






SP(k) – a hydrodynamical simulation-based model for the impact of baryon physics on the non-linear matter power spectrum

Jaime Salcido ,  Ian G. McCarthy ,  Juliana Kwan, Amol Upadhye and Andreea S. Font 

Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK

Accepted 2023 May 12. Received 2023 April 12; in original form 2022 September 25

ABSTRACT

Upcoming large-scale structure surveys will measure the matter power spectrum to approximately per cent level accuracy with the aim of searching for evidence for new physics beyond the standard model of cosmology. In order to avoid biasing our conclusions, the theoretical predictions need to be at least as accurate as the measurements for a given choice of cosmological parameters. However, recent theoretical work has shown that complex physical processes associated with galaxy formation (particularly energetic feedback processes associated with stars and especially supermassive black holes) can alter the predictions by many times larger than the required accuracy. Here we present SP(k), a model for the effects of baryon physics on the non-linear matter power spectrum based on a new large suite of hydrodynamical simulations. Specifically, the ANTILLES suite consists of 400 simulations spanning a very wide range of the ‘feedback landscape’ and show that the effects of baryons on the matter power spectrum can be understood at approaching the per cent level in terms of the mean baryon fraction of haloes, at scales of up to $k \lesssim 10 h \text{ Mpc}^{-1}$ and redshifts up to $z = 3$. For the range of scales and redshifts that will be probed by forthcoming cosmic shear measurements, most of the effects are driven by galaxy group mass haloes ($M \sim 10^{13-14} M_{\odot}$). We present a simple PYTHON implementation of our model, available at <https://github.com/jemme07/pyspk>, which can be used to incorporate baryon effects in standard gravity-only predictions, allowing for marginalization over baryon physics within cosmological pipelines.

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

1 INTRODUCTION

Measurements of the growth of large-scale structure (LSS) provide an important test of our cosmological theoretical framework (Bond, Efstathiou & Silk 1980; Peebles 1980; Davis et al. 1985; Kaiser 1987; Peacock & Dodds 1994). They are independent of, and complementary to, constraints from analyses of fluctuations in the cosmic microwave background (CMB) and geometric probes, such as Type Ia supernovae (SNe) and baryon acoustic oscillations. The different LSS tests (e.g. galaxy clustering, Sunyaev–Zel’dovich power spectrum, cosmic shear, CMB lensing, redshift-space distortions, etc.) characterize the clustering of matter on different scales and its evolution over cosmic time. On the largest scales where the fluctuations are small and gravity is the sole force of relevance, perturbation theory is sufficiently accurate to calculate the clustering of matter. However, at present most LSS tests probe well into the non-linear regime, since that is typically where most of the observational signal and much of the cosmological sensitivity originates (e.g. Amon & Efstathiou 2022). The standard approach is therefore to adopt the (semi-)analytical ‘halo model’ formalism (Ma & Fry 2000; Peacock & Smith 2000; Seljak 2000; Cooray & Sheth 2002; Smith et al. 2003; Mead et al. 2016; Acuto et al. 2021), which is often

calibrated using aspects of large N -body cosmological simulations, or to use such N -body simulations to directly correct linear theory in an empirical fashion (e.g. Takahashi et al. 2012; Heitmann et al. 2016; Lawrence et al. 2017; DeRose et al. 2019; Euclid Collaboration 2019; Angulo et al. 2021).

These approaches would be fully adequate if the matter in the Universe were composed entirely of dark matter. However, baryons contribute a non-negligible fraction of the matter density ($\Omega_b/\Omega_m \approx 0.157 \pm 0.001$; Planck Collaboration VI 2020) and work based on cosmological *hydrodynamical* simulations has shown that feedback processes associated with galaxy formation can have a relatively large effect (typically many times larger than the statistical precision of upcoming surveys) on the matter distribution on scales of up to a few tens of megaparsecs (van Daalen et al. 2011; Schneider & Teyssier 2015; Mummery et al. 2017; Springel et al. 2018; van Daalen, McCarthy & Schaye 2020). It is therefore crucially important that we model these effects as accurately as possible, as they will introduce significant biases in the inferred cosmological parameters from upcoming surveys if no action is taken (e.g. Semboloni et al. 2011; Schneider et al. 2020; Castro et al. 2021; Debackere, Schaye & Hoekstra 2021).

One way to tackle this challenging problem is to modify the gravity-only predictions with simple analytical prescriptions for baryon physics that have some number of associated free parameters and to jointly constrain the cosmological and feedback parameters through comparisons to LSS observables. Examples of this approach

* E-mail: J.SalcidoNegrete@ljmu.ac.uk (JS); I.G.McCarthy@ljmu.ac.uk (IGM)

include the halo model HMCODE of Mead et al. (2016) (see also Mead et al. 2021) or the ‘baryonification’ method of Schneider & Teyssier (2015), Aricò et al. (2021) and Giri & Schneider (2021), which directly modifies the outputs of gravity-only simulations by adjusting the radial position of particles within haloes in order to mimic the effects of baryons on small scales. Some of the strengths of these approaches include (i) they are considerably computationally cheaper than running full cosmological hydrodynamical simulations; (ii) the baryon prescriptions are generally flexible and easy to adjust; and (iii) they can be straightforwardly incorporated within existing pipelines based on gravity-only simulations or the halo model.

Recent cosmic shear surveys have employed these methods in an attempt to account for and measure, baryonic effects. For example, the fiducial Kilo Degree Survey (KiDS) analysis employs a halo model prescription where the effects of baryons are incorporated by marginalizing over a phenomenological ‘bloating’ parameter that modifies the concentrations of dark matter haloes (e.g. Asgari et al. 2021; Heymans et al. 2021; Tröster et al. 2021). The Dark Energy Survey (DES) has taken a different approach, by introducing scale cuts to remove small-scale measurements that are most affected by baryonic effects (e.g. DeRose et al. 2019; Krause et al. 2021; Amon et al. 2022; Secco et al. 2022). This approach is motivated by the concern that errors in the modelling of baryonic effects could lead to biased cosmological parameter estimates. However, introducing scale cuts comes at the expense of losing valuable information on small scales, which can be important for constraining cosmological parameters and testing extensions to the standard model. More recently, the full DES data (including small-scale measurements) have been reanalysed using the baryonification approach to account for the impact of baryon physics (Aricò et al. 2023; Chen et al. 2023). Both the reanalysis of the DES data and a recent joint analysis of KiDS cosmic shear and Sunyaev–Zel’dovich effect data (Tröster et al. 2022) using a more physical halo model (Mead et al. 2020) detect the impact of baryon physics even in current data at the approximately $2\text{--}3\sigma$ level. These studies conclude that the implied level of suppression of the power spectrum is consistent with the predictions of recent calibrated hydrodynamical simulations such as BAHAMAS (see below).

While the halo model and baryonification approaches have some important strengths, there are also some disadvantages to these methods. Particularly, the modelling of baryon physics and its back reaction on dark matter is simplistic and generally not self-consistent and that there may be non-negligible degeneracies between the various baryonic ‘nuisance’ parameters themselves and also between those parameters and the cosmological parameters being varied. Marginalization over the uncertain baryon physics may therefore lead to a significant degradation of the cosmological constraining power. Of course, cosmological hydrodynamical simulations also have adjustable free parameters (as discussed immediately below), but the resulting diversity of outcomes in terms of the matter clustering is likely to be more constrained in hydrodynamical simulations than simple empirical models would allow. For example, the baryon fraction–halo mass relation cannot be arbitrarily steep in simulations because high-mass objects are assembled from the accretion of lower mass objects (e.g. Balogh et al. 2008) and haloes can recapture gas as they grow, too. Also, naturally emerging conditions such as convective and virial equilibrium place constraints on the radial distribution of matter within haloes, and so on. The upshot is that we expect many of the parameters that characterize mass distributions of groups and clusters to be physically correlated rather than independent of each other.

