

The halo bispectrum as a sensitive probe of massive neutrinos and baryon physics

Victoria Yankelevich¹  ¹★, Ian G. McCarthy¹  ¹★, Juliana Kwan,¹ Sam G. Stafford¹  ¹ and Jia Liu²

¹*Astrophysics Research Institute, Liverpool John Moores University, Liverpool L3 5RF, UK*

²*Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan*

Accepted 2023 February 16. Received 2023 February 14; in original form 2022 February 15

ABSTRACT

The power spectrum has been a workhorse for cosmological studies of large-scale structure. However, the present-day matter distribution is highly non-Gaussian and significant cosmological information is also contained in higher order correlation functions. Meanwhile, baryon physics (particularly active galactic nucleus feedback) has previously been shown to strongly affect the two-point statistics but there has been limited exploration of its effects on higher order functions to date. Here, we use the BAHAMAS suite of cosmological hydrodynamical simulations to explore the effects of baryon physics and massive neutrinos on the halo bispectrum. In contrast to matter clustering which is suppressed by baryon physics, we find that the halo clustering is typically enhanced. The strength of the effect and the scale over which it extends depends on how haloes are selected. On small scales ($k \gtrsim 1 h \text{ Mpc}^{-1}$, dominated by satellites of groups/clusters), we find that the bispectrum is highly sensitive to the efficiency of star formation and feedback, making it an excellent testing ground for galaxy formation models. We show that the effects of feedback and the effects of massive neutrinos are largely separable (independent of each other) and that massive neutrinos strongly suppress the halo bispectrum on virtually all scales up to the free-streaming length (apart from the smallest scales, where baryon physics dominates). The strong sensitivity of the bispectrum to neutrinos on the largest scales and galaxy formation physics on the smallest scales bodes well for upcoming precision measurements from the next generation of wide-field surveys.

Key words: methods: numerical – galaxies: haloes – large-scale structure of Universe.

1 INTRODUCTION

Providing a quantitative explanation for the formation and evolution of large-scale structure (LSS) in the Universe is one of the key tests of modern cosmological theories. The standard model of cosmology, the so-called Lambda cold dark matter (Λ CDM) model, has been remarkably successful in reproducing a wide range of observations, including the observed properties of the cosmic microwave background (e.g. Planck Collaboration VI 2020) and low-redshift probes, such as baryon acoustic oscillations (e.g. Eisenstein et al. 2005) and redshift–space distortions (see e.g. Alam et al. 2017). While Λ CDM has been shown to reproduce most current LSS observations well, the physical nature of both dark matter and dark energy remain elusive. Furthermore, there have been a number of recent mild tensions reported in the best-fitting parameter values for certain cosmological parameters, including the Hubble constant (see Verde, Treu & Riess 2019 for a review) and the LSS parameter $S_8 \equiv \sigma_8 \sqrt{\Omega_m}/0.3$ (where σ_8 is defined as the amplitude of the (linear) power spectrum on the scale of $8 h^{-1} \text{ Mpc}$ and Ω_m is the present-day matter density parameter), as derived from measurements of cosmic shear, galaxy clustering, and the abundance of massive galaxy clusters (see e.g.

discussion in McCarthy et al. 2018). Whether these tensions are hinting at the presence of (unaccounted for) systematic errors in some of the cosmological analyses, that there is new physics beyond the standard model, or that they represent moderately large statistical fluctuations is presently unclear.

The drive to test Λ CDM further, and to look for new physics beyond the standard model, requires that we make increasingly precise measurements of ‘tried and tested’ LSS statistics (e.g. two-point clustering of galaxy clustering, weak lensing, etc.), but also that we devise new tests of LSS which are capable of probing different aspects of the matter distribution, which in turn could hopefully break important degeneracies between the fitted cosmological parameters. The halo bispectrum (or three-point function of haloes/galaxies), which we focus on in this study, is one such example of an LSS test that is becoming increasingly observationally feasible (e.g. Peebles & Groth 1975; Gaztanaga 1994; Scoccimarro et al. 2001; Verde et al. 2002; Gaztañaga et al. 2005; Pan & Szapudi 2005; Kulkarni et al. 2007; McBride et al. 2011; Gil-Marín et al. 2015a,b, 2017; Slepian et al. 2017a,b; Gualdi et al. 2019a,b; Pearson & Samushia 2018; Veropalumbo et al. 2021) and for which previous theoretical studies have shown contains a significant source of cosmological information beyond what may be obtained from standard two-point clustering tests alone (e.g. Song, Taruya & Oka 2015; Chudaykin & Ivanov 2019; Yankelevich & Porciani 2019; Barreira 2020; Gualdi & Verde 2020; Hahn et al. 2020;

* E-mail: yankelevich.victoria@gmail.com (VY); I.G.McCarthy@ljmu.ac.uk (IM)

Heinrich & Doré 2020; Agarwal et al. 2021; Eggemeier et al. 2021; Hahn & Villaescusa-Navarro 2021; Moradinezhad Dizgah et al. 2021).

As in the case of most two-point statistics, much of the observational signal of the bispectrum is expected to come from quasi-linear and non-linear scales. Thus, ‘beyond linear’ methods for evaluating the growth of structure, such as perturbation theory and full cosmological simulations, are generally required to predict the bispectrum. Examples of previous bispectrum work based on perturbation theory and cosmological simulations include Bernardeau et al. (2002), Crocce & Scoccimarro (2006), Bernardeau, Crocce & Scoccimarro (2008), Pietroni (2008), Bernardeau, Crocce & Scoccimarro (2012), Taruya et al. (2012), Angulo et al. (2015), Baldauf, Mercolli & Zaldarriaga (2015), Lazanu et al. (2016), Steele & Baldauf (2021), and Alkhanishvili et al. (2022).

In addition to the extra complexity of following the gravitational evolution of matter when density fluctuations become large, there is also the challenge of accurately accounting for non-gravitational physics that comes into play. Specifically, feedback processes associated with the formation of stars and the growth of black holes are well-established, both empirically and via theoretical models and full cosmological hydrodynamical (hereafter hydro) simulations, to have a strong effect on the overall matter distribution on scales of up to a few tens of Mpc (e.g. van Daalen et al. 2011; Schneider & Teyssier 2015; Mummery et al. 2017; Chisari et al. 2019; van Daalen, McCarthy & Schaye 2020). Previous studies that focused on two-point statistics have shown that, unless these effects are properly modelled, they will result in biases in the derived cosmological parameters (e.g. Semboloni et al. 2011).

Given that galaxy formation physics has been shown to significantly affect classical two-point statistics, it is reasonable to expect that it should also affect the three-point correlation function/bispectrum. But how large are the effects? Is the bispectrum more or less sensitive to these processes than previously studied LSS quantities? And how do we mitigate against these effects (if the goal is cosmological constraints) or exploit them (if the goal is to better understand galaxy formation)? These are the questions we focus on in this study.

Note that, thus far, there has been remarkably little attention devoted to the effects of baryon physics on the halo bispectrum using cosmological simulations. This is plausibly explained by the fact that very large cosmological volumes (box sizes of at least several hundred Mpc on a side) are required to reliably measure the bispectrum. Combined with the constraint that the simulations must also use full hydrodynamics and a model for galaxy formation in order to self-consistently incorporate the impact of baryons, there are very few simulations that presently meet these criteria. In fact, we are aware of only two suites of hydro simulations of sufficient volume for this purpose: the BAHAMAS simulations (McCarthy et al. 2017, 2018) and the Magneticum simulations (Bocquet et al. 2016; Dolag, Komatsu & Sunyaev 2016). However, to date neither have been used to examine the halo/galaxy bispectrum. Semboloni, Hoekstra & Schaye (2013) presented a study of baryonic effects on the *matter* bispectrum using the OWLS simulations (Schaye et al. 2010). Their study also showed how the combination of the power spectrum and bispectrum can be used to break degeneracies between baryonic and cosmological effects. Foreman et al. (2020) have used BAHAMAS to also explore the effects of baryons on the matter bispectrum, finding that baryon physics does significantly affect the matter bispectrum and that the effects differ in magnitude and scale dependence compared to their effects on the matter power spectrum. Aricò et al. (2021) used the measurements of Foreman et al. (2020) to show that their

‘baryonification’ method for altering the outputs of collisionless (DM-only, DMO) simulations can be used to simultaneously model the effects of baryons on the matter power spectrum and the reduced matter bispectrum. We note here that, while the matter bispectrum is related to the halo bispectrum (and that the former is considerably easier to measure in simulations, given a large number of tracer particles), they are clearly distinct quantities, as haloes/galaxies are biased tracers of the matter distribution. Indeed, we will show later that the impact of baryons on halo clustering is qualitatively different (opposite sign) to that on the matter distribution. Note that, observationally, the halo/galaxy bispectrum may be measured from galaxy redshift surveys, whereas the matter bispectrum can be inferred from a combination of galaxy and weak lensing (cosmic shear) surveys (e.g. Fu et al. 2014).

In this paper, we use the BAHAMAS suite of cosmological hydro simulations to explore the impact of baryon physics (particularly active galactic nucleus, AGN, feedback) on the halo bispectrum. We also examine the role of massive neutrinos, noting that previous theoretical studies have highlighted the sensitivity of the bispectrum to this component (e.g. Hahn et al. 2020; Chen, Upadhye & Wong 2021; Hahn & Villaescusa-Navarro 2021), and the potential degeneracy between baryons and neutrinos.

The paper is structured as follows. In Section 2, we describe the cosmological simulations used in this study. In Section 3, we describe our methods for measuring the halo power spectrum and bispectrum and, for the latter, discuss the dependence on triangular configuration. In Section 4, we present our main results, showing the dependence of the bispectrum on baryon physics, halo mass, and neutrino mass, and we explore the degeneracy between baryons and neutrinos. In Section 5, we summarize and discuss our findings.

2 BAHAMAS

In this study, we use the BAHAMAS¹ cosmological simulations (McCarthy et al. 2017, 2018) to explore the impact of baryons and neutrinos on the halo bispectrum. BAHAMAS is a suite of large-volume hydro simulations that simultaneously explores variations in baryonic physics and cosmology, including extensions to Λ CDM such as massive neutrinos (Mummery et al. 2017), dynamical dark energy (Pfeifer et al. 2020), and a running scalar spectral index (Stafford et al. 2020). Here, we use seven runs from McCarthy et al. (2018) which explore variations in AGN feedback and the summed mass of the neutrinos in the context of a WMAP nine-yr cosmology (Hinshaw et al. 2013). Table 1 presents the cosmological parameters of BAHAMAS simulations used in this work.

Each of the hydro runs used in this study has a box size of $400 h^{-1}$ Mpc and includes 2×1024^3 particles, equally split between baryons and dark matter, which have a mass ratio equal to $\Omega_b/\Omega_{\text{cdm}}$. (As described below, the neutrinos are not followed using particles.) In addition, for each hydro run, there is a corresponding collisionless simulation which, as described in van Daalen et al. (2020), is represented by two collisionless fluids (one for the baryons and one for the dark matter).

