lavaux & Jasche 2016), spanning the BORG posterior. Within this simulation volume, discretised on to a 256^3 grid, the BORG algorithm uses a particle-mesh gravity solver (Jasche & Lavaux 2019) to evolve the DM particle field from an initial redshift of $z = 1000$ to the present day. This constrains the ICs on a grid of cell length $2.65 \text{ Mpc } h^{-1}$.

The density field is well constrained in the volume in which the 2M++ survey has high completeness, a spherical volume with radius $\sim 155 \text{ Mpc } h^{-1}$ around the MW. In this subvolume the ICs, propagated linearly to $z = 69$, are augmented with white noise on a 2048^3 grid, giving a spatial resolution of $0.33 \text{ Mpc } h^{-1}$ and a minimum DM particle mass of $3.09 \times 10^9 \text{ M}_\odot h^{-1}$. Surrounding this zoom-in region, there is a buffer region in which the DM particle mass increases for distances further from the MW. This buffer extends approximately $10 \text{ Mpc } h^{-1}$ in order to smoothly connect the high resolution central region to the outer region with the same grid resolution as that of BORG. The adaptive mesh refinement code RAMSES (Teyssier 2002) is then used to evolve the DM density field to $z = 0$. Only the central augmented region is refined (reaching level 18 by $z = 0$, a spatial resolution of $2.6 \text{ kpc } h^{-1}$), so it is only there that we consider haloes. The cosmology of BORG (Jasche & Lavaux 2019), and hence the simulations, is primarily taken from the Planck Collaboration I (2014) best-fitting results including Wilkinson Microwave Anisotropy Probe (WMAP) polarization, high multipole moment, and baryonic acoustic oscillation (BAO) data, with the exception of H_0 which is taken from the 5-yr WMAP results combined with Type Ia supernovae and BAO data (Hinshaw et al. 2009): $T_{\text{CMB}} = 2.728 \text{ K}$, $\Omega_{\text{m}} = 0.307$, $\Omega_{\Lambda} = 0.693$, $\Omega_{\text{b}} = 0.04825$, $H_0 = 70.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\sigma_8 = 0.8288$, and $n = 0.9611$.

To quantify the effects of the constraints, we compare with unconstrained simulations. Ideally, one would run another suite of ~ 100 simulations identical to CSiBORG except with random ICs. As this is computationally very expensive (the full CSiBORG suite required $\sim 3\text{M CPU-hr}$), we instead investigate the publicly available Quijote simulations. Each of these is an N -body simulation evolved forward in time from unconstrained ICs at $z = 127$ to $z = 0$, using the GADGET-III code (Springel et al. 2008). The volume of each Quijote box is $1 (\text{Gpc } h^{-1})^3$. We use the 100 realizations with 1024^3 resolution, giving a particle mass of $8.72 \times 10^{10} \text{ M}_\odot h^{-1}$, and their ‘fiducial cosmology’ ($\Omega_{\text{m}} = 0.3175$, $\Omega_{\Lambda} = 0.6825$, $\Omega_{\text{b}} = 0.049$, $H_0 = 67.11 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\sigma_8 = 0.834$, and $n = 0.9624$), with the phases of the ICs varying randomly between realizations at fixed power spectrum. This subset is chosen as its resolution is high enough to compare to CSiBORG over a broad mass range. In future, the development of suites of matched constrained and unconstrained simulations would eliminate any potential biases due to differences between the simulations other than their ICs.

To isolate the effect of the IC constraints, we must mitigate the impact of the different cosmological parameters used by CSiBORG and Quijote. We do this by running two simulations with unconstrained ICs drawn from a Gaussian density field at $z = 69$ with the same resolution, box size, and refined subvolume as the CSiBORG simulations. One of these has the cosmology of CSiBORG and the other that of Quijote. Again, we consider only haloes located in the central $155 \text{ Mpc } h^{-1}$ ‘zoom-in’ sphere. These boxes link the two simulation suites, since they have identical resolution and ICs and differ only in their cosmological parameters.

3 METHOD

In this section, we describe the methods used for halofinding, the statistics that we compare between the simulations described above and the procedure for matching haloes between realizations.

3.1 Halofinding

We run three halofinders on the CSiBORG boxes and the two additional unconstrained simulations: PHEW, FOF, and HOP. For

Quijote, we use only the FOF algorithm that has already been run on it by Villaescusa-Navarro et al.

The watershed halofinder PHEW was run on-the-fly with RAMSES for the CSiBORG simulations. This algorithm splits the density field into basins by locating watersheds (boundaries between regions which share a common local minimum density via steepest descent). These basins are sequentially merged according to user-defined thresholds on the density of saddle points between basins. We use the standard threshold value from Bleuler et al. of $200\rho_c$ to merge subhaloes into haloes.

For FOF, we adopt a linking length parameter of $b = 0.2$. This matches the choice made for the Quijote simulations by Villaescusa-Navarro et al. FOF links together particles within a given ‘linking distance’ of each other, defined as b times the mean interparticle separation. It is known that FOF tends to connect extraneous particles to haloes along nearby filaments of the density field (Eisenstein & Hut 1998), creating more massive structures than other algorithms. This is particularly true for haloes which are merging near $z = 0$ as FOF is more likely to cluster them into a single object than other algorithms, an effect similar to the ‘bridged’ haloes discussed by Lukić et al. (2009). Warren et al. (2006) also find that FOF haloes are biased towards larger masses for haloes with low particle number and propose a correction to the particle number of the form $N_{\text{corrected}} = N_0(1 - N_0^{-0.6})$ for this effect, with N_0 the original number of DM particles in the halo. Marginally decreasing the particle number in this way was found to eliminate the effect of resolution on halo mass. Lukić et al. have shown that corrections of this sort must depend on halo concentration as well as particle number. For the purpose of quantifying the variance between realizations, however, this level of precision of the HMF is not necessary and we do not use such corrections.