Full cosmological hydrodynamical simulations could therefore be employed as a means of self-consistently incorporating the impact of baryons on LSS when constraining cosmology. However, this has so far not proved possible due to their computational expense, noting that a typical MCMC chain in cosmology can require $\sim 10^5$ evaluations whereas currently available suites of simulations typically only contain a handful of realizations. Furthermore, a major obstacle in directly simulating the impact of baryon physics on the matter clustering is that current simulations do not resolve all of the physical scales necessary to capture such processes in an *ab initio* way. Consequently, so-called subgrid models are required to include these effects and they often have considerable uncertainties, which are not unlike the uncertainties in simple baryon models discussed above (although the impact of subgrid models is often bounded by physical constraints, whereas the empirical models may not be, as discussed above). Indeed, previous simulation work has shown that variations of the parameters associated with the efficiencies of feedback processes even within plausible bounds can lead to relatively large differences in the predicted properties of galaxy groups and clusters (e.g. Puchwein, Sijacki & Springel 2008; McCarthy et al. 2010; Le Brun et al. 2014; Planelles et al. 2014; Oppenheimer et al. 2021) which dominate the matter clustering (van Daalen & Schaye 2015; Mead et al. 2020). A consequence of these variations in the predicted properties of groups/clusters is relatively large study-to-study variations in the predicted impact of baryons on the matter power spectrum (see the simulation comparisons in Chisari et al. 2019; van Daalen et al. 2020), in spite of the simulations being more constrained than phenomenological models.

As discussed recently in Oppenheimer et al. (2021), the variations in the predicted properties of groups/clusters and the impact of baryons on the matter power spectrum from hydrodynamical simulations in the literature is not unexpected. It is a consequence of not being able to derive the efficiencies for the relevant feedback processes from first principles (see discussion in Schaye et al. 2015). The efficiency of feedback in simulations must therefore generally be calibrated in order to ensure they reproduce particular observed quantities, after which the realism of the simulations may be tested against independent quantities. The approach of the BAHAMAS programme (McCarthy et al. 2017, 2018, see also the recent FABLE simulations; Henden et al. 2018; Henden, Puchwein & Sijacki 2020), was to explicitly calibrate the feedback efficiencies so that they reproduce the observed baryon fractions of galaxy groups. Aside from an explicit dependence on the universal baryon fraction (White et al. 1993), Ω_b/Ω_m , which is tightly constrained by the CMB, the baryon fractions of groups should be insensitive to changes in cosmology and therefore represents a fairly ideal quantity on which to calibrate the feedback. Note also that since the growth of fluctuations is fundamentally a gravitational process, by ensuring the simulations have the correct baryon fractions on the scale of groups/clusters, the impact of baryons on $P(k)$ ought to be strongly constrained by this approach. van Daalen et al. (2020) have recently confirmed this simple picture by demonstrating that the differences in the predicted impact of baryons on the present-day $P(k)$ from different simulations can be understood at approximately the per cent level up to $k \approx 1 \ h \text{ Mpc}^{-1}$ in terms of the differences in baryon fraction in the various simulations at a mass scale of $\sim 10^{14} \text{ M}_\odot$. We highlight here that the simulations analysed in that study varied by more than a factor of a thousand in mass resolution, used different hydro solvers and subgrid physics implementations, assumed different baseline cosmologies, and some (specifically BAHAMAS) varied the initial conditions to explore

the potential impact of cosmic variance. None of these variations were found to significantly affect the impact of baryons on $P(k)$ after differences due to the group baryon fractions were factored out.

The results of van Daalen et al. (2020) are very promising and potentially offer a path forward for incorporating the impact of baryons on the matter power spectrum from cosmological hydrodynamical simulations. Note that with a mapping between the baryon fraction and the impact of baryons $P(k)$, there is no longer a necessity to calibrate the hydrodynamical simulations to high precision to match some particular data set. Instead, one can properly account for the uncertainties in the feedback/subgrid effects on $P(k)$ by, for example, conservatively marginalizing over the uncertainties in the observed baryon fractions. Furthermore, with a simple mapping, the correction to the gravity-only clustering could be straightforwardly included in cosmological pipelines because it can be rapidly evaluated. However, before these aims can be achieved a number of limitations must first be overcome. First and foremost, the quantitative link between halo baryon fraction and the suppression of the power spectrum must be established for a wider range of models. van Daalen et al. (2020) used a small number (≈ 10) of publicly available simulations which likely do not bracket the full range of possible behaviours. Furthermore, the mapping between baryon fraction and $P(k)$ was established only at $z = 0$ and up to a maximum wavenumber of $k \approx 1 \, h \, \text{Mpc}^{-1}$, both of which are insufficient for current and future LSS tests.

In this study, we overcome these limitations by presenting a new large suite of cosmological hydrodynamical simulations designed specifically to build a mathematical model for the suppression of the matter power spectrum with the baryon fractions of galaxy groups as its input. The ANTILLES suite contains 400 cosmological hydrodynamical simulations that vary both the important parameters that characterize the efficiencies of stellar and active galactic nucleus (AGN) feedback in the simulations as well as the employed hydrodynamics scheme, and span an unprecedentedly wide range of behaviours in terms of baryon fractions and $P(k)$ modifications. We develop a model that can reproduce the effects of baryons on $P(k)$ to typically better than ≈ 2 per cent precision up to $k = 10 \, h \, \text{Mpc}^{-1}$ out to a redshift of $z = 3$.

This study is structured as follows. In Section 2, we describe the new suite of cosmological hydrodynamical simulations. In Section 3, we present $SP(k)$, an empirical model that provides the mapping between the observable baryon fractions of groups/clusters and the suppression of the matter power spectrum, $P(k)$. In Section 4, we test $SP(k)$ against the independent BAHAMAS simulations. In Section 5, we discuss some limitations of this work and in Section 6 we summarize our findings.

2 SIMULATIONS

To quantify the potential impact of baryon physics on the non-linear matter power spectrum, our simulations must satisfy a number of requirements. First, the simulated volumes must be sufficiently large to contain a representative sample of the haloes that contribute most significantly to the matter power spectrum. They must also be of sufficiently high resolution to resolve the range of scales over which cosmological measurements are made and so that we resolve the locations of important feedback processes (specifically, the simulations contain the haloes from which the majority of the energetic feedback originates). Finally, we need to explore a wide range of feedback possibilities, which we dub the ‘feedback landscape’, such that key properties such as the baryon fractions

Table 1. The cosmological parameters for the simulations used in this study. We adopt a flat Λ CDM cosmology with WMAP 9-year-based cosmological parameters (Hinshaw et al. 2013). Ω_m , Ω_Λ , Ω_b are the average densities of matter, dark energy, and baryonic matter in units of the critical density at redshift $z = 0$; H_0 is the Hubble constant, σ_8 is the square root of the linear variance of the matter distribution when smoothed with a top-hat filter of radius $8 \, h^{-1} \, \text{cMpc}$, and n_s is the scalar power-law index of the power spectrum of primordial adiabatic perturbations.

Cosmological parameter	Value
Ω_m	0.2793
Ω_Λ	0.7207
Ω_b	0.0463
$h \equiv H_0/(100 \, \text{km s}^{-1} \, \text{Mpc}^{-1})$	0.7
σ_8	0.821
n_s	0.972

(stellar and gas fractions) span a wide range and conservatively bracket current and hopefully future observational constraints.

With these requirements in mind, we have created a new bespoke suite of 400 simulations, the ANTILLES suite. Each simulation has a box size of $100 \, \text{Mpc} \, h^{-1}$ on a side, which van Daalen et al. (2020) have shown is sufficiently large to capture the relative effects of baryons on the matter power spectrum (see their appendix A). Our simulations adopt the same mass resolution as BAHAMAS, which McCarthy et al. (2017, 2018) have shown is sufficiently high for various LSS tests (see also appendix A of van Daalen et al. 2020). Thus, each simulation has 256^3 baryon and dark matter particles (each), corresponding to (initial) masses¹ of $m_g = 1.09 \times 10^9 \, \text{M}_\odot$ and $m_{\text{dm}} = 5.51 \times 10^9 \, \text{M}_\odot$, respectively, given our choice of cosmology (below). The gravitational softening is fixed to $4 \, h^{-1} \, \text{kpc}$ in physical coordinates below $z = 3$ and in comoving coordinates at higher redshifts.

We adopt a flat Λ CDM cosmology consistent with the WMAP 9-year results (Hinshaw et al. 2013). The cosmological parameters used in these simulations are listed in Table 1. As we focus on the relative impact of baryons on the matter power spectrum, we do not expect the precise choice of cosmology to be important for our purposes. We nevertheless discuss the possible dependence of our results on cosmology in Section 5.