The Boltzmann code CAMB² (Lewis, Challinor & Lasenby 2000; Howlett et al. 2012) was used to compute the transfer functions and a modified version of N-GENIC was used to create the initial conditions at a starting redshift of $z = 127$. N-GENIC was modified to include second-order Lagrangian perturbation theory corrections and support

¹<https://www.astro.ljmu.ac.uk/~igm/BAHAMAS/>

²<http://camb.info/>

Table 1. Cosmological parameter values for the BAHAMAS simulations used here. For all the models a WMAP nine-based cosmology is adopted. The columns are: (1) model name; (2) summed mass of the three active neutrino species (we adopt a normal hierarchy for the individual masses); (3) subgrid AGN heating temperature; (4) Hubble’s constant; (5) present-day baryon density; (6) present-day dark matter density; (7) present-day neutrino density, computed as $\Omega_\nu = M_\nu/(93.14 \text{ eV } h^2)$; (8) spectral index of the initial power spectrum; (9) amplitude of the initial matter power spectrum at a CAMB pivot k of $2 \times 10^{-3} \text{ Mpc}^{-1}$; (10) present-day (linearly evolved) amplitude of the matter power spectrum on a scale of $8 h^{-1} \text{ Mpc}$ (note that we use A_s rather than σ_8 to compute the power spectrum used for the initial conditions, thus the ICs are ‘CMB normalized’). In addition to the cosmological parameters, we also list the following simulation parameters: (11) dark matter particle mass; (12) initial baryon particle mass.

(1) Model	(2) M_ν (eV)	(3) ΔT_{heat} (K)	(4) H_0 ($\text{km s}^{-1} \text{ Mpc}^{-1}$)	(5) Ω_b	(6) Ω_{cdm}	(7) Ω_ν	(8) n_s	(9) A_s (10^{-9})	(10) σ_8	(11) M_{DM} ($10^9 h^{-1} \text{ M}_\odot$)	(12) $M_{\text{bar, init}}$ ($10^8 h^{-1} \text{ M}_\odot$)
Fiducial	0.0	$10^{7.8}$	70.00	0.0463	0.2330	0.0	0.9720	2.392	0.8211	3.85	7.66
Low AGN	0.0	$10^{7.6}$	70.00	0.0463	0.2330	0.0	0.9720	2.392	0.8211	3.85	7.66
High AGN	0.0	$10^{8.0}$	70.00	0.0463	0.2330	0.0	0.9720	2.392	0.8211	3.85	7.66
0.06	0.06	$10^{7.8}$	70.00	0.0463	0.2317	0.0013	0.9720	2.392	0.8069	3.83	7.66
0.12	0.12	$10^{7.8}$	70.00	0.0463	0.2304	0.0026	0.9720	2.392	0.7924	3.81	7.66
0.24	0.24	$10^{7.8}$	70.00	0.0463	0.2277	0.0053	0.9720	2.392	0.7600	3.77	7.66
0.48	0.48	$10^{7.8}$	70.00	0.0463	0.2225	0.0105	0.9720	2.392	0.7001	3.68	7.66

for massive neutrinos.³ Separate transfer functions, computed by CAMB, are used for each individual component (baryons, neutrinos, and CDM) when producing the initial conditions. Note also that the same random phases are used for each of the simulations, such that comparisons between different runs are not subject to cosmic variance.

The simulations were carried out with a version of the Lagrangian TreePM-SPH code GADGET3 (last described in Springel 2005), which was modified to include subgrid physics as part of the OWLS project (Schaye et al. 2010). The gravitational softening is fixed to $4 h^{-1} \text{ kpc}$ (in physical coordinates below $z = 3$ and in comoving coordinates at higher redshifts) and the SPH smoothing is done using the nearest 48 neighbours.

To include the effects of massive neutrinos, BAHAMAS uses the semilinear algorithm developed by Ali-Haïmoud & Bird (2013). The algorithm evaluates neutrino perturbations on the fly at every time-step using a linear perturbation integrator, which is sourced from the full non-linear baryons + CDM potential and is added to the total gravitational force. The dynamical responses of the neutrinos to the baryons + CDM and of the baryons + CDM to the neutrinos are mutually and self-consistently included.

The hydro simulations include subgrid prescriptions for metal-dependent radiative cooling, star formation, and stellar evolution (including chemical enrichment from Type II and Ia supernovae and asymptotic giant branch stars). The simulations include stellar feedback and prescriptions for supermassive black hole growth and AGN feedback. See Schaye et al. (2010) and references therein for a detailed description of the subgrid implementations.

As described in McCarthy et al. (2017), the parameters controlling the efficiencies of the stellar and AGN feedback were adjusted so that the simulations reproduce the present-day galaxy stellar mass function (GSMF) for $M_* > 10^{10} \text{ M}_\odot$ and the amplitude of the gas mass fraction–halo mass relation of groups and clusters, as inferred from high-resolution X-ray observations. These quantities were selected to ensure that the collapsed structures in the simulations have the correct baryon content in a global sense, which van Daalen et al. (2020) have shown is critical for capturing the impact of baryons on the matter power spectrum. We point out that the parameters governing the feedback efficiencies were *not* recalibrated when varying the cosmologi-

cal parameters away from the fiducial WMAP nine-yr cosmology (with massless neutrinos). However, as shown in McCarthy et al. (2018), the internal properties of collapsed structures (stellar masses, gas masses, etc.) are, to first order, insensitive to the variations in cosmology. Nevertheless, we explore the potential degeneracy between baryon physics and neutrino free-streaming in Section 4.2.

3 BISPECTRUM ESTIMATION

3.1 Power spectrum and bispectrum definitions

The power spectrum and the bispectrum are the Fourier transforms of two- and three-point correlation functions, respectively, of the dimensionless overdensity $\delta(\mathbf{x})$ parameter, which is defined as

$$\delta(\mathbf{x}) = \frac{\rho(\mathbf{x})}{\bar{\rho}} - 1, \quad (1)$$

where $\rho(\mathbf{x})$ is a continuous random field which gives the local density per unit comoving volume in the expanding Universe and $\bar{\rho}$ is the mean density with $\bar{\rho} = \langle \rho(\mathbf{x}) \rangle$ (the brackets $\langle \dots \rangle$ denotes the averages taken over an ideal ensemble of realizations).

We define the power spectrum and the bispectrum, respectively, as

$$\langle \delta(\mathbf{k}) \delta(\mathbf{k}') \rangle = (2\pi)^3 P(\mathbf{k}) \delta_{\text{D}}(\mathbf{k} + \mathbf{k}'), \quad (2)$$

$$\langle \delta(\mathbf{k}_1) \delta(\mathbf{k}_2) \delta(\mathbf{k}_3) \rangle = (2\pi)^3 B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \delta_{\text{D}}(\mathbf{k}_{123}), \quad (3)$$

where $\delta_{\text{D}}(\mathbf{k})$ is the three-dimensional Dirac delta function with $\mathbf{k}_{123} = \mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3$, which implies that the bispectrum is defined only for closed triangles of wavevectors.

The halo power spectrum and the halo bispectrum measured from numerical simulation or observational data can be significantly affected by shot noise. Assuming the distribution of haloes derives from Poisson sampling, the measured spectra (which is a sum of the true spectra and shot noise terms) can be defined as (e.g. Oddo et al. 2020, and references therein)

$$\tilde{P}(\mathbf{k}) = P(\mathbf{k}) + \frac{1}{n_{\text{h}}}, \quad (4)$$

$$\tilde{B}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) + [P(\mathbf{k}_1) + P(\mathbf{k}_2) + P(\mathbf{k}_3)] \frac{1}{n_{\text{h}}} + \frac{1}{n_{\text{h}}^2}, \quad (5)$$

³<https://github.com/sbird/S-GenIC>

where n_h is the halo number density (e.g. Matarrese, Verde & Heavens 1997). Note that the assumption of a Poisson distribution for the bispectrum shot noise term is accurate to first-order only, since the clustering of haloes is not strictly Poissonian. We discuss in detail the role of Poisson and non-Poisson shot noise in Section 3.4.

3.2 Halo power spectrum and bispectrum measurements

In order to measure halo power spectra and bispectra, we modify the publicly available BSKIT⁴ code of Foreman et al. (2020). The original code measures the matter power spectrum and matter bispectrum from the BAHAMAS, ILLUSTRISTNG⁵, ILLUSTRIS,⁶ and EAGLE⁷ simulations. BSKIT is built upon the NBODYKIT (Hand et al. 2018) simulation analysis package, and uses the standard FFT-based bispectrum estimator (Scoccimarro 2000; Sefusatti et al. 2016; Tomlinson, Jeong & Kim 2019). By default, the BAHAMAS module of BSKIT reads in the positions and masses of particles in order to compute the matter power spectrum and bispectrum. We modify it for this study to read in the positions (defined here as the centre of potential) and masses of haloes as opposed to particles. When computing the halo power spectra and bispectra with BSKIT by default we weight each halo equally. In Appendix A, we explore the effects of weighting haloes by their mass and show that such a weighting scheme tends to reduce the impact of baryons on the bispectrum, as higher mass haloes, which are less affected by ejective AGN feedback, contribute more to the computed spectra.

In this work, we analyse all self-gravitating substructures, often referred to as ‘subhaloes’, regardless of whether they are central or satellite subhaloes. Such a selection is closer to a galaxy-based selection than selecting central subhaloes only since all subhaloes of sufficient mass are expected to host galaxies. Without risk of confusion, going forward we will collectively refer to such self-gravitating substructures as just ‘haloes’. For BAHAMAS, a standard friends-of-friends (FoF) algorithm (Davis & Peebles 1983) was run to identify FoF groups and the SUBFIND (Springel et al. 2001; Dolag et al. 2009) algorithm was then used to identify all self-gravitating substructures (haloes).

Note that, because the simulations use the same random phases in the initial conditions, it is possible to identify the same haloes in both the hydro run and its corresponding collisionless run. This allows us to isolate the impact of baryonic and cosmological effects in a straightforward way. We refer to the case where we analyse a common set of haloes between multiple simulations as a ‘matched’ sample. An ‘unmatched’ sample refers to the case where we simply select haloes based on some criteria (e.g. total mass) from each simulation without regard for whether they correspond to the same haloes.⁸ Observationally, it is often the case that haloes/galaxies are selected by, for example, their total masses, thus it is also useful to analyse the results for the unmatched case.

To construct matched samples of haloes, we use the unique dark matter particle IDs, using the 50 most bound particles in each halo of a reference simulation (e.g. the collisionless run). Whichever halo in the target simulation (e.g. the hydro simulation) contains the largest fraction of these dark matter particles from the reference simulation

is considered that haloes match. Matches are performed bijectively, however, and only haloes that are matched both backwards and forwards are retained in the final matched sample.