Finally, we run HOP. This algorithm first generates an estimate of the local density at each particle position using a kernel over a number N_{dens} of neighbouring particles. Then it associates with each particle the highest density particle from its nearest $N_{\text{hop}} - 1$ neighbours. This association is repeated until a local maximum of the density field is reached. All particles above a certain density threshold ρ_t which lead to the same local maximum are then grouped together. We run HOP with parameters $N_{\text{hop}} = N_{\text{dens}} = 20$ and $\rho_t = 80$, which produces similar results to FOF with $b = 0.2$ (see Tweed et al. 2009 for a summary of these parameters in the AdaptaHOP (Aubert, Pichon & Colombi 2004) algorithm, which reduces to HOP when substructure is neglected).

In order to reduce resolution errors in halo properties we require that haloes found by any halofinder must contain at least 100 particles. This corresponds to a minimum halo mass of $3.09 \times 10^{11} M_\odot h^{-1}$ for CSiBORG and $8.72 \times 10^{12} M_\odot h^{-1}$ for Quijote. However, the lower resolution of the Quijote simulations means that this constraint would severely limit the mass range across which haloes are found compared to CSiBORG. Hence, we include haloes from the Quijote suite with a minimum of 20 DM particles when it is desirable to extend the comparison with CSiBORG to lower masses.

3.2 The HMF and halo–halo cross-correlation function

We compare halo distributions via two main summary statistics: the HMF, defined by

$$n(M) \equiv \frac{dn}{d \log M}, \quad (1)$$

where \log is the base-10 logarithm throughout, and the two-point halo–halo cross-correlation function between pairs of simulation

boxes. Denoting the halo density field by $\rho(\mathbf{x})$ and its mean by $\bar{\rho}$, we measure the cross-correlation

$$\xi_{ij}^{(hh)}(\mathbf{r}) = \frac{1}{V} \int d^3\mathbf{x} \delta_i(\mathbf{x}) \delta_j(\mathbf{x} - \mathbf{r}), \quad (2)$$

where the subscripts label the pair of simulation boxes being compared and

$$\delta_i(\mathbf{x}) \equiv \frac{\rho_i(\mathbf{x})}{\bar{\rho}_i} - 1, \quad (3)$$

is the halo overdensity field. Assuming an isotropic universe, $\xi(\mathbf{r}) = \xi(r)$ and we can estimate the cross-correlation by counting pairs of haloes separated by distances which fall into a given set of radial distance bins and adopt the Landy–Szalay estimator (Landy & Szalay 1993)

$$\hat{\xi}_{ij}^{(hh)}(r) = \frac{D_i D_j - D_i R - D_j R + RR}{RR}, \quad (4)$$

where D_i represents the catalogue of data points (haloes) in simulation i and R is a catalogue of random points. That is, $D_i D_j(r)$ is calculated as the number of pairs of data points with separation, r , falling in a given distance bin, and similarly for the other terms. We use 100 times as many random points as data points so that the uncertainty in the estimator itself (which is at the Poisson level) is negligible compared to the variance between realizations. We increase the statistical significance by averaging the cross-correlation function over all pairs of realizations in a given suite (CSiBORG or Quijote), and denote the averaged quantity simply as $\langle \xi(r) \rangle$.

The HMF is calculated in each box using 20 bins spaced evenly in $\log(M/(M_\odot h^{-1}))$. Using a larger number of bins has no impact on the results at the low-mass end but decreases the visibility of trends at the high-mass end where shot noise variance is large as there are fewer haloes. For each suite we then find the mean and standard deviation of the HMF, as well as the correlation coefficients between the halo counts in each mass bin (normalised covariance), over all the realizations.

The variance of statistics calculated from N -body simulations, often called ‘cosmic variance’, has several origins (see Colombi et al. 2000; Colombi, Bouchet & Schaeffer 1994; Moster et al. 2011; and references therein). The most important are (a) shot noise due to the finite number of objects in the simulation, (b) finite-volume uncertainty due to underlying density modes on scales larger than the simulation volume, and (c) edge effect uncertainty, which is related to the geometry of the survey in that objects near its edge carry less weight than those far from the edge (Szapudi & Colombi 1996). The latter two sources of variance are negligible compared to shot noise at scales approaching homogeneity ($\sim 500 \text{ Mpc } h^{-1}$; Schneider et al. 2016). The ICs of CSiBORG will not influence finite-volume or edge-effect uncertainties over the entire simulation box, but the density field constraints will reduce shot noise. By reducing variance in the density modes at and above the IC scale ($\sim 2.65 \text{ Mpc } h^{-1}$), the constraints do however induce correlations on smaller scales that effectively reduce finite-volume variance for them too. We will quantify the extent to which constraints determine structure both within and beyond the range of scales on which they are directly imposed.

Since the full Quijote volume of $1 \text{ Gpc } h^{-1}$ is well beyond the homogeneity scale, those boxes are not subject to such large finite-volume uncertainties and are therefore dominated by shot noise. However, given that the refined $155 \text{ Mpc } h^{-1}$ region of the CSiBORG simulations is far below it, we cannot use the same method to isolate the shot noise contribution to the CSiBORG HMF variance. We proceed by using the full Quijote volume to verify that the shot

noise uncertainty is Poissonian, allowing us to infer the shot noise contribution within the refined region.