The Boltzmann code CAMB² (Lewis, Challinor & Lasenby 2000) was used to compute the transfer functions which were supplied to a modified version of the N-GENIC³ code to include second-order Lagrangian Perturbation Theory to create the initial conditions at a starting redshift of $z = 127$. When producing initial conditions for hydrodynamical simulations, we use the separate transfer functions computed by CAMB for each individual component (i.e. baryons and dark matter). Similarly, in order to avoid any offset in the amplitude of the matter power spectrum of the hydro simulations with respect to their dark matter-only counterpart at large scales, for the dark matter-only version of the simulations, we generated two separate fluids, one with the dark matter transfer function and the other with the baryon transfer function (Valkenburg & Villaescusa-Navarro 2017; van Daalen et al. 2020). Additionally, the same random phases were used to generate each set of initial conditions. Hence, comparisons

¹Unless otherwise stated, throughout the paper masses are specified in M_\odot (not $h^{-1} \, \text{M}_\odot$).

²<https://camb.info/>

³<https://github.com/sbird/S-GenIC>

made between the different simulations are not subject to cosmic variance complications.

The simulation suite was run with a modified version of the GADGET-3 smoothed particle hydrodynamics (SPH) code (last described by Springel 2005) and includes a full treatment of gravity and hydrodynamics. Specifically, we use version of GADGET-3 that was modified for the EAGLE project (Schaye et al. 2015).

In order to examine the potential effects of different hydrodynamic schemes, the ANTILLES suite comprises a set of simulations that use the standard GADGET flavour of SPH, as well as set that uses the more recent state-of-the-art ANARCHY formulation. The improvements within ANARCHY include the use of the pressure–entropy formulation of SPH derived by Hopkins (2013), the artificial viscosity switch from Cullen & Dehnen (2010), an artificial conduction switch similar to that of Price (2008), the C_2 kernel of Wendland (1995), and the time-step limiters of Durier & Dalla Vecchia (2012).

Important non-gravitational processes, such as stellar and AGN feedback, that are not resolved by the simulations are implemented as subgrid physical models. A full description of these subgrid models can be found in Schaye et al. (2015). In summary:

(i) Radiative cooling and photoheating are implemented element by element as in Wiersma, Schaye & Smith (2009a), including the 11 elements found to be important, namely, H, He, C, N, O, Ne, Mg, Si, S, Ca, and Fe. Hydrogen reionization is implemented by switching on the full Haardt & Madau (2001) background at redshift $z = 7.5$.

(ii) Star formation is implemented stochastically following the pressure-dependent Kennicutt–Schmidt relation as in Schaye & Dalla Vecchia (2008). Above a density threshold $n_{\text{H}}^* = 0.1 \text{ cm}^{-3}$, which is designed to track the transition from a warm atomic to an unresolved cold molecular gas phase (Schaye 2004), gas particles have a probability of forming stars determined by their pressure. We use a single fixed density threshold for these simulations, as opposed to a metallicity-dependent threshold in EAGLE.

(iii) Time-dependent stellar mass loss due to winds from massive stars and asymptotic giant branch stars, core collapse SNe and Type Ia SNe, is tracked following Wiersma et al. (2009b).

(iv) Stellar feedback is implemented using the kinetic wind model of Dalla Vecchia & Schaye (2008), which differs from the thermal implementation used in EAGLE. This is motivated by resolution considerations, particularly that relative high resolution is required for efficient thermal feedback due to its stochastic implementation. We instead adopt a kinetic implementation, as also adopted in BAHAMAS.

(v) Seed black holes (BHs) of mass $M = 1 \times 10^6 M_{\odot}$ are placed in haloes with a mass greater than $2.75 \times 10^{11} M_{\odot}$ (corresponding to ≈ 50 DM particles) and tracked following the methodology of Springel, Di Matteo & Hernquist (2005) and Booth & Schaye (2009). Once seeded, BHs can grow via Eddington-limited gas accretion, at a rate which is proportional to the Bondi–Hoyle–Lyttleton rate, as well as through mergers with other BHs following Booth & Schaye (2009).

(vi) Feedback from AGNs is implemented following the stochastic heating scheme. A fraction of the accreted gas on to the BH is released as thermal energy with a fixed heating temperature into the surrounding gas following Booth & Schaye (2009).

In order to explore a wide feedback landscape that brackets current observational constraints (with their associated uncertainties) on the stellar and gas fractions, we systematically vary the main subgrid parameters governing the efficiencies of stellar and AGN feedback. In particular, the galaxy (star) formation efficiency is quite sensitive to variations of the wind velocity (v_w) and mass-loading (η_w) parameters.

Table 2. Subgrid parameters varied in this study. v_w and η_w are the wind velocity and mass-loading factor used in the kinetic wind stellar feedback model as per Dalla Vecchia & Schaye (2008). ΔT_{heat} is the temperature increase of gas particles during BH feedback events, and n_{heat} is the number of neighbouring gas particles to be heated. $n_{\text{H,BH}}^*$ is the density threshold above which the BH accretion rate is boosted in the Booth & Schaye (2009) accretion model due to the lack of resolution of the cold gas phase at high densities.

Parameter	Range
v_w (km s^{-1})	[50, 350]
η_w	[1, 10]
$\log_{10}(\Delta T_{\text{heat}} [\text{K}])$	[7, 8.5]
n_{heat}	[1, 30]
$\log_{10}(n_{\text{H,BH}}^* [\text{cm}^{-3}])$	[−3, −1]

These parameters have a significant effect on the star formation histories of galaxies and the shape of the galaxy stellar mass function (GSMF), especially at the low-mass end where stellar feedback is expected to dominate (e.g. Schaye et al. 2010; McCarthy et al. 2017). For AGN feedback, the two main parameters that determine how frequent and energetic the feedback events are, are the number of neighbouring gas particles to heat (n_{heat}) and the temperature increase of such particles (ΔT_{heat}). In particular, changes in ΔT_{heat} have a significant impact on the gas mass fraction of groups and clusters of galaxies (McCarthy et al. 2017), which can be understood by recognizing that the likelihood of significant gas ejection will be favourable if $\Delta T_{\text{heat}} \gg T_{\text{vir}}$. Finally, in the Booth & Schaye (2009) accretion model implemented in the simulations, the ‘boost factor’, α , relative to pure Bondi–Hoyle accretion is a power-law function of the local density for gas above a pivot point, $n_{\text{H,BH}}^*$, which is set to the star formation threshold $n_{\text{H}}^* = 0.1 \text{ cm}^{-3}$, as the simulations resolve lower densities, where no cold, molecular gas phase, is expected. At higher densities, the simulations lack the cold gas phase, which would boost the Bondi–Hoyle rate due to its sound speed dependence (Schaye 2004; Booth & Schaye 2009). We set the power-law exponent to $\beta = 2$ for high densities, and the boost factor is set to go to unity at low density, where the simulations are well resolved. Nevertheless, for gas at 10^4 K , at the particle resolution of the simulations, we only resolve the Jeans length for densities $< 10^{-3} \text{ cm}^{-3}$. Hence, we explore boosting the accretion rate at lower densities to compensate for the lack of resolution. We find that changing the pivot point, $n_{\text{H,BH}}^*$, can have a significant impact at the mass-scale at which AGN feedback becomes efficient, effectively changing the shape of the knee of the GSMF. The five parameters varied in this study, with their respective range of values, are provided in Table 2.

In order to ensure a full parameter space coverage, we employed a Latin hypercube sampling with multidimensional uniformity developed by Deutsch & Deutsch (2012). We used 200 nodes to uniformly sample the parameter space. Finally, we ‘mirrored’ these 200 simulations using the exact same subgrid physics implementation and sampled parameters, but using a standard GADGET flavour of SPH, as opposed to the more recent state-of-the-art ANARCHY formulation. This gives a total of 400 simulations in the ANTILLES sample.