The size of the simulation box and its mass resolution limits the range of scales over which the power spectra and bispectra can be measured. $k_f = 2\pi/L_{\text{box}} = 0.0157h \text{ Mpc}^{-1}$ corresponds to the so-called the fundamental k -mode of the BAHAMAS simulations, and is the longest wavelength mode that can be measured given the finite box size. In terms of the binning strategy between these limits, the choice of bin width, Δk , represents a trade-off between the accuracy of results and the computational cost. Note that the computational cost comes in two forms, which is the evaluation of the spectra from a grid (with finer bins requiring more evaluations) and the construction of the grid itself, which can be the most costly aspect depending on the required resolution. It has been shown before that the bispectrum is sensitive to the size of Δk in various ways (e.g. Yankelevich & Porciani 2019; Oddo et al. 2020). In many theoretical works, a binning of $\Delta k = k_f$ is used. However, such narrow bins may not be the best choice for measuring the bispectrum from simulations, as the computational cost scales inversely with the width of the k bins used. In addition to this, narrow k -bins also leads to additional numerical noise. The common solution to this problem is to have narrow (wide) k -bin for small (large) k . For example, Foreman et al. (2020) use $\Delta k = k_f$ for $k < 40k_f$ and $\Delta k = 6k_f$ for $40k_f < k < k_{\text{max}}$, where k_{max} is the maximum wavevector for which the analysis is done. We adopt $k_{\text{max}} = 3h \text{ Mpc}^{-1}$ throughout.

Fortunately, the computational resources to measure the halo bispectrum are lower than for the matter bispectrum in the example above, since the construction of the grid is considerably cheaper for the former. Therefore, when computing the halo bispectrum we use the same bin size of $\Delta k = 2k_f$ over the entire range of k -scales. However, note that when plotting power spectrum and bispectrum ratios in the figures described later, we apply some smoothing to the curves to reduce the impact of noise. Specifically, we rebin the spectra by a factor 3 and apply a Savitzky–Golay filter.

3.3 Choice of triangular configuration

A potentially important question is what triangular configuration k_1, k_2, k_3 should be adopted when computing the bispectrum. Results in the literature are often presented for the equilateral case only. However, there is no particularly strong motivation behind this choice, except perhaps for convenience. We present here a brief analysis of the dependence of (the baryons effects on) the halo bispectrum on the choice of triangular configuration.

In Fig. 1, we show the dependence of the effects of baryonic physics on the halo bispectrum as a function of triangular configuration (note that we discuss the physical origin of these effects in Section 4). We do this by plotting the ratio of the bispectrum measured from the fiducial hydro simulation with respect to the corresponding collisionless (DMO) simulation for the case with massless neutrinos.

For this test, we calculate the bispectrum for all haloes with masses above $10^{11.5} h^{-1} \text{ M}_\odot$. This choice, which roughly corresponds to a stellar mass of 10^{10} M_\odot , is dictated primarily by the mass resolution of the simulations. For the unmatched case, this mass cut is applied to both the DMO and hydro simulations separately, and thus somewhat different populations of haloes are selected in the two runs. For the matched case, however, we apply the mass cut to the DMO simulation and then select the corresponding matched haloes from the hydro simulation. For the purposes of this comparison we fix k_1 to $1.54 h \text{ Mpc}^{-1}$, since (i) there is an equilateral triangle with such side

⁴<https://github.com/sjforeman/bskit>

⁵<https://www.tng-project.org>

⁶<http://www.illustris-project.org>

⁷<http://icc.dur.ac.uk/Eagle/>

⁸In general they will not be the same, since feedback and cosmological effects alter the masses of haloes.

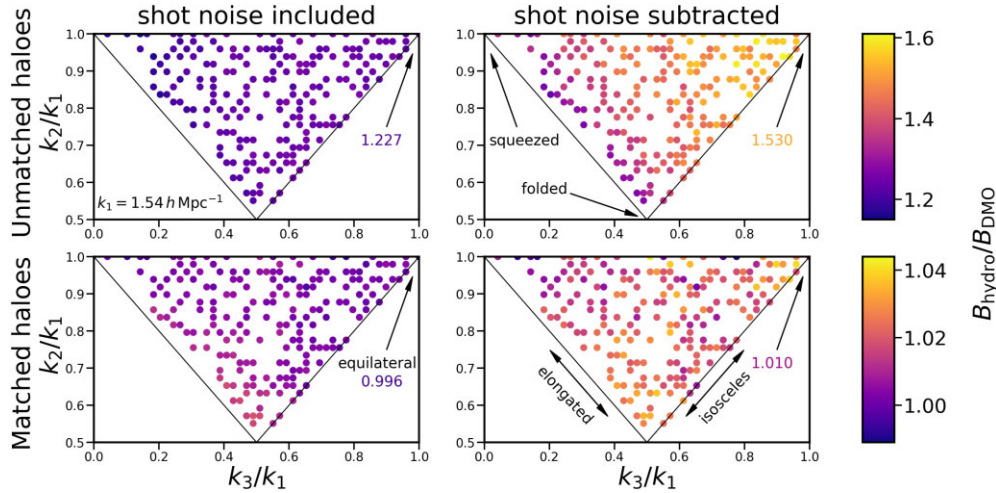


Figure 1. The colour-coded value of the ratio of the bispectrum, $B_{\text{hydro}}/B_{\text{DMO}}$, measured from the fiducial hydro and dark matter-only BAHAMAS simulations for different triangular configuration for the fixed value of $k_1 = 1.54 \, h \, \text{Mpc}^{-1}$. The ratio is presented for unmatched (top row) and matched (bottom row) haloes with masses $M \geq 10^{11.5} \, h^{-1} \, \text{M}_{\odot}$, with and without shot noise subtraction (right-hand and left-hand columns, respectively). The colour-coded numbers inside each subplot show the value of the ratio for the equilateral triangular configuration. The arrows indicate the positions of various triangular configurations.

length; (ii) there is a sufficient number of independent triangles such that the bispectrum is well measured; and (iii) it falls approximately half way between the full range of scales that we consider in the rest of the paper ($0.0157 \, h \, \text{Mpc}^{-1}$ to $3 \, h \, \text{Mpc}^{-1}$). In Appendix B, we explore how the choice of triangular configuration impacts the results for different choices of k_1 .

The dots in Fig. 1 (and in Figs B1 and B2) appear in an irregular way because the bispectrum measurements from the simulations do not contain all the possible triangular configurations for the chosen values of k_1 . We adopt the same order of wavevector length as in BSKIT which is $k_1 \geq k_2 \geq k_3$. For theoretical calculations all the triangles which follow triangles' condition on the sides, in our case, $k_1 \leq k_2 + k_3$ would exist and will fill all the $k_2/k_1 - k_3/k_1$ space. However, not all the configurations exist inside the simulation and therefore there are some gaps in the distribution.

The value of the $B_{\text{hydro}}/B_{\text{DMO}}$ in Fig. 1 is colour-coded using two colour bars: one for the unmatched haloes (top row) and another for the matched haloes (bottom row). The arrows pointing to the top right corner of each subplot indicate the ratio for equilateral triangle configuration. The right-hand (left-hand) column shows the results for the case where the shot noise has (not) been subtracted.

We can see from Fig. 1 that the impact of the choice of configuration appears to depend on the selection of haloes adopted (unmatched versus matched) and on whether the bispectrum has been shot noise subtracted. For the shot noise-subtracted case with unmatched haloes (arguably the most useful comparison, since the spectra should be shot noise subtracted and the unmatched case is closer to an observational selection), the ratio of the bispectrum ranges from 1.2 to 1.6, with the equilateral case yielding close to the maximal ratio. The fact that the ratio depends on the choice of triangular configuration suggests that it may be possible to extract additional information about the impact of baryons by computing the bispectra for a range of different configurations. We plan to explore this possibility further in future work and proceed with the standard equilateral configuration for the remainder of this study, noting that there is some sensitivity to the results on the choice of configuration.

3.4 Shot noise and resolution considerations

Here, we present an analysis of the impact of shot noise on our results and its dependence on the adopted triangular configuration. We also discuss the possible impact of non-Poisson shot noise as well as possible limitations due to finite mass and force resolution as well as finite volume effects. In short, we conclude that resolution and box size limitations are unlikely to be important for our results and conclusions, while the contribution of shot noise to the halo power spectra and bispectra becomes important (i.e. comparable to the true clustering signal) on scales of $k \approx 1 \, h \, \text{Mpc}^{-1}$. Thus, a degree of caution is warranted when examining the absolute power beyond this scale, although we expect the *relative* trends in $P(k)$ and $B(k_1, k_2, k_3)$ (e.g. as we increase feedback or the summed neutrino mass) to be robust. Readers who are mainly interested in the impact of baryon physics and neutrino free-streaming on the halo power spectrum and bispectrum may wish to skip ahead to Section 4.

We first consider the impact of Poisson shot noise. In Fig. 2, we show the absolute power spectrum and bispectrum (both DMO and hydro) with and without shot noise subtraction, which illustrates clearly the relative contribution of shot noise as a function of scale. In the case of the bispectrum, we show the two Poisson terms separately (see equation 5 for details). The plot also shows individual data points with a theoretical estimate of the errors in the Gaussian approximation for both power spectrum and bispectrum. For the theoretical Gaussian errors, we use a standard approach where the error is the square root of the covariance matrix. For details of how to calculate the covariance matrix, see equations 2.34 and 2.39 in Yankelevich & Porciani (2019).

Two vertical lines are added to Fig. 2 to indicate the 50 per cent shot noise contribution for both the power spectrum and bispectrum. (The k values at which Poisson shot noise contributes 1 per cent and 10 per cent are also presented in Table 2.) From this figure we conclude that the contribution from Poisson shot noise becomes comparable to that of the true clustering signal at scales of $k \approx 1 \, h \, \text{Mpc}^{-1}$. Thus, a degree of caution is warranted when examining the absolute power beyond this scale, although we expect the *relative* trends (e.g. as we increase feedback or the summed neutrino mass) to be robust.




Figure 2. The power spectrum and the bispectrum measured from the fiducial hydro (orange lines) and DMO (black lines) BAHAMAS simulations with (dash-dotted lines with errorbars) and without (solid lines) shot noise subtraction assuming equilateral triangles. The haloes mass cut is the same as in our fiducial analysis (e.g. Fig. 4) for the ‘unmatched’ selection. The error bars indicate the theoretically estimated Gaussian errors. In the left-hand panel, the dotted lines represent the Poisson shot noise for the power spectrum (the last term in equation 4). In the right-hand panel, the shot noise is shown according to the legend: 1 – full shot noise for DMO bispectrum (second and third terms in equation 5), 3 – part of shot noise for DMO bispectrum (third term in equation 5), 5 – part of shot noise for DMO bispectrum (second term in equation 5), all cyan colour. The purple lines 2, 4, 6 are the same as 1, 3, 5 but for the hydro bispectrum. The two vertical lines in both panels indicate the 50 per cent shot noise contribution for DMO (grey colour) and hydro (navy colour) power spectrum and bispectrum.