The logic is as follows. First, we measure the standard deviation σ_{155} of the HMF of both the CSiBORG and Quijote simulations within $V_{155} = 4\pi (155 \text{ Mpc}/h)^3/3$, such that finite-volume and edge-effect uncertainties are held constant. Secondly, we measure the Quijote HMF in the full $V_{1000} = (1 \text{ Gpc } h^{-1})^3$ volume for all the Quijote simulations and calculate its standard deviation σ_{1000} . This allows us to confirm that the HMF measured in the smaller volume is indeed consistent with that found over the whole box, and to estimate the isolated shot noise uncertainty. We can then verify that the shot noise uncertainty is Poissonian by computing the expected Poisson variance in V_{155} , namely $(N_{\text{halo}})^{-1/2}$, and comparing it with the inferred variance from the full Quijote volume, $\sigma_{1000}\sqrt{V_{155}/V_{1000}}$. N_{halo} is the average number of haloes in the $155 \text{ Mpc } h^{-1}$ region over all realizations. If these two measurements of the shot noise are consistent then we can be confident that the shot noise error in V_{155} is given by $(N_{\text{halo}})^{-1/2}$ and we can assume this for CSiBORG also. Thus, we isolate the shot noise contribution to both the Quijote and CSiBORG HMF variances and can identify whether the constraints on CSiBORG decrease the variance below the Poissonian expectation.

We compare the HMF to a standard function from the literature, Warren et al. (2006, hereafter W06). We choose this particular analytic form because it was originally fitted to the results of a FOF halofinder on a wide range of unconstrained N -body simulations. Other works (Despali et al. 2015) have shown that improved parametrizations and halo definitions lead to HMFs which are ~ 30 per cent discrepant from W06 at large halo masses. We do not, therefore, expect this curve to be fully compatible with our results, but it does allow us to compare the CSiBORG and Quijote HMFs to a standard often seen in the literature.

3.3 Matching haloes between realizations

We now describe a method for identifying haloes that are sufficiently similar in mass and location between all realizations of a simulation suite to be called the ‘same’ object. The haloes that are then present in *all* realizations are particularly robust and should potentially be prioritized in post-processing applications. For this, we use the FOF halo catalogues of CSiBORG and consider haloes with masses greater than $10^{13} \text{ M}_{\odot} h^{-1}$ to enable comparison with Quijote over the same mass range. The procedure is as follows:

- (i) Select a reference simulation and record the positions of all haloes in it.
- (ii) Select a halo in the reference simulation and calculate its size according to $R_{200} = (3M/(4\pi \cdot 200\rho_{\text{crit}}))^{1/3}$.
- (iii) In all other simulations, locate all haloes within a ‘search radius’ around the position of the reference halo selected in the previous step and select the one most similar in mass. If there are any simulations in which *no* haloes are located within the search radius around the reference halo position, discard that halo as it is not consistently present in all simulations. This results in a list of the haloes that are ‘matched’ to the reference halo in all the other simulations.
- (iv) Calculate the variance in the mass and position (i.e. distance from the reference halo position) of the matched haloes across all simulation boxes. This gives a measure of the success of the matching for this particular halo in the reference simulation.
- (v) Repeat from step (ii) until all the haloes in the reference simulation have been processed.

- (vi) Repeat from step (i) until all simulation boxes have been used as the reference simulation.

There are several possible ways to select the ‘search radius’ in step (iii). The simplest – a fixed distance in Mpc – would bias the fraction of matched haloes high at low mass relative to high mass because more massive, larger haloes would be expected to fluctuate more in position than lower mass ones. The next simplest is to scale the search radius with the halo radius, i.e. NR_{200} with constant N ; this is our fiducial scheme with $N = 5$. Although this avoids the previous bias, it may miss low-mass haloes that could be matched with a more realistically lenient scheme. This is because the deviation in position of haloes between the realizations is related to the haloes’ peculiar velocities. The velocity power spectrum peaks at larger radius than the matter power spectrum so that haloes peculiar velocities tend to be larger than their virial velocities (Sheth & Diaferio 2001; Hahn, Angulo & Abel 2015). Thus, one might expect lower mass haloes to deviate by more than $5R_{200}$ between the realizations, which would prefer a different function for the search radius. Our choice is conservative in the sense that it only associates low-mass haloes that are closer together than the expectation from the velocity field, and a custom matching procedure to suit the needs of the user may readily be applied to our public halo catalogues.

In our fiducial model, the exact ratio of the virial and search radii, N , is somewhat arbitrary but does not qualitatively affect the salient results: the variance in position and mass of the matched haloes, the difference in the number of matched haloes in CSiBORG and Quijote, and the dependence of this number on mass. We implement the procedure for $N = 3, 5, 10$ to verify this. Note that the search radius around a reference halo with mass $10^{13} \text{ M}_{\odot} h^{-1}$ is $5R_{200} \approx 1.8 \text{ Mpc } h^{-1}$ whereas for a very massive $10^{15} \text{ M}_{\odot} h^{-1}$ halo, it is approximately $8 \text{ Mpc } h^{-1}$. Hence, for the smallest haloes being considered, we expect far greater variance in the matching since the search radius is below the IC constrained scale. Nevertheless, identifying more successful matches at low masses in CSiBORG than Quijote is a strong indicator of the induced correlation on scales below those set by the CSiBORG constraints.

There is, however, a subtle residual bias in the matching procedure towards smaller haloes since within a given search radius in any realization, the probability to find a small halo is far greater than that of a large halo simply because there are more small haloes than large ones. Consider the case of a very massive reference halo. When searching for a match in another simulation, it is possible that there is no halo of similar mass within the search radius. In this case, the procedure is more likely to match this reference halo to a small halo which is found there. Hence, the mean mass of the matched haloes tends to be slightly lower than the mass of the reference halo. In our application, there are negligibly many ($\mathcal{O}(0.1 \text{ per cent})$) haloes where the mean mass of matched haloes differs by more than a factor of 5 from the reference halo mass and hence we have not imposed any filters on the matched haloes. In general, one may reduce this bias by imposing a threshold variance on the mass of the matched haloes, or on the ratio of the mean matched halo mass to the reference halo mass.