The GSMF, median gas fraction, and median total baryon fraction as a function of halo mass spanned by the simulations at redshift $z = 0.125$ is shown in Fig. 1. The figure shows how our parameter space exploration easily brackets the current observational constraints. Note that when comparing the simulations to the observations in the middle panel (gas mass fractions), one should compare the solid black curves, which represent the median relations of each

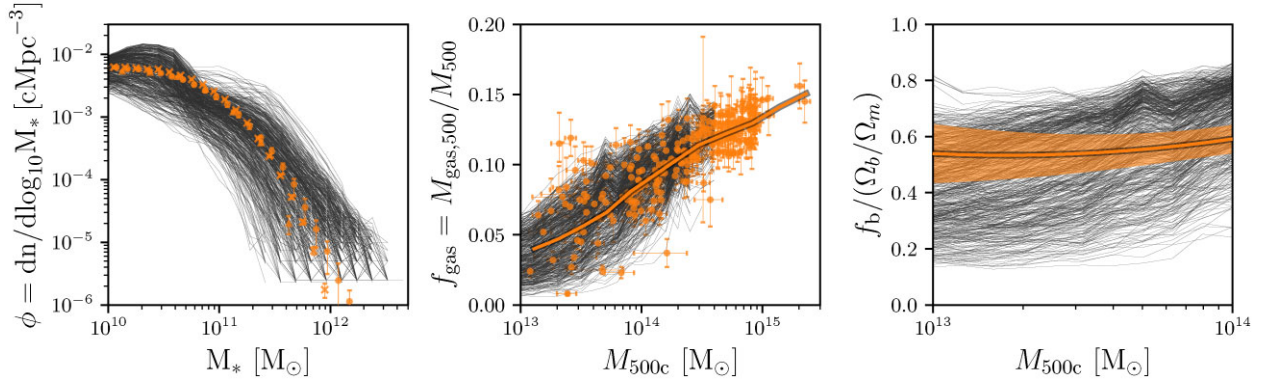


Figure 1. Galaxy stellar mass function (left), median gas fraction (middle), and median total baryon fraction (right) as a function of halo mass for all 400 simulations at redshift $z = 0.125$. Observational data with their associated uncertainties are shown in orange for the stellar mass function (Baldry et al. 2012; Bernardi et al. 2013; Moustakas et al. 2013; Driver et al. 2022), for the gas fractions of groups and clusters (Vikhlinin et al. 2006; Maughan et al. 2008; Pratt et al. 2009; Rasmussen & Ponman 2009; Sun et al. 2009; Lin et al. 2012; Gonzalez et al. 2013; Sanderson et al. 2013; Lovisari, Reiprich & Schellenberger 2015; Pearson et al. 2017), the latter of which are derived from resolved X-ray observations, and the fit to the median baryon fraction from the latest HSC-XXL weak gravitational lensing data from Akino et al. (2022), including a correction for the contribution of blue galaxies and the diffuse intracluster light. The light shaded region encloses the 1σ uncertainty. The orange line in the middle panel highlights the median $f_{\text{gas}}-M_{500c}$ relation from observations, including a correction from the inferred scaling relations from the X-ray selected sample of galaxy groups and clusters in Sereno et al. (2020). Comparing the solid black curves to the solid orange curve (i.e. median relations) in the middle panel, the simulations conservatively bracket the observed gas fractions, the observed galaxy stellar mass function (left-hand panel), as well as the median total baryon fraction (right-hand panel).

of the 400 simulations, to the solid orange curve, which represents the median relation from resolved X-ray observations. (The orange data points represent individual observed groups and clusters, we do not show individual simulated clusters for clarity, but we note the intrinsic scatter in the simulated relations is similar to the observed intrinsic scatter.) Thus, the simulations conservatively bracket both the observed gas fractions and the observed stellar mass function. If one adopts the statistical uncertainties on the median baryon mass fractions⁴ from the recent study of Akino et al. (2022), then our simulations span a range of approximately -7σ to $+6\sigma$ with respect to the observed baryon fraction at a mass scale of $\sim 10^{14} M_{\odot}$.

Note that our desire to span a much wider range of baryon fractions than is apparently allowed by current observations is motivated by two factors: (i) the low-redshift data we compare to may have non-negligible biases and (ii) the baryon fractions of higher redshift (e.g. $z \approx 0.5-1$) groups and clusters, which give rise to much of the lensing signal, is not well constrained at present by observations, thus we want a range of simulated behaviours that is wide enough to hopefully encapsulate future measurements of high-redshift systems as well. With regards to possible biases in current data, observational measurements are always subject to both random and systematic errors. For instance, the observationally derived stellar masses are subject to systematic errors originating from, e.g. stellar population modelling, spectral energy distribution fitting, surface brightness profile fitting, and corrections for dust extinction. Gas fraction measurements from X-ray data of groups and clusters are subject to uncertainties in, e.g. deviations from spherical

symmetry and hydrostatic equilibrium, and modelling of selection effects, particularly in the group regime.

Finally, we note that the prior ranges on each of the subgrid parameters were essentially selected by examining the baryon fraction results from a small number of test simulations, so it is no surprise that the resulting hypercube spans a large range of baryon fractions.

Note also that while we use the combined set of 400 simulations to construct our model for the suppression of the matter power spectrum (below), we have also explored analysing the two SPH suites separately. While for a given set of subgrid parameters the choice of SPH flavour can affect the resulting baryon fractions, the reaction in terms of $P(k)$ is independent of the choice of SPH flavour at a given baryon fraction. This is consistent with the findings of van Daalen et al. (2020) and allows us to construct a single (more accurate) model in terms of baryon fractions using the entire set of 400 ANTILLES simulations.

3 MODELLING BARYON PHYSICS EFFECTS ON $P(k)$

We begin by analysing the relative (fractional) impact of baryon physics on the total matter power spectrum for the 400 simulations. Fig. 2 shows the ratio of power spectra⁵ from the different simulations with respect to a DM-only counterpart at redshift $z \approx 0.1$. The wide range of variations in the feedback models in our simulations gives rise to a large diversity of impacts on $P(k)$, in agreement with previous studies that show that the effect of baryons on the power spectrum depends strongly on the adopted baryon physics (e.g. van Daalen et al. 2011, 2020; Mummery et al. 2017; Chisari et al. 2018, 2019; Springel et al. 2018). Given the selection of prior ranges for each subgrid parameter to conservatively bracket existing observational constraints and their corresponding uncertainties, our

⁴We note that for our simulations, we include all stellar and gas particles within a spherical overdensity radius. Hence, in order to make reasonable comparisons with the fits in Akino et al. (2022), we included an additional 15 per cent contribution to the total stellar masses from the contribution of blue galaxies, and 30 per cent additional stellar mass to the brightest cluster galaxies to account for the diffuse intracluster light (ICL; see Akino et al. 2022).

⁵We characterize the ratio in terms of the power spectrum, $P(k)$. But note that since $P(k) \propto \Delta^2(k)$, our results are equivalent to a ratio of $\Delta_{\text{hydro}}^2(k)/\Delta_{\text{DM}}^2(k)$.

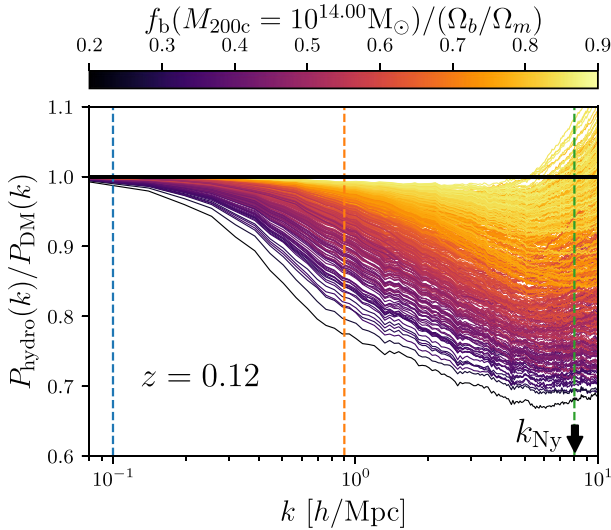


Figure 2. Fractional impact of baryons on the total matter power spectrum for the 400 models at redshift $z = 0.125$. Colour-coding represents the total baryon fraction of haloes of mass $M_{200c} = 10^{14} M_{\odot}$. Vertical dashed lines indicate three equally log-spaced scales, $k = 0.1, 0.9$, and $8.0 h \text{ Mpc}^{-1}$, shown in Fig. 3. The black arrow indicates the Nyquist frequency of the simulations k_{Ny} .

simulation suite encompasses baryonic effects that surpass observational bounds, with a range of parameters more extreme than those used in most cosmological simulations (see e.g. Tröster et al. 2022). The effect of AGN feedback is of particular importance, due to its ejective nature at early times (McCarthy et al. 2011). The removal of large quantities of gas not only affects the distribution of the remaining gas, but it also has a gravitational effect on the dark matter, causing it to expand as well. This effect is referred to as the ‘back reaction’ on baryons on the dark matter (e.g. van Daalen et al. 2011). The net result of baryon ejection and dark matter expansion is that $P(k)$ can be suppressed by a non-negligible amount (relative to forthcoming statistical errors from cosmic shear surveys) out to wavenumbers of $k \sim 0.1 h \text{ Mpc}^{-1}$, corresponding to physical scales of $\sim 30 \text{ Mpc } h^{-1}$.