Table 2. Values of wave-vector k at which the shot noise contributes 1 per cent, 10 per cent, or 50 per cent to the absolute power spectrum or bispectrum, and the values of the shot noise (SN) itself. Note that the units of shot noise are $h^{-3}\text{Mpc}^3$ for the power spectrum and $h^{-6}\text{Mpc}^6$ for the bispectrum.

	k , 1 per cent ($h \text{ Mpc}^{-1}$)	SN	k , 10 per cent ($h \text{ Mpc}^{-1}$)	SN	k , 50 per cent ($h \text{ Mpc}^{-1}$)	SN
P_{DMO}	0.063	88.71	0.346	88.71	1.288	88.71
P_{hydro}	0.094	78.91	0.377	78.91	1.539	78.91
B_{DMO}	0.064	2.24×10^6	0.064	2.24×10^6	0.879	5.06×10^4
B_{hydro}	0.036	3.74×10^6	0.064	2.09×10^6	1.476	2.58×10^4

Note that the shot noise contribution to the bispectrum depends on the adopted triangular configuration. Fig. 3 demonstrates that the shot noise ratio is smaller for the squeezed, elongated, and folder triangles, and is generally larger for isosceles and equilateral configurations. The largest shot noise case corresponds to the scalene configuration.

In our fiducial analysis in Section 4, we consider the Poissonian shot noise contribution only. However, on small scales we expect the pure Poissonian assumption to break down since haloes can be strongly clustered (e.g. satellites in a galaxy group). Previous work has shown that the non-Poisson contribution to the shot noise is generally subdominant to that of the Poissonian component. For example, using the halo model, Ginzburg, Desjacques & Chan (2017) have calculated that the non-Poisson shot noise is approximately half that of the Poisson shot noise contribution at largest k values (smallest scales) considered, a result which they confirmed using N -body simulations. At smaller k values (large scales), the contribution from non-Poisson shot noise is generally negligible. Thus, we expect that, at worse, our estimate of the shot noise is biased low by 50 per cent and, therefore, our estimate of the true intrinsic bispectrum signal small scales ($k \approx 1 h \text{ Mpc}^{-1}$) may be biased high up to 25 per cent. However, we do not expect our relative findings (i.e. ratios of bispectra) to be as adversely affected, since the absolute spectra will all be similarly affected.

We now turn our attention to resolution considerations. Note that the force resolution limit (taken to be the Plummer equivalent gravitational softening of $4 \text{ kpc } h^{-1}$) of the simulations is $k \approx$

$1507 h \text{ Mpc}^{-1}$, which is well beyond the scales we analyse. Thus, limitations due to force resolution are unimportant for our analyses. However, the force resolution is not the only resolution limitation when considering halo clustering. In particular, there are at least three relevant effects: halo exclusion, finite simulation volume, and the lower mass limit of the halo finder catalogue. Halo exclusion implies there will be no clustering below the halo radius because the haloes will be touching, which will also impact the noise on small scales (see Baldauf et al. 2013 for more details). In our analyses, however, we examine the clustering of subhaloes rather than FOF groups. Subhaloes identified with SUBFIND can be in close proximity to each other and can even overlap spatially (particles are assigned to a single subhalo based on an energy unbinding criterion). Thus, we do not anticipate halo exclusion effects to be relevant for our analyses. With regards to simulation volume effects, the finite volume can result in biased halo number densities for very large haloes, since they are rare. For the majority of this study, however, we cut off haloes at mass $10^{12.5} h^{-1} M_{\odot}$, which is a safe option even in moderate size boxes (see van Daalen et al. 2014). Even without this cut, our results are likely to be insensitive to the rarest and most massive haloes (which, in BAHAMAS, are $\sim 10^{15} M_{\odot}$), since there are orders of magnitude more haloes at lower masses which dominate the signal (see Fig. 5). Finally, with regards to the lower mass limit, we selected a lower cut-off of $10^{11.5} h^{-1} M_{\odot}$ corresponding to approximately 100 particles. These are therefore relatively well-resolved haloes, in terms of being able to robustly estimate their masses and positions. We have also




Figure 3. The relative contribution of the Poisson shot-noise to the bispectrum as a function of the triangular configuration. For consistency with Fig. B2 we use the same $k_1 = 2.92 h \text{ Mpc}^{-1}$.

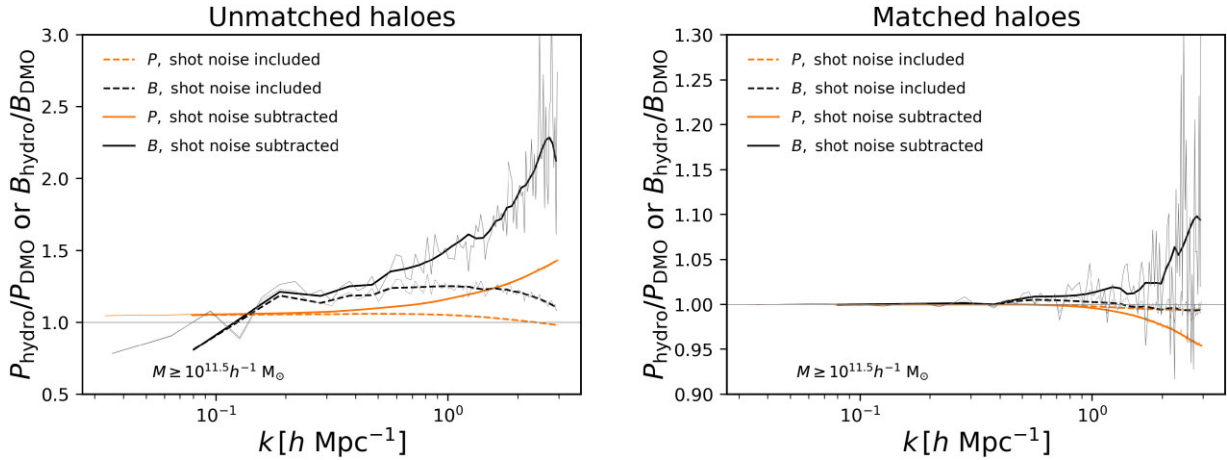


Figure 4. The ratio of the power spectrum (orange lines) and the bispectrum (black lines) measured from the fiducial hydro and DMO BAHAMAS simulations with (solid lines) and without (dashed lines) shot noise subtraction assuming equilateral triangles. All haloes (centrals and satellites) with masses exceeding $10^{11.5} h^{-1} M_\odot$ are selected in the left-hand panel (‘unmatched’ case). For the matched case (right-hand panel), haloes are selected from the DMO case according to the mass criterion above and the same haloes are identified in the hydro run (regardless of their mass), hence a common set of haloes are used. The thick lines show the ratio with the smoothing procedure applied, while thin lines show the unsmoothed data.

verified that the BAHAMAS halo mass function follows that of other results in the literature (e.g. Tinker et al. 2010) down to this limit.

4 RESULTS

In this section, we present our main results on the impact of baryon physics and neutrino free-streaming on the halo bispectrum. We first examine the impact of baryons for the massless neutrino case in Section 4.1 and then we explore the effects of neutrinos and their degeneracy with baryonic effects in Section 4.2.

4.1 Impact of baryon physics

In Fig. 4, we plot the ratio between the fiducial hydro and the DMO BAHAMAS simulations (both with massless neutrinos) of the halo power spectrum and bispectrum for equilateral triangles. We consider both the raw spectra computed by BSKIT which includes a shot noise contribution and shot noise-subtracted spectra, where we have assumed the shot noise follows a Poisson distribution as given in equations (4) and (5).

We note that when evaluating the shot noise contribution to the bispectrum in equation (5), the power spectra should be shot noise-subtracted (according to equation 4). Since the power spectrum differs between the hydro and DMO cases, the bispectrum shot noise term will also, in general, be different for the hydro and DMO cases,

even for a matched set of haloes. Furthermore, it is clear that the bispectrum shot noise is not a constant (i.e. it depends on scale), unlike that for the power spectrum. As can be seen in Fig. 4, shot noise does not greatly affect large scales, but starting from $k \gtrsim 0.1 h \text{ Mpc}^{-1}$ it begins to have an increasingly important effect. Therefore, we will subtract the shot noise by default for the rest of the paper and we focus on the shot noise-subtracted results.

Previous studies have shown that including galaxy formation physics in hydro simulations leads to a characteristic suppression in the non-linear matter power spectrum, which is due primarily to the ejection of baryons by AGN feedback (e.g. van Daalen et al. 2011, 2020; Mummery et al. 2017; Chisari et al. 2019; Foreman et al. 2020; Stafford et al. 2020). This effect was also shown to produce a suppression in the matter bispectrum by Foreman et al. (2020; see also Aricò et al. 2021). However, we show here that, at least at face value, haloes appear to show the opposite effect. In particular, for the unmatched case, the ‘hydro’ power spectrum and bispectrum are always greater than that of the corresponding ‘DMO’ simulation, implying that haloes selected to be above a certain mass appear to be more clustered in a hydro simulation than in the corresponding collisionless one. Note that the apparent slight suppression in the bispectrum at large scales (small k) is due to sampling noise from finite box size effects.

In the right-hand panel of Fig. 4, we consider the ratio of the hydro to the DMO case for a given (‘matched’) set of haloes

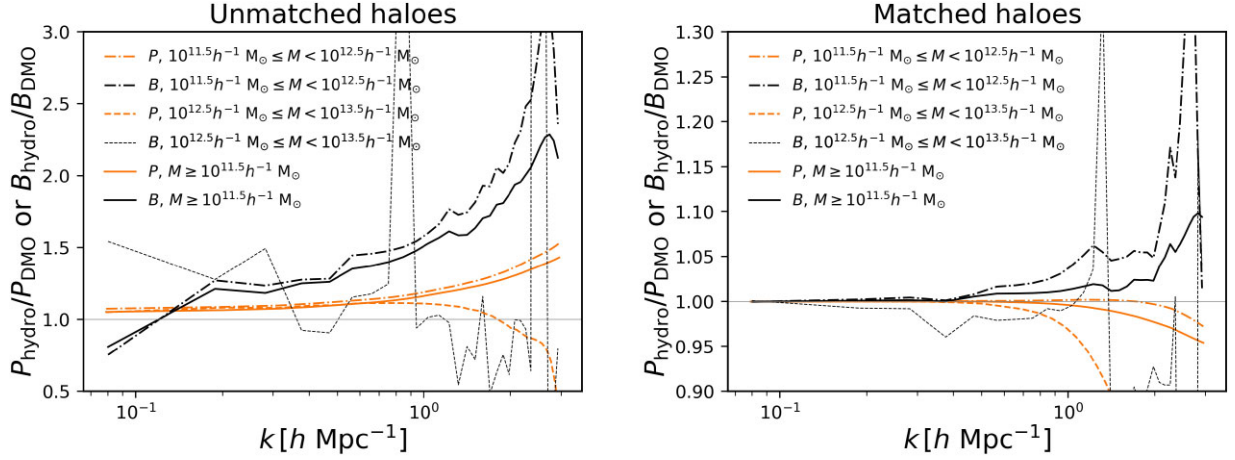


Figure 5. The ratio of the power spectrum (orange lines) and the bispectrum (black lines) measured from the hydro to DMO BAHAMAS simulations with shot noise subtraction for equilateral triangles. Different line styles denote different haloes mass cuts (solid lines demonstrate the same as at Fig. 4). The cut is done separately for hydro and DMO runs for the unmatched case. For the matched case, the cut is made according to DMO run for all lines.

common to the two simulations, rather than selecting haloes above a given mass limit (as in the left-hand panel). Here, we see that on relatively large scales ($k \lesssim 0.5 \, h \, \text{Mpc}^{-1}$) the ratio now goes to unity, indicating no appreciable difference in the clustering of haloes in the two simulations on large scales. The same conclusion was made by van Daalen et al. (2014) when examining the two-point correlation functions of a matched set of haloes in the OWLS simulations.