We show the result of this procedure by plotting the variance in the mass and position of all the matched haloes as a function of the reference halo mass, with the expectation that larger haloes will be more reliably matched between simulations because they are better constrained by the CSiBORG ICs. We also run this procedure on the Quijote suite in order to determine the proportion of haloes one would expect to find suitable matches for in the absence of IC constraints.




Figure 1. *Upper panel:* HMF of the pair of simulations with identical unconstrained ICs but differing in their cosmological parameters calculated with different halo finders. The solid lines use the CSiBORG cosmology and the dashed lines the Quijote cosmology. (Note that none of these boxes are part of the actual CSiBORG or Quijote suite; the labels ‘Quijote’ and ‘CSiBORG’ refer merely to the cosmology.) The black curves show the HMF from W06. *Lower panel:* Relative residual from W06 using the CSiBORG cosmology.

4 RESULTS

4.1 CSiBORG & Quijote HMFs

We begin our investigation of the HMF by quantifying the shift in it due to the different cosmologies of CSiBORG and Quijote. Fig. 1 shows the HMFs of the two unconstrained simulations performed using the CSiBORG resolution with identical ICs but different cosmologies, as described in Section 2, using the three halo finders. We show also a pair of curves using the HMF from W06, evaluated for both cosmologies, as a point of comparison.

It is apparent from this plot that, for a fixed cosmology, PHEW exhibits different behaviour to FOF and HOP: its HMF is > 10 per cent greater for masses between 10^{12} and $10^{14} M_\odot h^{-1}$ and markedly lower at very large masses. PHEW also shows fundamentally different behaviour to the other halo finders between the two cosmologies, as shown in Fig. 2. The PHEW HMF is larger in the CSiBORG cosmology than in the Quijote one, whereas the curves from W06, HOP, and FOF have the opposite behaviour. Since this work focuses on the variance between different realizations, we do not investigate further this discrepancy between PHEW and the other halo finders. In light of this result, we do, however, regard the FOF and HOP results as more reliable. In Fig. 1, we note also a decline in the HMF from all three halo finders at $M \lesssim 10^{12.5} M_\odot h^{-1}$ compared to the fiducial curves. This is due to a variety of numerical effects affecting haloes containing fewer than ~ 300 particles, including the starting redshift of the ICs, the order of the Lagrangian Perturbation Theory used to generate these (Hahn, Rampf & Uhlemann 2021; Michaux et al. 2021) and the gravity solver used to run the simulation (e.g. Schneider et al. 2016).

We now show the HMF of the CSiBORG and Quijote suites, averaged over all simulation boxes, alongside the result from W06 evaluated using the CSiBORG cosmology (Fig. 3). The Quijote curve is calculated in a $155 \text{ Mpc } h^{-1}$ spherical volume to allow comparison to the CSiBORG ‘zoom-in’ region, although we also show the Quijote HMF of the full $(1 \text{ Gpc } h^{-1})^3$ volume with a




Figure 2. Ratio of HMFs of the Quijote and CSiBORG cosmologies shown in Fig. 1 for the different halo finders. The black curve is the ratio found using the W06 HMF. In the axis label, the subscripts ‘C’ and ‘Q’ refer to the CSiBORG and Quijote cosmologies, respectively.




Figure 3. *Upper panel:* Mean HMFs of the CSiBORG suite using all three halo finders and of the Quijote suite using the FOF halo finder both in the full $(1 \text{ Gpc } h^{-1})^3$ volume (grey dashed) and the spherical $155 \text{ Mpc } h^{-1}$ volume (pink). The shaded regions around each curve show $\pm 1\sigma$ standard deviation estimated across all the simulations in a suite. The black curve shows the HMF from W06, calculated in the CSiBORG cosmology (solid line) and Quijote cosmology (dashed line). *Lower panel:* Relative residual from W06 using the CSiBORG cosmology. The light grey vertical dotted line indicates the mass at which haloes in the Quijote suite contain 100 particles.

dashed grey line and standard deviation band between realisations. The Quijote lines include all haloes with ≥ 20 particles to enable comparison to CSiBORG down to lower mass, although we caution that haloes with fewer than 100 particles ($M < 8.72 \times 10^{12} M_\odot h^{-1}$

for Quijote) have mass uncertainty at the ~ 30 – 50 per cent level due to imperfect resolution (Trenti et al. 2010).

The Quijote HMF begins to exceed the W06 curve corresponding to the same cosmology considerably for $M \gtrsim 10^{14.5} M_\odot h^{-1}$. This behaviour is well documented (Eisenstein & Hut 1998) and is a result of the greatly varying geometries of FOF haloes: the algorithm tends to link particles along filaments of the cosmic web, increasing the halo mass in a way which is not observed for algorithms that produce haloes with more consistent geometrical structure. We find that the CSiBORG HMFs overshoot W06 even further in this mass range, however, and also fall considerably below it at lower masses. At low masses, all halo finders produce a HMF ~ 20 per cent below W06 and at the largest masses probed ($\sim 10^{15.7} M_\odot h^{-1}$) the CSiBORG HMF is ~ 20 times (1.5σ) larger than that of Quijote. This discrepancy cannot be accounted for by a 5 per cent reduction in the value of Ω_m used to calculate the W06 HMF (which corresponds to the relative underdensity of the CSiBORG volume compared to the cosmic average). The BORG reconstruction of the $2M++$ volume therefore implies more massive structures, and fewer low-mass scale structures, than average. While this is not highly statistically significant at high mass due to cosmic variance (nor in the density field itself; McAlpine et al. 2022), it is at lower masses. These results may indicate an issue with the $2M++$ BORG ICs (J. Jasche & G. Lavaux, private communication), which does not, however, appreciably affect the variance between realizations, the focus of our work here.