In the following sections, we will present our model based on the correlation between the power spectrum suppression and the total baryon fraction of haloes of different masses. Note that in this study, halo masses are measured within both R_{200c} and R_{500c} , i.e. the radius within which the mean density is 200 or 500 times the critical density of the Universe, respectively. While both spherical overdensity apertures provide similar results, we find that using R_{200c} gives slightly better accuracy. Hence, in the following sections, we will present all of our results and plots in terms of R_{200c} , while also provide the best-fitting parameters for our model in both R_{200c} and R_{500c} .

3.1 The correlation between the power spectrum suppression and the total baryon fraction

Recently, van Daalen et al. (2020) presented an empirical model that uses the mean baryon fraction of group mass haloes ($M_{200c} = 10^{14} M_{\odot}$) as a predictor for the power suppression for scales $k \lesssim 1.0 h \text{ Mpc}^{-1}$, achieving an accuracy of ≈ 1 per cent. Fundamentally, the effect of baryon physics on the matter power spectrum is a gravitational process, which is why we expect the baryon fraction of haloes (which is a direct measure of the mass that has been removed)

that contribute most significantly to $P(k)$ to be a strong predictor of the effects on $P(k)$.

Taking our cue from van Daalen et al. (2020), in Fig. 2 we colour code each run by its median total baryon fraction of haloes of mass $M_{200c} = 10^{14} M_{\odot}$, and normalized by the universal baryon fraction (Ω_b/Ω_m). In agreement with van Daalen et al. (2020), the figure shows a clear gradient of the power spectrum suppression as a function of the total baryon fraction of group mass of haloes. At closer inspection, the correlation is not perfect at small scales of $k \geq 3 h \text{ Mpc}^{-1}$. Nevertheless, we will show below that a simple and accurate model for the $P(k)$ suppression can be formulated in terms of the median baryon fraction.

In this work, we aim to extend the van Daalen et al. (2020) model to smaller scales and higher redshifts, which we will achieve by taking into account the increasing importance of lower mass haloes as we push in both directions. In order to assess the strength of the correlation between the baryon fraction of haloes of different mass, we compute the power spectra at three equally spaced scales in $\log_{10} k$ ($k = 0.1, 0.9$, and $8.0 h \text{ Mpc}^{-1}$), which are shown as vertical dashed lines in Fig. 2. We show in Fig. 3 the fractional impact of baryons on the total matter power spectrum as a function the total baryon fraction for two different haloes of masses (10^{13} and $10^{14} M_{\odot}$) for these three scales. The colours correspond to the same scales of the vertical lines in Fig. 2. The two sets of SPH schemes used for the simulations are shown with different symbols.

Fig. 3 shows that the suppression of the power spectrum at a given scale correlates better with the total baryon fraction of haloes of different mass. For instance, the suppression at $k = 0.9 h \text{ Mpc}^{-1}$ (orange symbols) is better correlated with the baryon fraction of haloes of mass $10^{14} M_{\odot}$ (right) than with $10^{13} M_{\odot}$ (left). On the other hand, for small scales, $k = 8.0 h \text{ Mpc}^{-1}$ (green symbols), the suppression is better correlated with the baryon fraction of lower mass haloes (left). This dependence can be explained in terms of the differing relative contribution of haloes of different mass to the total matter power spectrum (e.g. Daalen & Schaye 2015; Mead et al. 2020). The strength of the correlation for very large scales, $k = 0.1 h \text{ Mpc}^{-1}$ (blue symbols), seems insensitive to choice of halo mass (for the two masses shown here) which is mainly just because feedback does not significantly alter the power spectrum on these scales, so there is little dynamic range in $P_{\text{hydro}}(k)/P_{\text{DM}}(k)$ to couple to the baryon fractions.

Note that while the selection of SPH flavor for a particular set of subgrid parameters may have an impact on the resulting baryon fractions, the response exhibited in $P(k)$ is invariant with regard to the selection of SPH flavor at a given baryon fraction. This is consistent with the findings of van Daalen et al. (2020).

In Fig. 4, we quantify the strength of the correlation between the suppression of the total matter power spectrum and the total baryon fraction of haloes of different mass as a function of scale k using the Spearman’s rank correlation coefficient, ρ_s . We calculate the correlation coefficient at 40 equally spaced scales⁶ in $\log_{10} k$ in the range $0.1 \leq k [h \text{ Mpc}^{-1}] \leq 10$. Each panel represents a different redshift in the simulations.

For each redshift, the lines are colour coded by the halo mass used to estimate the baryon fraction. We use a mass bin width of 0.1 dex around each indicated halo mass. To help guide the eye, we have highlighted the area where $\rho_s \geq 0.90$, i.e. where the correlation between the suppression of the total matter power spectrum and

⁶We re-bin the power spectra and ‘smooth’ them by calculating the median power for each bin.

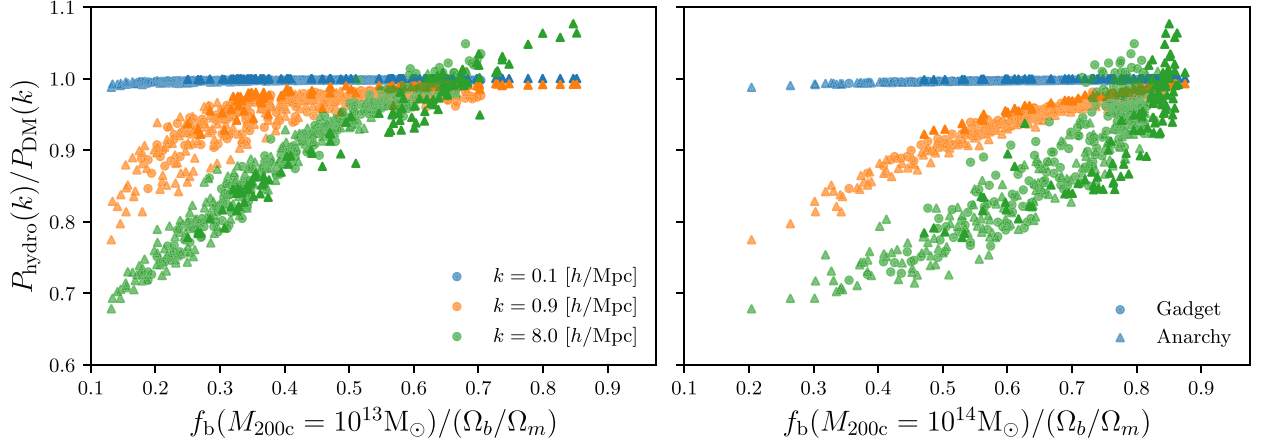


Figure 3. Fractional impact of baryons on the total matter power spectrum at $k = 0.1, 0.9$, and $8.0 \, h \, \text{Mpc}^{-1}$ as a function of the total baryon fraction of haloes of mass $M_{200c} = 10^{13} \, M_\odot$ (left) and $M_{200c} = 10^{14} \, M_\odot$ (right). Colours indicate the suppression at the three equally spaced scales in logarithmic scale ($\log_{10} k$) in Fig. 2. Each symbol represents 1 of the 400 models. The two sets of SPH schemes used for the simulations are shown with different symbols. The suppression of the power spectrum at a given scale correlates better with the total baryon fraction of haloes of different mass. For instance, the suppression at $k = 0.9 \, h \, \text{Mpc}^{-1}$ (orange symbols) is better correlated with the baryon fraction of haloes of mass $M_{200c} = 10^{14} \, M_\odot$ (right) than with $M_{200c} = 10^{13} \, M_\odot$ (left). The converse is true for $k = 8.0 \, h \, \text{Mpc}^{-1}$ (green symbols), while the strength of the correlation seems insensitive to the two masses shown for $k = 0.1 \, h \, \text{Mpc}^{-1}$ (blue symbols).