As we are examining a common set of haloes in this comparison, it implies that the differences for the unmatched case (left-hand panel) are primarily driven by a difference in the haloes that are selected when using a fixed mass cut in that case. In particular, AGN feedback removes baryons from haloes in the hydro case, reducing their overall masses (e.g. Cui, Borgani & Murante 2014; Cusworth et al. 2014; Velliscig et al. 2014; Schaller et al. 2015). Consequently, some haloes that would (in the absence of feedback) have been just above the mass cut are now below the threshold for selection and are therefore not selected. The net result is that introducing baryon physics (feedback) results in the selection of intrinsically (in the absence of feedback) more massive systems. Since more massive haloes are more strongly biased in their clustering with respect to the mean matter density (e.g. Tinker et al. 2010), the halo clustering is enhanced on large scales in the hydro simulation with respect to the DMO simulation when selecting haloes above a given mass limit. For a common set of haloes (as in the right-hand panel), however, there is no appreciable difference in the clustering on large scales.

On small scales ($k \gtrsim 1 \, h \, \text{Mpc}^{-1}$), however, the clustering is altered by hydrodynamics even for a common set of haloes. In particular, we find that the power spectrum is slightly suppressed (at the 5 per cent level up to $k \approx 3 \, h \, \text{Mpc}^{-1}$), whereas the bispectrum is enhanced (at the 10 per cent level up to $k \approx 3 \, h \, \text{Mpc}^{-1}$). On these scales, the clustering is almost certainly dominated by the so-called one-halo term, corresponding to satellites of massive haloes (groups and clusters). Thus, deviations from unity are likely signifying differences in how the satellites are spatially distributed in massive haloes in the hydro case with respect to the DMO case. Physically, such differences can potentially arise due to changes in the efficiency of tidal disruption, with gas ejection likely making satellites more susceptible to stripping at large (satellite-centric) scales, while star formation and adiabatic contraction would likely make disruption of the central regions of satellites more difficult to

strip (with respect to the DMO case). However, we cannot easily rule out that differences in the performance of the substructure finder (SUBFIND) between the DMO and hydro cases is also potentially playing a role here. To test this, higher resolution simulations and/or alternative substructure finders (e.g. phase space-based finders) could be employed, but we leave this for future work.

4.1.1 Halo mass dependence

In the analysis presented above, we explored the impact of baryon physics on the halo bispectrum at $z = 0$ when selecting all haloes above $10^{11.5} \, h^{-1} \, \text{M}_\odot$, corresponding roughly to a stellar mass selection of $M_* \gtrsim 10^{10} \, \text{M}_\odot$. Here, we attempt to examine the halo mass dependence of the results, noting that the effects of stellar and AGN feedback on the baryon fractions are expected to be halo mass dependent (e.g. Schaller et al. 2015). We therefore expect the effects of feedback on the halo bispectrum to be sensitive to the halo mass range examined. However, a limitation of this study is that very large samples of haloes are required to accurately measure the bispectrum, thus breaking the sample up into multiple mass bins will yield less statistically robust results. Nevertheless, we proceed here by examining the bispectrum (and power spectrum) in two adjacent mass bins of $10^{11.5} \, h^{-1} \, \text{M}_\odot \leq M < 10^{12.5} \, h^{-1} \, \text{M}_\odot$ and $10^{12.5} \, h^{-1} \, \text{M}_\odot \leq M < 10^{13.5} \, h^{-1} \, \text{M}_\odot$ and we compare these with our previous selection of $M \geq 10^{11.5} \, h^{-1} \, \text{M}_\odot$ (Fig. 5). Note that above the mass limit $10^{11.5} \, h^{-1} \, \text{M}_\odot$ there are 721 459 haloes (or 100 per cent) in the DMO simulation and that the lower mass bin ($10^{11.5} \, h^{-1} \, \text{M}_\odot \leq M < 10^{12.5} \, h^{-1} \, \text{M}_\odot$) contains 638 832 of these haloes (or 88.6 per cent) while the higher mass bin contains only 76 016 haloes (or 10.5 per cent).

In Fig. 5, we show the ratio of the halo power spectra and bispectra (hydro to DMO) for the different halo mass selections for both the unmatched (left) and matched (right) cases. The lower mass bin roughly follows the same trend as our total halo sample for the unmatched haloes, though the effects of baryons on both the power spectrum and bispectrum are slightly more pronounced in the low-mass bin compared to the total halo selection. The results for the higher mass bin ($10^{12.5} \, h^{-1} \, \text{M}_\odot \leq M < 10^{13.5} \, h^{-1} \, \text{M}_\odot$) are, unfortunately, too noisy to make any strong quantitative conclusions. Larger simulation volumes are required to examine

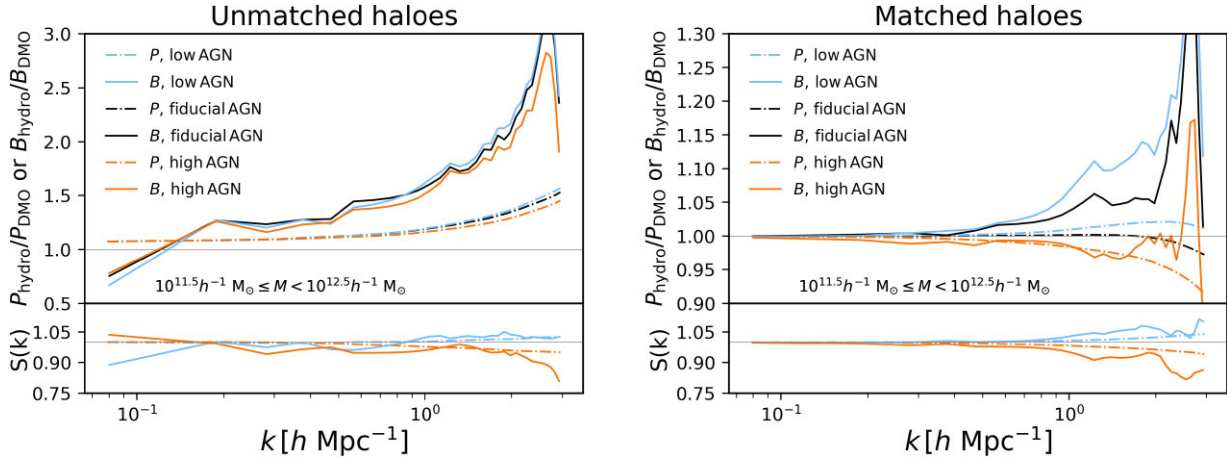


Figure 6. *Upper panels:* The ratio of the power spectrum (dot–dashed lines) and the bispectrum (solid lines) measured from the hydro to DMO BAHAMAS simulations with shot noise subtraction for equilateral triangles for halo masses in range $10^{11.5} h^{-1} M_{\odot} \leq M < 10^{12.5} h^{-1} M_{\odot}$. Different line colours represent variations the efficiency of AGN feedback. *Lower panels:* The ratio $S(k)$ of the power spectrum and the bispectrum ratios between low/high AGN feedback models with respect to the fiducial (tuned AGN) model.

the halo bispectrum of high-mass objects. Qualitatively, though, the results are consistent with a reduced impact due to baryon physics in comparison to the lower mass bin. One can also infer this from the comparison of the total selection and lower mass bin comparison, since the latter is more strongly affected by baryons and the only difference in the selections is that the total case also contains high-mass haloes.

For the case of matched haloes (i.e. where haloes are selected according to some mass criterion in the DMO simulations and bijectively matched to the hydro simulation), shown in the right-hand panel of Fig. 5, we find much the same behaviour as in the right-hand panel of Fig. 4. Specifically, for a given set of haloes, baryon physics does not significantly alter the halo power spectrum or bispectrum on large scales ($k \lesssim 0.5 h \text{ Mpc}^{-1}$) independent of halo mass. On small scales (the one-halo regime), however, the spectra are altered and the effects are largest for the lower mass bin (in agreement with the left-hand panel) and we again find that the bispectrum is enhanced due to baryon effects whereas the power spectrum is suppressed (similar to the matter power spectrum).

4.1.2 Feedback variations

In addition to the fiducial calibrated BAHAMAS model, McCarthy et al. (2018) introduced two additional simulations which varied the strength of the AGN feedback: ‘low AGN’ and ‘high AGN’. Specifically, they varied the subgrid AGN heating temperature parameter by ± 0.2 dex around the fiducial $\log_{10}[\Delta T_{\text{heat}}] = 7.8$, which approximately brackets the upper and lower envelopes of the observed gas fractions of galaxy groups and clusters, as inferred from resolved X-ray observations (see fig. 3 of McCarthy et al. 2018). Note that ΔT_{heat} is the most important parameter in the model for determining the gas fractions of groups/clusters. Higher values of ΔT_{heat} correspond to more energetic (and more bursty) AGN feedback episodes and vice-versa.

We show in Fig. 6 the effect of varying ΔT_{heat} on the halo power spectrum and bispectrum for a halo mass selection of $10^{11.5} h^{-1} M_{\odot} \leq M < 10^{12.5} h^{-1} M_{\odot}$, once again through the ratio of these statistics measured from the hydro simulations computed with respect to the DMO case. For the unmatched case, the power spectrum

and bispectrum ratios for all three AGN models show an enhancement due to feedback. Interestingly, it is the ‘low AGN’ model which shows the largest effect particularly on small scales, whereas for the *matter* power spectrum and bispectrum the opposite is true (e.g. fig. 13 of Foreman et al. 2020, see also van Daalen et al. 2020). In the case of the matter clustering, the interpretation is straightforward: more energetic feedback leads to increased baryon ejection which reduces the overall matter clustering. For the halo clustering, particularly on small scales which are dominated by the one-halo term (satellites of groups/clusters), it is plausible that particularly energetic feedback leads to satellites being more susceptible to environmental processing (tidal heating/stripping), reducing the power on small scales compared to a run with less energetic feedback. It is interesting that in the matched halo case (right-hand panel of Fig. 6), the fiducial and low AGN runs show enhanced halo bispectra (and power spectra, for low AGN) relative to the DMO case, whereas the high AGN shows suppressed halo clustering. This suggests there is a relatively fine line between baryon effects enhancing the clustering on small scales relative to DMO (e.g. due to star formation and adiabatic contraction of the satellites subhaloes) and enhanced tidal stripping due to baryon ejection.