We also find differences between the results from different halo finders. The three halo finders produce consistent results at large masses, whereas at the low-mass end there is a 7σ difference between the HMFs of FOF and HOP. Despite being close to the FOF result at the smallest halo masses, in the intermediate-mass range the PHEW HMF has a characteristic local maximum in the Warren residual reminiscent of the trend found in the unconstrained simulations in Fig. 1. We also note that at large masses, the HOP HMF rises consistently above those of FOF and PHEW (which agree very closely above $10^{14} M_\odot h^{-1}$). This behaviour is not seen to such an extent in the unconstrained simulations shown in Fig. 1 and the halo finder comparison project of Knebe et al. (2011) found the opposite behaviour, with the FOF HMF approximately 30 per cent greater than that of HOP at the bright end. This could indicate that the relative behaviour of the halo finders is a function of the ICs.

4.2 Variance of the HMF

We show in Fig. 4 the fractional standard deviation of the HMF across all simulation boxes of the Quijote and CSiBORG suites (for all halo finders) using the same $155 \text{ Mpc } h^{-1}$ spherical volume for all simulations. We again include all Quijote FOF haloes down to 20 particles to increase the range of comparison. The fractional deviation of the Quijote HMF is approximately a factor of 2 larger than that of CSiBORG across the entire mass range. At the lowest masses, below $1.30 \times 10^{13} M_\odot$, the fractional deviation of the Quijote suite is up to $2.6 \times$ larger than that of CSiBORG due to the poorer identification of haloes with < 100 particles.

We compare to the shot noise variance of the Quijote HMF in the $155 \text{ Mpc } h^{-1}$ volume, calculated in two separate ways. First, we calculate the expected fractional standard deviation due to Poissonian shot noise, $\langle N_{\text{halo}} \rangle^{-1/2}$ (using the Quijote HMF to estimate $\langle N_{\text{halo}} \rangle$) with a dashed black line. Secondly, as discussed in Section 3.2, we estimate the shot-noise contribution in the $155 \text{ Mpc } h^{-1}$ volume by normalizing the standard deviation of the HMF calculated using the entire Quijote volume across all simulations. This is shown




Figure 4. *Upper panel:* Fractional standard deviation of the HMF of the CSiBORG and Quijote suites across all simulations in the central $155 \text{ Mpc } h^{-1}$ spherical volume. The dashed pink line is the shot noise contribution to the standard deviation in the $155 \text{ Mpc } h^{-1}$ sphere estimated by normalizing the standard deviation measured over the entire Quijote volume (see Section 3.2). The expected Poisson shot noise calculated in the $155 \text{ Mpc } h^{-1}$ sphere for the Quijote HMF is shown with a dashed black line, and for the CSiBORG FOF HMF with a solid black line. *Lower panel:* Relative residuals of the fractional standard deviation with respect to the Poissonian shot noise expectation for the corresponding HMF. The light grey vertical dotted line indicates the mass at which haloes in the Quijote suite contain 100 particles.

with a dashed pink line. These two methods align closely across all halo masses, which gives confidence that the excess deviation in the Quijote HMF calculated in the $155 \text{ Mpc } h^{-1}$ volume above the shot-noise contribution is, indeed, due to finite-volume effects, and also that the shot-noise error is Poissonian. This also confirms that the dominant source of error at large scales is shot noise since the Poissonian shot noise uncertainty (dashed black) and the measured standard deviation in the $155 \text{ Mpc } h^{-1}$ volume (solid pink) are negligibly different above $10^{15} M_\odot h^{-1}$. We note in particular that since the Quijote simulations are unconstrained, their fractional deviation is never lower than shot noise uncertainty.

We also show the expected Poissonian shot noise uncertainty in the CSiBORG HMF, $\langle N_{\text{halo}} \rangle^{-1/2}$, by using the FOF HMF to estimate $\langle N_{\text{halo}} \rangle$, with a solid black line. The lower panel of the figure shows the residual of the measured deviation of the HMF (from a given halo finder) with respect to the corresponding Poissonian shot noise uncertainty. For example, the blue line shows the residual of the measured variance of the HOP HMF with respect to $\langle N_{\text{halo}} \rangle^{-1/2}$, where $\langle N_{\text{halo}} \rangle$ is estimated using the HOP HMF, and similarly for the other curves. We find a stark difference between the ‘excess’ variance above Poisson shot noise at low masses, which is due to residual finite-volume variance in the relatively small volume used. Since all HMFs are measured in the same region, the finite-volume contribution to the uncertainty is similar between the Quijote and CSiBORG simulations (they are not exactly equal as the largest




Figure 5. Heatmap of the correlation coefficient matrix of the number of haloes in each mass bin for the FOF halofinder in the CSiBORG (left) and Quijote (right) suites. The final row and column on the right are left blank, as there are no haloes in the final mass bin in the Quijote suite.




Figure 6. Mean cross-correlation function between all pairs of simulations in the CSiBORG and Quijote suites. Only haloes with $M > 10^{13} M_{\odot} h^{-1}$ are included. The vertical dashed line shows the IC constrained scale ($2.65 \text{ Mpc } h^{-1}$). The error bars show $\pm 1\sigma$ deviation measured in the correlation functions between all pairs of realizations.

wavelengths of the density field modes differ between the Quijote and CSiBORG suite) and therefore the \sim two-fold reduction in variance across all masses is a result of the IC constraints in CSiBORG alone.