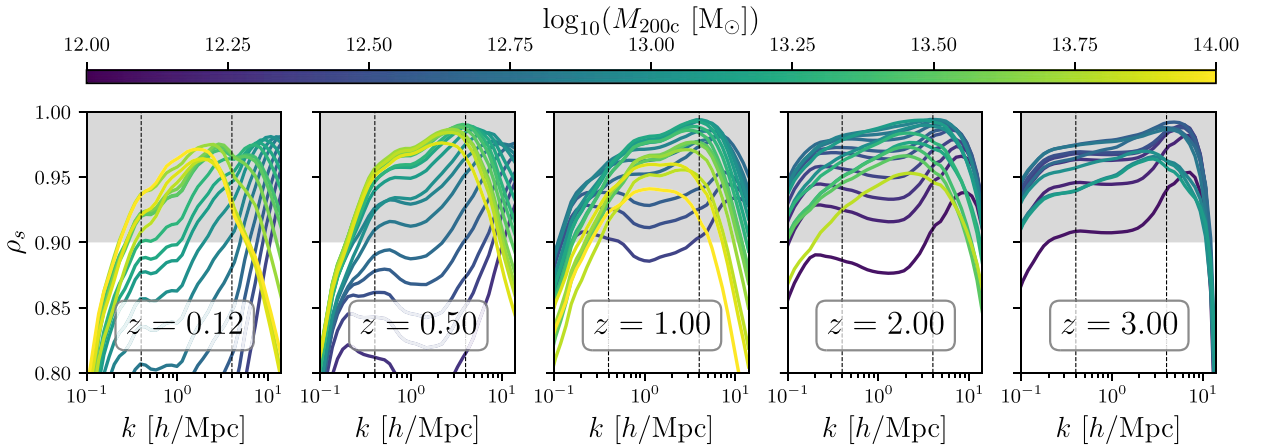


Figure 4. Spearman's rank correlation coefficient ρ_s , between the suppression of the total matter power spectrum and the total baryon fraction of haloes of different mass as a function of scale k . The five panels represent different redshifts and colour-coding represents different halo masses M_{200c} . The grey shaded area indicates where $\rho_s \geq 0.90$. Independent of redshift, for small scales ($k \gtrsim 4 \, h \, \text{Mpc}^{-1}$), the baryon fraction of progressively lower mass haloes provide better correlation for increasing k . For intermediate scales ($0.4 \lesssim k [h \, \text{Mpc}^{-1}] \lesssim 4$), haloes with a characteristic mass (which evolves with redshift) provide the best correlation. For $k \lesssim 0.4 \, h \, \text{Mpc}^{-1}$, the strength of the correlation decreases rapidly regardless of the halo mass.

the total baryon fraction is exceptionally strong, and it can be well described using a monotonic function.

To guide the eye, we have added three vertical lines in Fig. 4 to roughly divide small, intermediate, and large scales:

(i) *Small scales* ($k \gtrsim 4 \, h \, \text{Mpc}^{-1}$): The baryon fraction of progressively lower mass haloes provides a better correlation with the suppression for increasing k . The converse is also true; i.e. the strength of the correlation using higher mass haloes decreases for increasing k .

(ii) *Intermediate scales* ($0.4 \lesssim k [h \, \text{Mpc}^{-1}] \lesssim 4$): Haloes with a characteristic mass (which evolves with redshift) provide the best correlation.

(iii) *Large scales* ($k \lesssim 0.4 \, h \, \text{Mpc}^{-1}$): The strength of the correlation decreases rapidly regardless of the halo mass. This is expected, as for (very)large scales, the suppression of the total matter power spectrum should be close to unity independent of the baryon fraction of haloes (non-monotonic, see Fig. 3).

3.2 The optimal halo mass

As we aim to determine the halo masses that best predict the suppression of the power spectrum at each scale, we define the *optimal mass*, $\hat{M}_{k,200c}$, as the halo mass used to measure the baryon fraction that maximizes ρ_s as a function of scale and redshift:

$$\hat{M}_{k,200c}(k, z) = M_{200c} \Big|_{\max\{\rho_s(k, z)\}}. \quad (1)$$

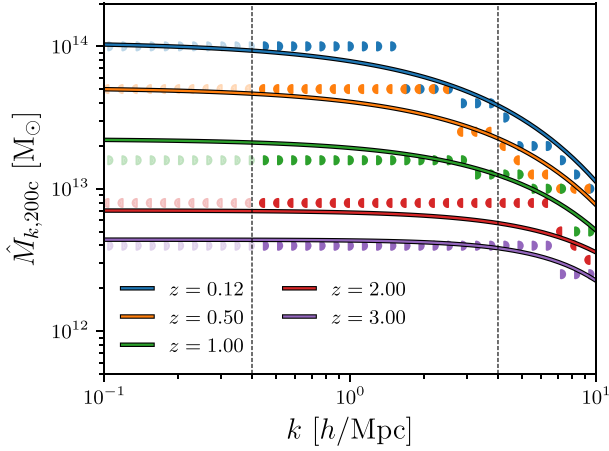


Figure 5. Optimal mass $\hat{M}_{k,200c}$ as a function of scale k at five different redshifts. Coloured symbols represent the halo mass that provides the best correlation between the power spectrum suppression and the baryon fraction of such haloes at each k bin. Solid lines show the best fit using equation (2). For small scales ($k \gtrsim 4 h \text{ Mpc}^{-1}$), the optimal mass decreases for increasing k . For intermediate scales ($0.4 \lesssim k [h \text{ Mpc}^{-1}] \lesssim 4$), haloes with a characteristic mass provide the best correlation (i.e. the optimal mass plateaus). For large scales, the suppression of the total matter power spectrum should approach unity, independent of the baryon fraction of haloes (see Figs 3 and 4). Therefore, in order to avoid noise while fitting, we assume that \hat{M}_k remains constant for $k \leq 0.4 h \text{ Mpc}^{-1}$ and data points are shown in faint colours.

Fig. 5 shows $\hat{M}_{k,200c}$ as a function of scale k for the five redshifts shown in Fig. 4. We used alternating filled semicircles to allow us to show overlapping points from different redshift bins. For each redshift, the shape of the relation roughly reflects the three distinct behaviours described above: at small scales ($k \gtrsim 4 h \text{ Mpc}^{-1}$), the baryon fraction of haloes decreasing in mass provides a better correlation for increasing k ; at intermediate scales ($0.4 \lesssim k [h \text{ Mpc}^{-1}] \lesssim 4$), haloes with a characteristic mass provide the best correlation (i.e. the optimal mass plateaus); and at large scales, the strength of the correlation decreases rapidly regardless of the halo mass, as the total matter power spectrum suppression is close to unity, independent of the baryon fraction of haloes (see Figs 3 and 4). In order to avoid noise while fitting, without loss of generality, we assume that \hat{M}_k remains constant for $k \leq 0.4 h \text{ Mpc}^{-1}$ (data points are shown in faint colours). Note that this is also consistent with the results of van Daalen et al. (2020), i.e. to a good approximation, the total baryon fraction for a single halo mass provides enough information to predict the power spectrum suppression due to baryons for $k \lesssim 1.0 h \text{ Mpc}^{-1}$.

It is worth noting that the optimal mass exhibits a monotonic behaviour as function of redshift. Therefore, we approximate $\hat{M}_{k,200c}$ using the following exponential plateau functional form,

$$\log_{10}(\hat{M}_{k,200c}(k, z)) = \alpha(z) - [\alpha(z) - \beta(z)]k^{\gamma(z)}, \quad (2)$$

where $\alpha(z)$ is the asymptote of the function as $k \rightarrow 0$, $\beta(z)$ is the value for $k = 1$, and $\gamma(z)$ characterizes the rate of change. We impose a simple polynomial ansatz for the redshift dependence of these parameters as follows:

$$X(z) = \sum_{i=0}^2 X_i(1+z)^i, \quad (3)$$

where $X = \{\alpha, \beta, \gamma\}$.