The upshot of the sensitivity of the halo clustering due to baryon effects is that the simulations probably cannot be used to robustly predict whether there should be an enhancement or a suppression relative to DMO simulations on small scales (for a common set of haloes), as the answer appears to depend sensitively on the balance between feedback and environmental physics. This is in contrast to the findings of Foreman et al. (2020; see e.g. their fig. 13), in which the matter bispectra of the three BAHAMAS hydro simulations consistently show a suppression relative to the DMO case, whereas our measurements of the *halo* bispectrum show either a suppression or an enhancement depending on the AGN feedback strength. This potentially complicates the interpretation of observational measurements, since a probe such as the cosmic shear bispectrum would observe a suppression while the galaxy three-point function might be suppressed or enhanced. A joint analysis of the matter and halo bispectra would therefore be useful for simultaneously constraining the effects of feedback and environmental physics on small scales. In this regard, it is noteworthy that the effects are typically much stronger in the bispectrum than in the power spectrum (consistent

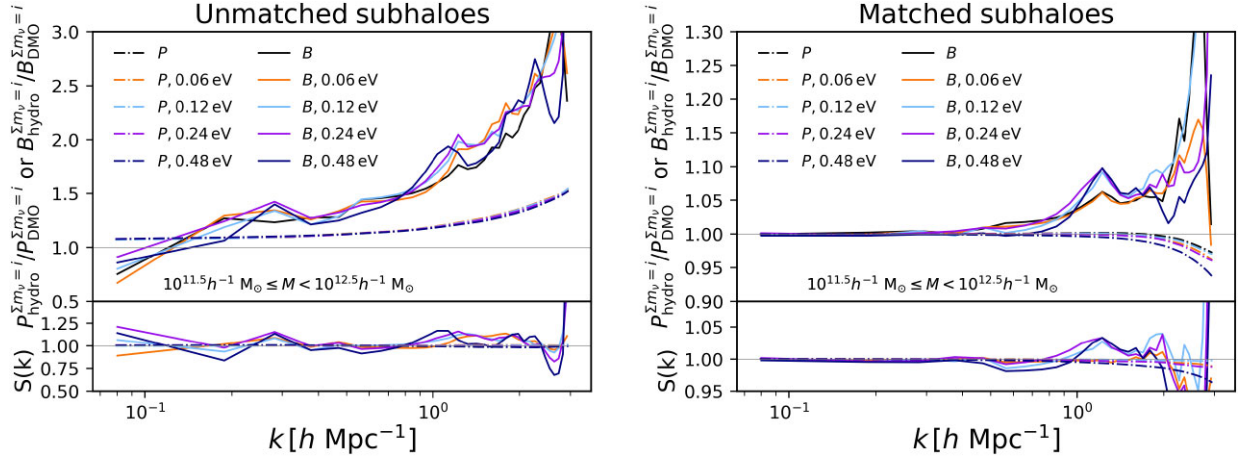


Figure 7. *Upper panels:* The ratio of the power spectrum (dot–dashed lines) and the bispectrum (solid lines) measured from the hydro to DMO BAHAMAS simulations with shot noise subtraction for equilateral triangles for halo masses in range $10^{11.5} h^{-1} M_{\odot} \leq M < 10^{12.5} h^{-1} M_{\odot}$. Different line colours represent different summed neutrino masses Σm_{ν} . Note that the summed neutrino mass is varied simultaneously here for the hydro and DMO runs. *Lower panels:* As in lower panels of Fig. 6 but ratio $S(k)$ is of the massive neutrino cases with respect to the fiducial massless neutrino case.

with the findings of Foreman et al. (2020), suggesting that future measurements of the small-scale bispectrum are a promising tool for studying these effects.

4.2 Impact of neutrino free-streaming

Of the various possible extensions to the standard Λ CDM model, massive neutrinos are perhaps the most well motivated, as the results of atmospheric and solar oscillation experiments imply that the three active species of neutrinos have a *minimum* summed mass, Σm_{ν} , of at least 0.06 eV (0.1 eV) when adopting a normal (inverted) hierarchy (see Lesgourgues & Pastor 2006 for a review). At present, cosmological studies place the strongest upper limits on the summed neutrino mass. Galaxy clustering studies, in particular, have a rich history of placing important constraints on the summed neutrino mass (e.g. Tegmark et al. 2004; Cole et al. 2005; Sánchez et al. 2006; Beutler et al. 2014; Abbott et al. 2018), with its sensitivity coming from the fact that neutrinos suppress the growth of structure (i.e. reduce the clustering amplitude relative to a massless neutrino cosmology) on scales smaller than the associated free-streaming scale. In this study, we use the massive neutrinos extension of BAHAMAS (Mummery et al. 2017; McCarthy et al. 2018) to explore the joint effects of baryon physics and neutrinos, and their possible degeneracy, on the halo power spectrum and bispectrum. In addition to the massless neutrino cases explored above, BAHAMAS has four variations of summed neutrino mass, with $\Sigma m_{\nu} = 0.06, 0.12, 0.24$, or 0.48 eV (Table 1). As mentioned in Section 2, neutrinos are implemented in the simulations using the semilinear method of Ali-Haïmoud & Bird (2013), which is sometimes referred to as a ‘linear response’ method. For further details of the implementation within BAHAMAS, we refer the reader to McCarthy et al. (2018).

In Fig. 7, we show the ratio of the power spectrum and the bispectrum (hydro to DMO) where the summed neutrino mass is systematically varied for both the hydro and DMO runs (simultaneously). In essence, this comparison explores whether the relative impact of baryon physics on the halo clustering is dependent on the summed neutrino mass. If the ratio varies significantly between the simulations, this implies that the impact of baryon physics is dependent on the summed neutrino mass, whereas if the ratio does not vary between the simulations then the effects of baryon physics

are independent of (i.e. separable from) those of massive neutrinos. As shown in Fig. 7 (focusing on the bottom subpanels, in particular, which show the ratio with respect to the massless neutrino case), there is no evidence for a strong systematic dependence on the summed neutrino mass, with the results being very similar to the massless case. The noise in the halo bispectrum measurements means we can only place an upper limit on the potential effect but, for the matched case, the impact of baryons on the halo bispectrum agrees amongst all of the neutrino simulations to typically a few per cent accuracy (and to higher accuracy for the halo power spectrum). Thus, the impact of baryons and neutrinos appear to be separable effects for the halo power spectrum and bispectrum, as previously found for the matter clustering in Mummery et al. (2017) and van Daalen et al. (2020).

In Fig. 8, we explore the combined effects of baryon physics and neutrino free-streaming on the halo clustering, by taking the ratio of the hydro simulations of varying neutrino mass to the DMO run with *massless* neutrinos. Thus, both feedback and neutrinos will impact this metric. In fact, there are four effects present when considering the left-hand panel of Fig. 8 (unmatched haloes); as both feedback and free-streaming affect the halo mass (i.e. alter the halo selection) and both feedback and free-streaming affect the clustering (although we have shown that the clustering is only affected on small scales by baryon physics after we control for the halo mass alteration). Neutrinos alter the halo mass because they represent a form of dark matter that is effectively unable to collapse, whereas all of the dark matter can collapse in the massless neutrino case. Mummery et al. (2017) have shown that it is the largest haloes (massive clusters) that are most strongly affected in terms of halo mass.

In the left-hand panel of Fig. 8, we see that the halo power spectra and bispectra are typically enhanced with respect to the DMO with massless neutrinos case. Increasing the summed neutrino mass tends to reduce this enhancement, implying that the free-streaming reduction in clustering amplitude is a more significant effect than the reduction in halo mass that is due to neutrinos. However, for this range of neutrino masses, the change in halo mass due to baryon ejection is sufficient to overcome the impact of neutrinos, which is why there is an overall enhancement in the clustering rather than a suppression. The one exception to this is the $\Sigma m_{\nu} = 0.48$ eV case, which shows a slight suppression on quasi-linear scales. Here, the neutrino suppression is sufficiently strong to overcome the enhancement due to baryon ejection.

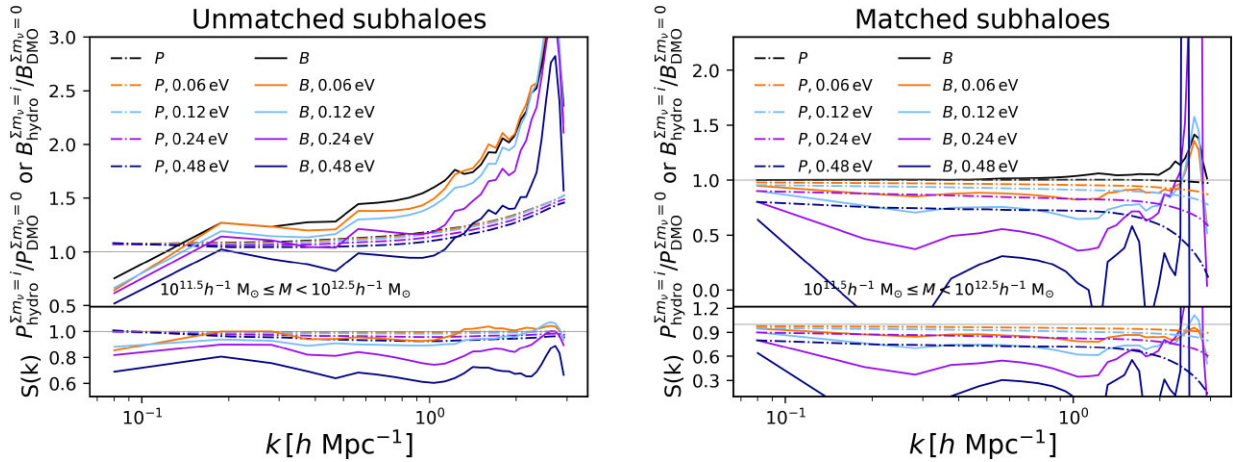


Figure 8. As in Fig. 7 except here the DMO case corresponds to massless neutrinos. Thus, both baryon physics and neutrinos will alter the ratio of the hydro simulation with respect to the DMO run.