The true constraining power of the CSiBORG ICs is best evidenced at large masses, $M \gtrsim 10^{14} M_{\odot} h^{-1}$, where the HMFs measured by all three halofinders have variance below that of the Poissonian expectation. At these halo masses, uncertainty is dominated by shot noise and the CSiBORG variance beats Poisson by as much as 50 per cent.

4.3 Covariance of halo masses

Fig. 5 shows the full correlation coefficient matrix of halo counts in CSiBORG and Quijote, again going down to 20 particles in the latter. This is the covariance matrix normalized by the standard deviation and hence measures correlations rather than the magnitude

of the scatter:

$$\frac{\text{cov}(N_i, N_j)}{\sigma_{N_i} \sigma_{N_j}} = \frac{\langle (N_i - \langle N_i \rangle)(N_j - \langle N_j \rangle) \rangle}{\sigma_{N_i} \sigma_{N_j}}, \quad (5)$$

where N_i is the number of haloes in the i th mass bin and the averages are taken over all realizations. We use the FOF halofinder for both suites for consistency, although the results for HOP and PHEW are extremely similar for CSiBORG.

The dominant effect for the unconstrained Quijote suite is that the mass in the central $155 \text{ Mpc } h^{-1}$ region in which we identify haloes varies by ~ 8 per cent between realizations. Realizations with an upward fluctuation of this mass contain more haloes at all masses, introducing a global positive covariance. This is particularly significant for less massive haloes which respond more uniformly to the fluctuations in the total mass of the volume as they are much more numerous. There are $\mathcal{O}(1)$ haloes in the most massive bins in each realization and haloes with masses in excess of $10^{15} M_{\odot} h^{-1}$ make up less than 10 per cent of the total mass in haloes. Hence, even in a realization with an 8 per cent increase in the total mass of the $155 \text{ Mpc } h^{-1}$ volume, the number of haloes in each of these most massive bins is unlikely to vary systematically. This reduces the covariance between the most massive haloes in the Quijote suite, which, due to statistical fluctuations, is sometimes found to be negative.

Including the IC constraints in the CSiBORG suite, the dominant effect is instead that a halo lost from one mass bin in one realization is likely to be found in a nearby one in any other realization. The total mass in haloes varies by only ~ 2 per cent between realizations in CSiBORG, making this effect subdominant to that of haloes ‘matching’ between realizations with only a small mass difference. This predominantly affects nearby mass bins and massive haloes for which the ICs are most constraining; for widely separated bins, or at lower mass, the correlation may again become positive due to the effect of correlated mass fluctuations (although at a weaker level than in Quijote due to the lower amplitude of these fluctuations). These results complement those on the variance of the HMF, and illustrate the extent to which haloes may be taken to match among

the constrained realizations, discussed in more detail in Section 4.5.

4.4 Cross-correlation function

Fig. 6 shows the two-point halo–halo cross-correlation function between haloes of mass in excess of $10^{13} M_{\odot} h^{-1}$ in the central $155 \text{ Mpc } h^{-1}$ spherical volume, averaged over all pairs of realizations in the CSiBORG and Quijote suites. The solid line shows the average over all pairs of realizations, and the shaded band the standard deviation among them. The different simulation boxes of the Quijote suite are fully uncorrelated, as expected. The CSiBORG suite, however, demonstrates strong positive cross-correlation, increasing on smaller scales. This result is salient as the trend continues down to scales below the IC constrained scale ($\sim 2.65 \text{ Mpc } h^{-1}$). This implies that constraints on large ($> \text{Mpc } h^{-1}$) scales induce correlation on smaller ones, despite having considerably larger variance at these (formally) unconstrained small scales. This is because smaller-scale modes reside in larger ones, inheriting their constraints to some extent.

4.5 Matching haloes between realizations

We show in Fig. 7 the fractional standard deviation of the distance between the matched and reference haloes, and of the mass of matched haloes, when we carry out the halo matching procedure described in Section 3.3 on CSiBORG using the FOF halofinder. We find a clear decrease in the relative variance for massive haloes compared to lower mass ones, particularly in the matched halo mass. Averaged over all haloes of all simulations, for a search radius of $5R_{200}$, approximately 1.4 per cent of haloes are matched throughout all 101 realizations. However, the fraction of haloes that are successfully matched varies with the mass of the reference halo, as seen in Fig. 8. For example, for the intermediate search radius of $5R_{200}$, over 50 per cent of all haloes with masses between 10^{15} and $10^{16} M_{\odot} h^{-1}$ are matched throughout, whereas only a negligible fraction of those with masses $< 10^{14} M_{\odot} h^{-1}$ are present in all realizations. This clearly demonstrates the effectiveness of the constraints for pinning down massive haloes.

To make sure our matching procedure depends on the strength of the IC constraints alone, we apply it also to the Quijote suite. For $N = 5$, we find no matched haloes across all simulations at any mass, as we would expect from unconstrained realizations. Hence, the result that 1.4 per cent of all haloes found in a given simulation (and approximately half of the largest haloes) are matched successfully throughout the *entire* CSiBORG suite is a statistically significant consequence of the IC constraints. Naturally, unconstrained modes beneath the scale of the constraints lead to structure formation differing between realizations. Our halo matching procedure, as well as the cross-correlation function discussed above, is a strong diagnostic of the strength of the IC inference from BORG, and provides a natural way to quantify the improvement in the strength of the constraints that will be afforded by future inferences. In the public halo catalogues, we include a flag that indicates when a halo is matched to a halo in all other CSiBORG realizations when the realization in question is taken as the reference, using our fiducial matching scheme ($N = 5$).