Table 3. Best-fitting parameters for the optimal mass in equation (2), expressed either in terms of M_{200c} or M_{500c} . For the exponential plateau functional form, α is the asymptote of the function as $k \rightarrow 0$, β is the value for $k = 1$, and γ characterizes the rate of change. The redshift dependence of the parameter is given in equation (3).

Spherical overdensity		Parameter		
		$i = 0$	$i = 1$	$i = 2$
M_{200c}	α_i	15.243	− 1.243	0.148
	β_i	14.969	− 1.099	0.129
	γ_i	0.800	− 0.017	0.061
M_{500c}	α_i	14.783	− 0.999	0.120
	β_i	14.620	− 0.913	0.108
	γ_i	0.967	− 0.031	0.026

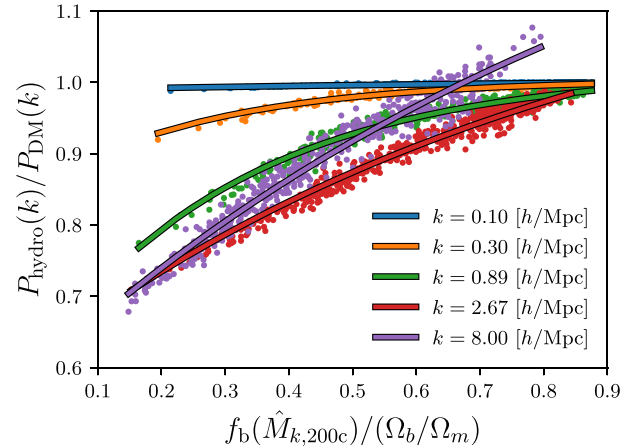


Figure 6. Suppression of the total matter power spectrum at redshift $z = 0.125$ as a function of the total baryon fraction using the optimal mass of haloes $\hat{M}_{k,200c}$ for each corresponding scale k . Colours indicate the suppression at five equally spaced scales in logarithmic scale ($\log_{10}k$). Each symbol represents 1 of the 400 models. Solid lines show the best fit using equation (4) without a scale and redshift dependence, i.e. each scale shown has been fitted using a simple exponential plateau function.

We use a simple least-squares estimator to fit the simulation data and we provide the best-fitting parameter values in Table 3. The solid lines in Fig. 5 show the best fit using equation (2).

In Fig. 6, we use the optimal mass $\hat{M}_{k,200c}(k, z)$ from equation (2) to compute the suppression of the total matter power spectrum as a function of the total baryon fraction at redshift $z = 0.125$. Note that we do not assume any functional form for the $f_b - M_{\text{halo}}$ relation. Rather, we use a piecewise cubic Hermite interpolating polynomial using the narrow binned data from the simulations to compute the baryon fractions at the required optimal mass.

The figure shows that using $\hat{M}_{k,200c}(k, z)$ significantly reduces the scatter of the relationship at all scales shown, from $k = 0.1 h \text{ Mpc}^{-1}$, up to the one-dimensional Nyquist frequency of the simulations, $k_{\text{Ny}} = \pi N/L \approx 8 h \text{ Mpc}^{-1}$, where N is the cube root of the total number of particles, and L is the length of the cubic box. Given the wide range of physical models and the different hydrodynamic schemes used in this study, the fact that such a tight correlation is achieved solely based on the baryon fraction of haloes of different masses is remarkable. One might have expected that a thorough characterization of the mass density profiles around haloes (as opposed to a simple aperture baryon fraction) would have been required to obtain this level of precision. For example, Debackere, Schaye & Hoekstra (2020) have shown that the behaviour of the

profiles between r_{500} and r_{200} can affect the matter power spectrum if the profiles are allowed to vary significantly over this range. Evidently, there is a very strong physical correlation between the integrated baryon fractions and the shape of the profile over this range in our simulations, such that most of the physical effect on the power spectrum is encoded in the integrated baryon fractions alone.

3.3 *SP(k) – a model for the suppression of P(k)*

We find that, similar to equation (2), an exponential plateau functional form provides a good approximation to the fractional impact of baryons on the total matter power spectrum:

$$P_{\text{hydro}}(k)/P_{\text{DM}}(k) = \lambda(k, z) - [\lambda(k, z) - \mu(k, z)] \exp[-\nu(k, z) \tilde{f}_b], \quad (4)$$

where \tilde{f}_b is the baryon fraction at the optimal halo mass normalized by the universal baryon fraction, i.e.

$$\tilde{f}_b = f_b(\hat{M}_{k,200c}(k, z))/(\Omega_b/\Omega_m), \quad (5)$$

and where $\lambda(k, z)$ is the asymptote of the function for large \tilde{f}_b values, $\mu(k, z)$ is the suppression for $\tilde{f}_b = 0$, and $\nu(k, z)$ is the rate of change parameter. In the left-hand panel of Fig. 7, we show an illustration of the suppression of the total matter power spectrum as a function of \tilde{f}_b , as parametrized by equation (4).

We use this functional form to fit the 400 simulations at each scale individually. The best fit for each scale at $z = 0.125$ is shown as solid lines in Fig. 6. Equation (4), without a scale and redshift dependence, provides a good approximation to the full range of scales up to the Nyquist frequency. Furthermore, we approximate each of the parameters in equation (4) based on their behaviour as a function of scale using an exponential, logistic and lognormal functional forms (respectively):

$$\lambda(k, z) = 1 + \lambda_a(z) \exp(\lambda_b(z) \log_{10}(k)), \quad (6)$$

$$\mu(k, z) = \mu_a(z) + \frac{1 - \mu_a(z)}{1 + \exp(\mu_b(z) \log_{10}(k) + \mu_c(z))}, \quad (7)$$

$$\nu(k, z) = \nu_a(z) \exp\left(-\frac{(\log_{10}(k) - \nu_b(z))^2}{2\nu_c(z)^2}\right). \quad (8)$$

We model the evolution of each parameter as a polynomial function in redshift, using equation (3), with $X = \{\lambda_a, \lambda_b, \mu_a, \mu_b, \mu_c, \nu_a, \nu_b, \nu_c\}$ accordingly.

In order to reduce the sensitivity to outliers in our wide range of baryon fractions in the 400 simulations, for our parameter estimation utilized a Huber loss function defined as

$$L_\delta(y, \hat{y}) = \begin{cases} \frac{1}{2}(y - \hat{y})^2 & \text{for } |y - \hat{y}| \leq \delta \\ \delta(|y - \hat{y}| - \frac{1}{2}\delta) & \text{otherwise,} \end{cases} \quad (9)$$

where y_i is the measured variable $P_{\text{hydro}}(k)/P_{\text{DM}}(k)$ and \hat{y}_i is the predicted values from our model using equations (4)–(8). The value of δ was chosen to represent the median absolute deviation. We minimized the loss function equation (9) at five anchor redshifts ($z = 0.125, 0.5, 1.0, 2.0$, and 3.0). The best-fitting parameters are provided in Table 4.

We present a publicly available PYTHON implementation of our model as `py-SP(k)`. All relevant information about installation and usage can be found at <https://github.com/jemme07/pyspk>.