In the right-hand panel of Fig. 8, we do the same comparison but now for a common set of haloes, selected from the DMO massless neutrino run and bijectively matched to each of the hydro simulations with varying neutrino masses. As explained before, by examining a common set of haloes we are removing the impact of changes in halo mass, due to both feedback and neutrino free-streaming, on the halo selection. Here, we see a strong suppression in the clustering which is present on virtually all scales plotted, apart from the very smallest scales where the enhancement due to baryon physics begins to dominate. In fact, the suppression will be present out to approximately the free-streaming scale, which corresponds to $k_{fs} \approx 0.05 h \text{ Mpc}^{-1}$ for $\Sigma m_\nu = 0.06 \text{ eV}$ (see equation 9 of Ali-Haïmoud & Bird 2013). In common with the impact of feedback, we find that the bispectrum is affected much more significantly than is the power spectrum, although we note that estimates of the bispectrum are more affected by noise (sample variance) than is the power spectrum. In principle, though, upcoming precision measurements of the bispectrum may be able to place strong constraints on both cosmology and models of galaxy formation simultaneously, by using small scales to probe feedback and environmental processing and large scales to help constrain the neutrino mass. Certainly the combination of the power spectrum and bispectrum will allow for much more powerful constraints than may be obtained from the power spectrum alone.

Finally, in the above discussion we have neglected the potential degeneracies between the summed neutrino mass, σ_8 , and baryons. It is known, for example, that using clustering data alone which is restricted to quasi-linear and non-linear scales (e.g. as in current cosmic shear surveys) that Σm_ν and σ_8 can be strongly degenerate. Because the BAHAMAS neutrino simulations are ‘CMB normalized’, varying Σm_ν leads to a variation in σ_8 as well. To explore the degeneracy in detail would require additional simulations which hold σ_8 fixed as the neutrino mass is varied (or vice-versa), which is beyond the scope of this study. We note, however, that with the increasing quality of weak lensing and spectroscopic galaxy clustering surveys, one can exploit the fact that neutrino free-streaming, variations in σ_8 , and the impact of baryons on the matter and halo clustering all have different scale and redshift dependencies (see e.g. Mummery et al. 2017 who compared the scale and redshift dependencies of baryons and massive neutrinos on various clustering statistics). We therefore expect upcoming clustering studies to be able to place much stronger constraints on these effects than is currently possible.

5 SUMMARY AND CONCLUSIONS

The primary goal of this paper was to demonstrate the power of the halo bispectrum as a tool for testing cosmological models and galaxy formation physics (particularly AGN feedback). To this end, we measured the halo power spectrum and the halo bispectrum from the BAHAMAS hydro and dark matter simulation suites. Aside from the fiducial calibrated BAHAMAS hydro simulation (which was calibrated to reproduce the galaxy stellar mass function and the gas fractions of galaxy groups/clusters), the suite contains runs which vary the strength of the AGN feedback and the summed mass of neutrinos, allowing us to explore the sensitivity of the halo clustering to these effects. We modified the publicly available matter clustering code BSKIT (Foreman et al. 2020) to measure the halo clustering in BAHAMAS.

The main results of our study may be summarized as follows:

- (i) The bispectrum measured from the simulations in general depends on the adopted triangular configuration. It is common place to adopt an equilateral configuration mostly for convenience. We explored the dependence of our main results (the relative impact of baryon physics on the halo bispectrum) on the adopted triangular configuration in Fig. 1 (see also Appendix B), finding that the equilateral choice tends to maximize the impact of baryons. We adopted the equilateral choice for the remainder of the analysis but note that the results are somewhat sensitive to the choice of triangular configuration. In the future, it may be possible to exploit this sensitivity as a further test of baryon and cosmological physics.
- (ii) When selecting haloes above a given mass, baryon physics (particularly AGN feedback) tends to *enhance* the halo power spectrum and bispectrum relative to a collisionless (DMO) case (see left-hand panel of Fig. 4). This enhancement is driven by the change in halo mass due to AGN feedback. Specifically, haloes just above the selection threshold in the DMO case have their masses reduced due to gas ejection and therefore not selected in the hydro simulation case. Consequently, the mean bias of haloes selected in the hydro case is larger than that of the DMO case, giving rise to an apparent enhancement in the clustering. Comparing the clustering of a *common set* of haloes in the hydro and DMO cases (see right-hand panel of Fig. 4) demonstrates that baryon physics does not appreciably affect the clustering on large scales of $k \lesssim 0.5 h \text{ Mpc}^{-1}$. On small scales, dominated by the one-halo term (satellites of massive groups and clusters), the power spectrum and bispectrum can

be significantly affected (≈ 10 per cent level) and, interestingly, we find for the fiducial calibrated BAHAMAS model that the bispectrum is enhanced with respect to the DMO case, while the power spectrum is suppressed.

(iii) We find that the effects described above are a function of halo mass, with lower masses ($\sim 10^{12} M_{\odot}$) being more significantly affected by baryon physics (see Fig. 5).

(iv) We explored the sensitivity of the results to variations in galaxy formation physics by using two variations from the BAHAMAS suite ('low AGN', 'high AGN'). Interestingly, we find that reducing the efficiency of AGN feedback results in an increased enhancement in the halo bispectrum on small scales (see right-hand panel of Fig. 6), which is plausibly due to the fact satellites are more resilient against tidal disruption in this case. Increasing the efficiency of AGN feedback can actually change the sign of the effect, such that the halo bispectrum is suppressed with respect to the DMO case. This suggests that there is a very fine balance between star formation/adiabatic contraction (which tend to make haloes less susceptible to tidal heating/stripping) and gas ejection (which make them more susceptible). Small-scale measurements of the halo bispectrum therefore offer a means of simultaneously probing feedback and environmental physics.

(v) Using the massive neutrinos extension of BAHAMAS, we explored the joint effects of neutrino free-streaming and baryon physics (and their potential degeneracy) on the halo clustering. We showed that the effects of baryon physics are largely independent of (separable from) the effects of neutrinos (see Fig. 7). For a common set of haloes, we find neutrinos strongly suppress the halo bispectrum on all but the smallest scales probed in this study (the smallest scales are dominated by baryon physics which can either enhance or suppress the bispectrum, as discussed above). See the right-hand panel of Fig. 8.

Our results have demonstrated that the halo power spectrum and particularly the halo bispectrum are highly sensitive to neutrino effects on the largest scales and baryon physics (AGN feedback and star formation) on the smallest scales. This bodes well for upcoming precision measurements of the bispectrum for surveys like Euclid, LSST, and DESI. To capitalize on these measurements and the implied sensitivity what is required now is a large suite of very large volume cosmological simulations that systematically and simultaneously explore variations in cosmological parameters and galaxy formation physics. This is the aim of ongoing work.

ACKNOWLEDGEMENTS

The authors thank the referees for constructive suggestions which have improved the paper. They also thank Andreea Font, Jaime Salcido, and Amol Upadhye for helpful feedback on the paper and acknowledge helpful discussions with Davit Alkhanishvili, William Coulton, Alexander Eggemeier, ChangHoon Hahn, Andrea Oddo, and Joop Schaye. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no 769130). SGS acknowledges the Science and Technology Facilities Council (STFC) doctoral studentship. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC High Performance Computing (HPC) Facility. The equipment was funded by the Department for Business, Energy & Industrial Strategy (BEIS) capital funding via STFC capital grants ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University and STFC

operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The data underlying this article may be shared on reasonable request to the corresponding author.

REFERENCES

- Abbott T. M. C. et al., 2018, *Phys. Rev. D*, 98, 043526
- Agarwal N., Desjacques V., Jeong D., Schmidt F., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 021
- Alam S. et al., 2017, *MNRAS*, 470, 2617
- Ali-Haïmoud Y., Bird S., 2013, *MNRAS*, 428, 3375
- Alkhanishvili D., Porciani C., Sefusatti E., Biagetti M., Lazanu A., Oddo A., Yankelevich V., 2022, *MNRAS*, 512, 4961
- Angulo R. E., Foreman S., Schmittfull M., Senatore L., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 039
- Aricò G., Angulo R. E., Hernández-Monteagudo C., Contreras S., Zennaro M., 2021, *MNRAS*, 503, 3596
- Baldauf T., Seljak U., Smith R. E., Hamaus N., Desjacques V., 2013, *Phys. Rev. D*, 88, 083507
- Baldauf T., Mercolli L., Zaldarriaga M., 2015, *Phys. Rev. D*, 92, 123007
- Barreira A., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 031
- Bernardeau F., Colombi S., Gaztañaga E., Scoccimarro R., 2002, *Phys. Rep.*, 367, 1
- Bernardeau F., Crocce M., Scoccimarro R., 2008, *Phys. Rev. D*, 78, 103521
- Bernardeau F., Crocce M., Scoccimarro R., 2012, *Phys. Rev. D*, 85, 123519
- Beutler F. et al., 2014, *MNRAS*, 444, 3501
- Bocquet S., Saro A., Dolag K., Mohr J. J., 2016, *MNRAS*, 456, 2361
- Chen J. Z., Upadhye A., Wong Y. Y. Y., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 065
- Chisari N. E. et al., 2019, *Open J. Astrophys.*, 2
- Chudaykin A., Ivanov M. M., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 034
- Cole S. et al., 2005, *MNRAS*, 362, 505
- Crocce M., Scoccimarro R., 2006, *Phys. Rev. D*, 73, 063519
- Cui W., Borgani S., Murante G., 2014, *MNRAS*, 441, 1769
- Cusworth S. J., Kay S. T., Battye R. A., Thomas P. A., 2014, *MNRAS*, 439, 2485
- Davis M., Peebles P. J. E., 1983, *ApJ*, 267, 465
- Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
- Dolag K., Komatsu E., Sunyaev R., 2016, *MNRAS*, 463, 1797
- Eggemeier A., Scoccimarro R., Smith R. E., Crocce M., Pezzotta A., Sánchez A. G., 2021, *Phys. Rev. D*, 103, 123550
- Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
- Foreman S., Coulton W., Villaescusa-Navarro F., Barreira A., 2020, *MNRAS*, 498, 2887
- Fu L. et al., 2014, *MNRAS*, 441, 2725
- Gaztanaga E., 1994, *MNRAS*, 268, 913
- Gaztañaga E., Norberg P., Baugh C. M., Croton D. J., 2005, *MNRAS*, 364, 620
- Gil-Marín H., Noreña J., Verde L., Percival W. J., Wagner C., Manera M., Schneider D. P., 2015a, *MNRAS*, 451, 539
- Gil-Marín H. et al., 2015b, *MNRAS*, 452, 1914
- Gil-Marín H., Percival W. J., Verde L., Brownstein J. R., Chuang C.-H., Kitaura F.-S., Rodríguez-Torres S. A., Olmstead M. D., 2017, *MNRAS*, 465, 1757
- Ginzburg D., Desjacques V., Chan K. C., 2017, *Phys. Rev. D*, 96, 083528
- Gualdi D., Verde L., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 041
- Gualdi D., Gil-Marín H., Schuhmann R. L., Manera M., Joachimi B., Lahav O., 2019a, *MNRAS*, 484, 3713
- Gualdi D., Gil-Marín H., Manera M., Joachimi B., Lahav O., 2019b, *MNRAS*, 484, L29
- Hahn C., Villaescusa-Navarro F., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 029

Hahn C., Villaescusa-Navarro F., Castorina E., Scoccimarro R., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 040