5 DISCUSSION AND CONCLUSION

We analyse the CSiBORG suite of 101 constrained realizations of the $2M++$ volume using ICs drawn from the BORG posterior, identifying

haloes within $155 \text{ Mpc } h^{-1}$ of the MW using three halofinders (PHEW, FOF, and HOP). We compare the resulting HMFs and halo–halo cross-correlation functions with those of a similar subset of the unconstrained Quijote simulations to quantify the extent to which the local universe constraints reduce variance between the realizations. This indicates the reduction in the uncertainty of halo statistics afforded by the BORG algorithm.

By using the entire Quijote simulation volume to eliminate finite-volume effects, we isolate the shot-noise contribution to the variance in the HMF, and show that the IC constraints of CSiBORG lead to a fractional standard deviation which is ~ 2 times lower than that of the unconstrained Quijote suite over all masses, and up to 50 per cent lower than Poisson shot-noise uncertainty at large masses $\gtrsim 10^{15} M_{\odot} h^{-1}$. The similarity of haloes between CSiBORG realizations is reflected in the covariance matrix of halo masses, which is near-uniformly positive in Quijote but typically negative for nearby mass bins in CSiBORG. We also investigate the effect of the constraints on the halo–halo cross-correlation across all pairs of simulation boxes and find positive correlation on all scales, increasing towards short distance. The fact that this trend continues below the IC scale of $2.65 \text{ Mpc } h^{-1}$ implies that large-scale constraints increase consistency in structures on much smaller, formally unconstrained scales. That is, variations in the smaller scale modes of the density field are limited by the constraints on the larger modes that they inhabit. Hence, IC constraints imposed on a given set of scales percolate down to introduce effective constraints, albeit with larger variance, on smaller scales.

We develop a halo matching procedure to identify structures which are particularly well localized throughout the entire CSiBORG simulation suite. The variance in the masses and positions of the matched haloes with respect to a given reference halo decreases strongly with increasing mass, indicating that the CSiBORG constraints have greater effect for more massive haloes. We demonstrate that the qualitative results of this matching are independent of the details of its implementation, and that our method finds no matches within the unconstrained Quijote suite. We publicly release the CSiBORG halo catalogues from the PHEW, FOF, and HOP algorithms for analyses requiring knowledge of the local halo field, including its uncertainty given state-of-the-art constraints.

In order to quantify the impact of constrained ICs on the halo population we rely on halofinder algorithms, and our results indicate significant differences between them. In future, the application of 6D algorithms such as Rockstar (Behroozi, Wechsler & Wu 2012) or Velociraptor (Elahi et al. 2019) to suites like CSiBORG and Quijote would improve halofinding at the massive end because they are less susceptible to erroneously combining haloes that are merging near $z = 0$.

We anticipate that constrained DM-only and hydrodynamical simulations will be of great value to future studies of the local universe. Future galaxy surveys with high completeness out to scales comparable with the homogeneity scale across large parts of the sky will reduce the finite-volume uncertainties that dominate in small simulation volumes. In parallel, the continued development of algorithms like BORG will integrate the ever-improving observational efforts with the numerical techniques required to pin down the distribution of DM on cosmological scales. Incorporating additional information from, for example, peculiar velocities (Boruah, Lavaux & Hudson 2021) or cosmic shear (Porqueres et al. 2021) would strengthen the constraints and further reduce variance in the statistics we consider here. An investigation of the effects of ICs imposed on multiple scales would also be illuminating (cf. Sawala et al. 2022 for the local group).




Figure 7. Matched haloes distribution for a search radius of $5R_{200}$. *Left-hand panel:* Fractional standard deviation in the position of matched haloes with respect to the reference halo position, $\Delta R = |r_{\text{ref}} - r|$, where r is the matched halo position, as a function of the mass of the reference halo. *Right-hand panel:* Fractional standard deviation in the mass of matched haloes as a function of the mass of the reference halo. The dashed black lines trace the mean fractional standard deviation of the matched haloes in each mass bin.




Figure 8. Cumulative fraction of CSiBORG haloes matched between all realizations for three different search radii NR_{200} ($N = 3, 5, 10$). The shaded region shows $\pm 1\sigma$ standard deviation measured across all reference realizations. The percentages at the far right by each line indicate the total fraction of haloes that are matched throughout the entire suite.

ACKNOWLEDGEMENTS

We thank Jens Jasche, Guilhem Lavaux, and Stephen Stopyra for helpful discussions. We also acknowledge the authors of the `Halomod` (Murray et al. 2021) and `Corrfunc` (Sinha & Garrison 2020) PYTHON libraries, whose codes were used to generate the fiducial HMF curves and the cross-correlation functions, respectively. HD was supported by a Junior Research Fellowship at St John’s College, Oxford, a McWilliams Fellowship at Carnegie Mellon University and a Royal Society University Research Fellowship (grant no. 211046). MLH was funded by the President’s PhD Scholarships. This work was done within the Aquila Consortium.

This work used the DiRAC Complexity and DiRAC@Durham facilities, operated by the University of Leicester IT Services and Institute for Computational Cosmology, which form part of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment is funded

by BIS National E-Infrastructure capital grants ST/K000373/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, STFC DiRAC Operations grant ST/K0003259/1, and Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National E-Infrastructure.

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 693024).

DATA AVAILABILITY

Our CSiBORG halo catalogues for the PHEW, FOF, and HOP halo finders are publicly available [here](#). The Quijote halo catalogues are available [here](#) (Villaescusa-Navarro et al. 2020). Other data will be made available upon reasonable request.

REFERENCES

- Aubert D., Pichon C., Colombi S., 2004, *MNRAS*, 352, 376
- Bartlett D. J., Desmond H., Ferreira P. G., 2021, *Phys. Rev. D*, 103, 023523
- Behroozi P. S., Wechsler R. H., Wu H.-Y., 2012, *ApJ*, 762, 109
- Bistolas V., Hoffman Y., 1998, *ApJ*, 492, 439
- Bleuler A., Teyssier R., Carassou S., Martizzi D., 2015, *Comput. Astrophys. Cosmol.*, 2, 5
- Boruah S. S., Lavaux G., Hudson M. J., 2021, preprint ([arXiv:2111.15535](https://arxiv.org/abs/2111.15535))
- Colombi S., Bouchet F. R., Schaeffer R., 1994, *A&A*, 281, 301
- Colombi S., Szapudi I., Jenkins A., Colberg J., 2000, *MNRAS*, 313, 711
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- Desmond H., Hutt M. L., Devriendt J., Slyz A., 2022, *MNRAS*, 511, L45
- Despali G., Giocoli C., Angulo R. E., Tormen G., Sheth R. K., Baso G., Moscardini L., 2015, *MNRAS*, 456, 2486
- Eisenstein D. J., Hut P., 1998, *ApJ*, 498, 137
- Elahi P. J., Cañas R., Poulton R. J. J., Tobar R. J., Willis J. S., Lagos C. d. P., Power C., Robotham A. S. G., 2019, *PASA*, 36, e021
- Gottloeber S., Hoffman Y., Yepes G., 2010, preprint ([arXiv:1005.2687](https://arxiv.org/abs/1005.2687))
- Hahn O., Angulo R. E., Abel T., 2015, *MNRAS*, 454, 3920
- Hahn O., Rampf C., Uhlemann C., 2021, *MNRAS*, 503, 426
- Hinshaw G. et al., 2009, *ApJS*, 180, 225
- Hoffman Y., Ribak E., 1991, *ApJ*, 380, L5

- Jasche J., Lavaux G., 2019, *A&A*, 625, A64
- Jasche J., Wandelt B. D., 2012, *MNRAS*, 425, 1042
- Jasche J., Wandelt B. D., 2013, *MNRAS*, 432, 894
- Jasche J., Kitaura F. S., Wandelt B. D., Enßlin T. A., 2010, *MNRAS*, 406, 60
- Klypin A., Hoffman Y., Kravtsov A. V., Gottlöber S., 2003, *ApJ*, 596, 19
- Knebe A. et al., 2011, *MNRAS*, 415, 2293–2318
- Kolatt T., Dekel A., Ganon G., Willick J. A., 1996, *ApJ*, 458, 419
- Kravtsov A. V., Klypin A., Hoffman Y., 2002, *ApJ*, 571, 563
- Lacey C., Cole S., 1993, *MNRAS*, 262, 627
- Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
- Lavaux G., Hudson M. J., 2011, *MNRAS*, 416, 2840
- Lavaux G., Jasche J., 2016, *MNRAS*, 455, 3169
- Lavaux G., Jasche J., Leclercq F., 2019, preprint ([arXiv:1909.06396](https://arxiv.org/abs/1909.06396))
- Lukić Z., Reed D., Habib S., Heitmann K., 2009, *ApJ*, 692, 217
- McAlpine S. et al., 2022, *MNRAS*, 512, 5823
- Michaux M., Hahn O., Rampf C., Angulo R. E., 2021, *MNRAS*, 500, 663
- Moster B. P., Somerville R. S., Newman J. A., Rix H.-W., 2011, *ApJ*, 731, 113
- Murray S. G., Diemer B., Chen Z., Neuhold A., Schnapp M., Peruzzi T., Blevins D., Engelman T., 2021, *Astronomy and Computing*, 36, 100487
- Planck Collaboration I, 2014, *A&A*, 571, A1
- Porqueres N., Heavens A., Mortlock D., Lavaux G., 2021, *MNRAS*, 502, 3035
- Press W. H., Schechter P., 1974, *ApJ*, 187, 425
- Sawala T., McAlpine S., Jasche J., Lavaux G., Jenkins A., Johansson P. H., Frenk C. S., 2022, *MNRAS*, 509, 1432
- Schneider A. et al., 2016, *J. Cosmol. Astropart. Phys.*, 2016, 047
- Sheth R. K., Diaferio A., 2001, *MNRAS*, 322, 901
- Sinha M., Garrison L. H., 2020, *MNRAS*, 491, 3022
- Sorce J. G. et al., 2015, *MNRAS*, 455, 2078
- Springel V. et al., 2008, *MNRAS*, 391, 1685
- Szapudi I., Colombi S., 1996, *ApJ*, 470, 131
- Teyssier R., 2002, *A&A*, 385, 337
- Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, *ApJ*, 688, 709
- Trenti M., Smith B. D., Hallman E. J., Skillman S. W., Shull J. M., 2010, *ApJ*, 711, 1198
- Tweed D., Devriendt J., Blaizot J., Colombi S., Slyz A., 2009, *A&A*, 506, 647
- Tweed D., Yang X., Wang H., Cui W., Zhang Y., Li S., Jing Y. P., Mo H. J., 2017, *ApJ*, 841, 55
- van de Weygaert R., Hoffman Y., 2000, in Courteau S., Willick J., eds, *ASP Conf. Ser. Vol. 201, Cosmic Flows Workshop*. Astron. Soc. Pac., San Francisco, p. 169
- Villaescusa-Navarro F. et al., 2020, *ApJS*, 250, 2
- Wang H., Mo H. J., Yang X., Jing Y. P., Lin W. P., 2014, *ApJ*, 794, 94
- Wang H. et al., 2016, *ApJ*, 831, 164
- Warren M. S., Abazajian K., Holz D. E., Teodoro L., 2006, *ApJ*, 646, 881

This paper has been typeset from a \LaTeX file prepared by the author.