Hand N., Feng Y., Beutler F., Li Y., Modi C., Seljak U., Slepian Z., 2018, *AJ*, 156, 160

Heinrich C., Doré O., 2020, *Phys. Rev. D*, 102, 123549

Hinshaw G. et al., 2013, *ApJS*, 208, 19

Howlett C., Lewis A., Hall A., Challinor A., 2012, *J. Cosmol. Astropart. Phys.*, 2012, 027

Kulkarni G. V., Nichol R. C., Sheth R. K., Seo H.-J., Eisenstein D. J., Gray A., 2007, *MNRAS*, 378, 1196

Lazanu A., Giannantonio T., Schmittfull M., Shellard E. P. S., 2016, *Phys. Rev. D*, 93, 083517

Lesgourgues J., Pastor S., 2006, *Phys. Rep.*, 429, 307

Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473

McBride C. K., Connolly A. J., Gardner J. P., Scranton R., Newman J. A., Scoccimarro R., Zehavi I., Schneider D. P., 2011, *ApJ*, 726, 13

McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *MNRAS*, 465, 2936

McCarthy I. G., Bird S., Schaye J., Harnois-Deraps J., Font A. S., van Waerbeke L., 2018, *MNRAS*, 476, 2999

Matarrese S., Verde L., Heavens A. F., 1997, *MNRAS*, 290, 651

Moradinezhad Dizgah A., Biagetti M., Sefusatti E., Desjacques V., Noreña J., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 015

Mummery B. O., McCarthy I. G., Bird S., Schaye J., 2017, *MNRAS*, 471, 227

Oddo A., Sefusatti E., Porciani C., Monaco P., Sánchez A. G., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 056

Pan J., Szapudi I., 2005, *MNRAS*, 362, 1363

Pearson D. W., Samushia L., 2018, *MNRAS*, 478, 4500

Peebles P. J. E., Groth E. J., 1975, *ApJ*, 196, 1

Pfeifer S., McCarthy I. G., Stafford S. G., Brown S. T., Font A. S., Kwan J., Salcido J., Schaye J., 2020, *MNRAS*, 498, 1576

Pietroni M., 2008, *J. Cosmol. Astropart. Phys.*, 2008, 036

Planck Collaboration VI, 2020, *A&A*, 641, A6

Sánchez A. G., Baugh C. M., Percival W. J., Peacock J. A., Padilla N. D., Cole S., Frenk C. S., Norberg P., 2006, *MNRAS*, 366, 189

Schaller M. et al., 2015, *MNRAS*, 451, 1247

Schaye J. et al., 2010, *MNRAS*, 402, 1536

Schneider A., Teyssier R., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 049

Scoccimarro R., 2000, *ApJ*, 544, 597

Scoccimarro R., Feldman H. A., Fry J. N., Frieman J. A., 2001, *ApJ*, 546, 652

Sefusatti E., Crocce M., Scoccimarro R., Couchman H. M. P., 2016, *MNRAS*, 460, 3624

Seljak U., Hamaus N., Desjacques V., 2009, *Phys. Rev. Lett.*, 103, 091303

Semboloni E., Hoekstra H., Schaye J., van Daalen M. P., McCarthy I. G., 2011, *MNRAS*, 417, 2020

Semboloni E., Hoekstra H., Schaye J., 2013, *MNRAS*, 434, 148

Slepian Z. et al., 2017a, *MNRAS*, 468, 1070

Slepian Z. et al., 2017b, *MNRAS*, 469, 1738

Song Y.-S., Taruya A., Oka A., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 007

Springel V., 2005, *MNRAS*, 364, 1105

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726

Stafford S. G., McCarthy I. G., Crain R. A., Salcido J., Schaye J., Font A. S., Kwan J., Pfeifer S., 2020, *MNRAS*, 493, 676

Steele T., Baldauf T., 2021, *Phys. Rev. D*, 103, 023520

Taruya A., Bernardeau F., Nishimichi T., Codis S., 2012, *Phys. Rev. D*, 86, 103528

Tegmark M. et al., 2004, *Phys. Rev. D*, 69, 103501

Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlöber S., 2010, *ApJ*, 724, 878

Tomlinson J., Jeong D., Kim J., 2019, *AJ*, 158, 116

van Daalen M. P., Schaye J., Booth C. M., Dalla Vecchia C., 2011, *MNRAS*, 415, 3649

van Daalen M. P., Schaye J., McCarthy I. G., Booth C. M., Dalla Vecchia C., 2014, *MNRAS*, 440, 2997

van Daalen M. P., McCarthy I. G., Schaye J., 2020, *MNRAS*, 491, 2424

Velliscig M., van Daalen M. P., Schaye J., McCarthy I. G., Cacciato M., Le Brun A. M. C., Dalla Vecchia C., 2014, *MNRAS*, 442, 2641

Verde L., Treu T., Riess A. G., 2019, *Nat. Astron.*, 3, 891

Verde L. et al., 2002, *MNRAS*, 335, 432

Veropalumbo A. et al., 2021, *MNRAS*, 507, 1184

Yankelevich V., Porciani C., 2019, *MNRAS*, 483, 2078

APPENDIX A: HALO MASS-WEIGHTED SPECTRA

In the main analysis, when computing the halo power spectra and bispectra we treat all haloes equally (same weighting). Here, we explore an alternative weighting scheme, which is to weight each halo by its mass. It was shown previously by Seljak, Hamaus & Desjacques (2009) that the effects of shot noise can be significantly reduced in the halo power spectrum by adopting such a weighting scheme. In BSKIT there is an option to include particle weight into analysis when computing the grids. We simply substitute in the halo mass. For the matched case we used DMO halo masses for both DMO and hydro grids.

We demonstrate our findings in Fig. A1. While we see no obvious advantage in terms of the shot noise for the bispectrum, it is interesting to note that the effects of baryon physics are reduced significantly when weighting haloes by their mass instead of equally weighting each halo. This is likely because higher mass haloes are less significantly affected by ejective feedback. Thus, if the aim is

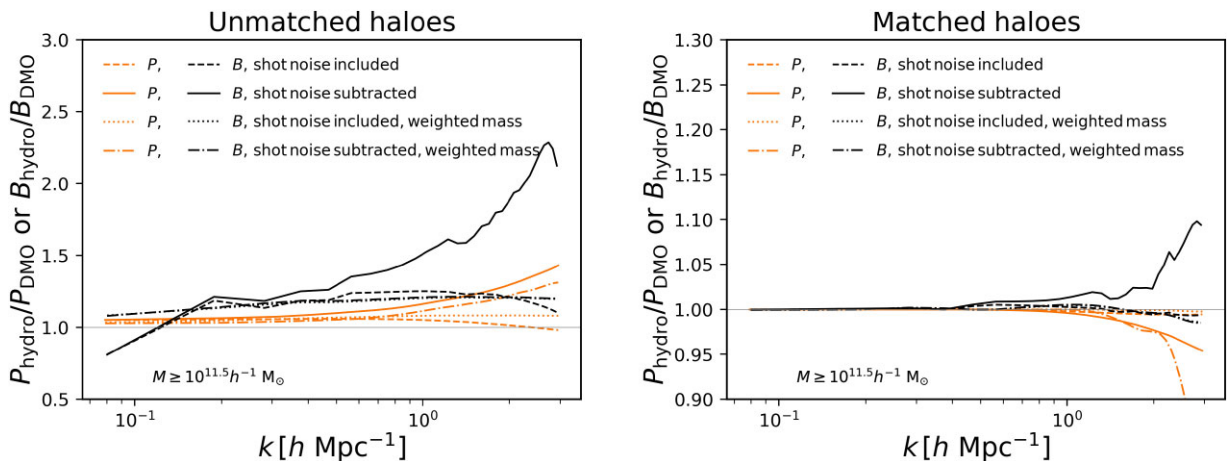


Figure A1. The same as in Fig. 4 but including cases where the haloes are weighted by their masses when computing the power spectra and bispectra.

to mitigate against the effects of baryons, employing a halo mass weighting scheme goes some way towards achieving this objective.

APPENDIX B: TRIANGULAR CONFIGURATION DEPENDENCE AS A FUNCTION OF SCALE

Here, we extend the analysis presented in Section 3.3 for different values of k_1 . In Section 3.3, we adopted $k_1 = 1.54 h \text{ Mpc}^{-1}$, which is approximately half way between between the minimum and maximum scales we can sample in the simulation box. Here, we

study the dependence of the bispectrum ratio for different triangular configurations on large scales ($k_1 = 1.01 h \text{ Mpc}^{-1}$; see Fig. B1) and for small scales ($k_1 = 2.92 h \text{ Mpc}^{-1}$; see Fig. B2). Again the values of k_1 are chosen in a way that the equilateral configuration $k_1 = k_2 = k_3$ must exist. For small values of k_1 , the number of possible triangles in the box is low. We therefore plot in Fig. B1 the first wavevector value where the number of triangles is representative.

Overall, we see similar trends to that presented in Fig. 1, indicating that the dependence on the adopted triangular configuration is similar for different choices of k_1 .

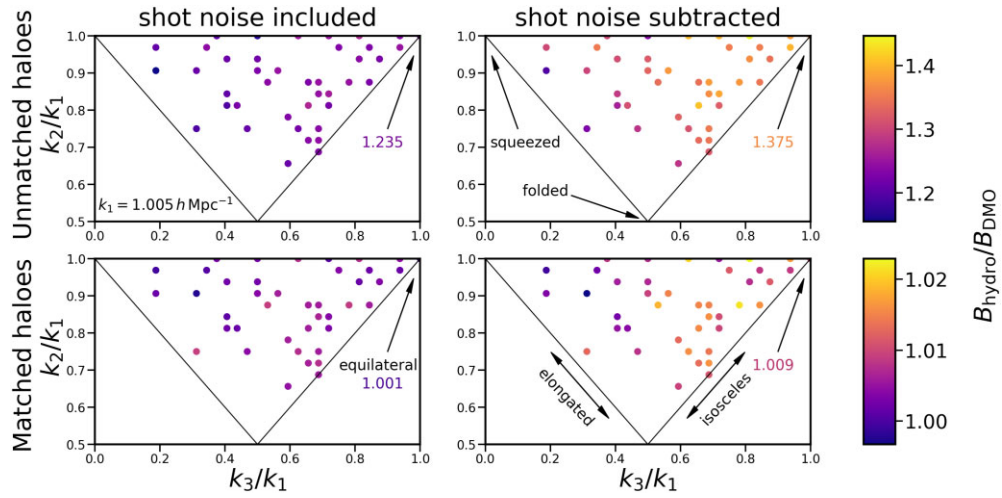


Figure B1. The same as in Fig. 1 but for a lower value of $k_1 = 1.01 h \text{ Mpc}^{-1}$.




Figure B2. The same as in Figs 1 and B1 but for a higher value of $k_1 = 2.92 h \text{ Mpc}^{-1}$.