The choice of functional form in equation (4) provides insight into the effect of baryons on the total matter power spectrum. In the right-hand side of Fig. 7, we show the scale and redshift dependence of the best-fitting parameters in Table 4. The figure shows that for very large

scales ($k \sim 0.1 h \text{ Mpc}^{-1}$), both λ and μ are close to unity, especially at high redshift, i.e. the suppression of the matter power spectrum is a flat function of the baryon fraction of haloes. In other words, at very large scales, baryons do not significantly affect the matter power spectrum. As we examine progressively smaller scales, on the one hand, λ is an increasing function of k , i.e. models with large median baryon fraction ($\tilde{f}_b \rightarrow 1$) exhibit weaker suppression values, or even a power spectrum enhancement for small scales ($k \gtrsim 4 h \text{ Mpc}^{-1}$). On the other hand, μ is a decreasing function of k , i.e. models with low median baryon fraction ($\tilde{f}_b \rightarrow 0$) present stronger suppression, especially for scales $k \gtrsim 1 h \text{ Mpc}^{-1}$. Finally, the rate of change parameter, ν , dictates how quickly the suppression changes from the value of μ at low baryon fractions, to λ at large baryon fractions. For very large scales, the value of ν is irrelevant, as both λ and μ are close to unity. For scales $0.2 \lesssim k [h \text{ Mpc}^{-1}] \lesssim 1$, $\nu \gg 1$, i.e. there is a rapid transition from strong suppression values at low baryon fractions, to weak suppression values at large baryon fractions. This is indeed the case if we examine the behaviour the suppression for that scales range in Fig. 6 (orange and green lines). There is a rapid transition giving rise to a ‘curved’ relationship that asymptotes to the value of λ at high baryon fractions. For smaller scales ($k \gtrsim 1 h \text{ Mpc}^{-1}$), $\nu \rightarrow 0$, i.e. there is a slow transition from strong suppression values at low baryon fractions, to weak suppression (or enhancement) values at large baryon fractions. Examining Fig. 6 again, this slow transition gives rise to a more linear relation (rather than curved) for small scales (red and purple lines).

3.4 Model accuracy

The level of agreement between the simulations and *SP(k)* is shown in Fig. 8. To appreciate the performance of the model, in the top panel, we show the fractional impact of baryons on the total matter power spectrum against data for 8 randomly selected models from the 400 simulations at redshift $z = 0.1$ spanning a wide range of baryon fractions. Solid lines show the power spectra computed directly from the simulations. Our model is not restrictive to a particular shape of the baryon fraction–halo mass relation, and in order to show its flexibility, we show in dashed lines our best-fitting model using binned data for the f_b – M_{halo} relation obtained from the simulations, while in dotted lines we use a simple power-law functional form fit to the same relation in the halo mass range of interest (see also Section 3.5).

While the data from the simulations show complex features, the overall behaviour is captured well by our model. By construction, the suppression goes to unity for low values of k . The characteristic ‘spoon-like’ shape of the power spectrum suppression, the amplitude of which depends on the baryon fraction of haloes, is well captured by the model. In the middle panel, we show the ratio between the measurements of the suppression as measured in the hydrodynamical simulations, to our best-fitting model for all 400 models. The solid blue line shows the median ratio as a function of k , while the blue dashed lines enclose the 15.9 and 84.1 percentiles (one sigma). In the bottom panel, we show the mean absolute percentage error (MAPE), defined as

$$\text{MAPE} = \frac{100 \text{ per cent}}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right|, \quad (10)$$

where y_i is the measured variable $P_{\text{hydro}}(k)/P_{\text{DM}}(k)$ and \hat{y}_i is the predicted values from our model. The figure shows that the model is able to reproduce the wide range of baryonic effects explored here remarkably well. For all the scale range up to the Nyquist frequency,

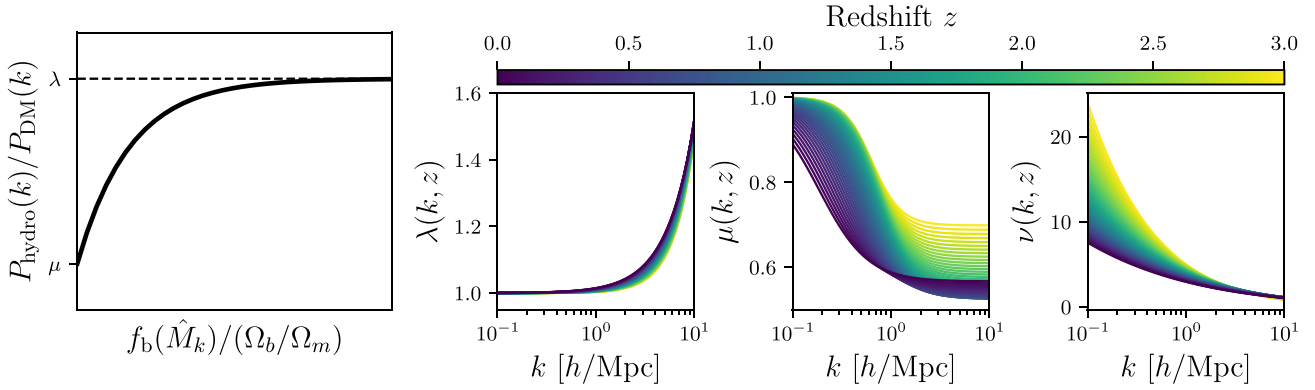


Figure 7. Left: Illustration of the suppression of the total matter power spectrum as a function of \tilde{f}_b as parametrized by equation (4). Right: Scale and redshift dependence of the best-fitting parameters in Table 4. $\lambda(k, z)$ is the asymptote of the function in equation (4) for large \tilde{f}_b values, modelled using an exponential functional form equation (6), $\mu(k, z)$ is the suppression for $\tilde{f}_b = 0$, modelled using a logistic functional form equation (7), and $\nu(k, z)$ is the rate of change, modelled with lognormal functional form equation (8).

Table 4. Best-fitting parameters for power spectrum suppression in equation (4). For the exponential plateau functional form, λ is the asymptote of the function for large k values, μ is the suppression for $\tilde{f}_b = 0$, and ν characterizes the rate of change. The k scale and redshift dependence of the parameter are given in equation (3) and equations (6)–(8).

Spherical overdensity		Parameter		
		$i = 0$	$i = 1$	$i = 2$
M_{200c}	$\lambda_{a,i}$	0.021	−0.007	0.000
	$\lambda_{b,i}$	3.087	0.452	0.001
	$\mu_{a,i}$	0.693	−0.169	0.042
	$\mu_{b,i}$	3.161	0.861	0.011
	$\mu_{c,i}$	5.532	−3.086	0.508
	$\nu_{a,i}$	413.009	311.639	37.891
	$\nu_{b,i}$	−11.243	−0.344	0.334
	$\nu_{c,i}$	3.476	−0.018	−0.082
M_{500c}	$\lambda_{a,i}$	0.019	−0.007	0.000
	$\lambda_{b,i}$	2.956	0.620	−0.001
	$\mu_{a,i}$	0.715	−0.192	0.049
	$\mu_{b,i}$	3.385	0.965	−0.068
	$\mu_{c,i}$	4.457	−2.191	0.454
	$\nu_{a,i}$	478.864	429.887	249.256
	$\nu_{b,i}$	−11.227	−0.558	0.448
	$\nu_{c,i}$	3.499	−0.084	−0.092

the MAPE is ≈ 1 per cent. For large scales, the model recovers the power spectrum suppression to better than 1 per cent accuracy, while for the smallest scales considered, the error for most models within <5 per cent, with an average of <2 per cent for the entire range of baryon fractions.

We accomplish this significant reduction in the scatter for the power spectrum suppression–baryon fraction relation in our model by selecting the optimal mass that maximizes ρ_s for each scale. Nevertheless, for the smallest scales that can be resolved in the simulations (close to the Nyquist frequency, $k_{Ny} \approx 8 h \text{ Mpc}^{-1}$), while the model can recover a mostly monotonic relation, there is still significant scatter around it, as shown with purple dots in Fig. 6. As the scatter increases with scale k , this gives rise to the heteroscedastic behaviour of the errors, i.e. the error in the model increases with scale (dashed blue lines). None the less, our model is an unbiased estimator of the true baryonic effects, as the ratio between the measurements of the suppression as measured in the hydrodynamical simulations to our best-fitting model is centred around unity at all scales.

In Fig. 9, we use the individual model errors for the 400 simulations to compute a 2D mean, and maximum, absolute percentage error distribution as a function of scale and baryon fraction for different redshifts. The baryon fraction is computed using the optimal mass of haloes for each corresponding scale and redshift using equation (2). The mean and maximum errors are computed over linearly spaced baryon fraction bins and logarithmically spaced k bins between 0.1 and $10.0 h \text{ Mpc}^{-1}$. Only pixels with five or more data points are shown. The top panels show that our model has a mean error of $\lesssim 2$ per cent for the entire scale and redshift range. Furthermore, as shown in the bottom panels of Fig. 9, our model is less accurate for increasing k . This can be seen as the increasing contours of maximum absolute percentage error towards the right-hand side of each panel at the bottom of the figure. The maximum absolute percentage error reaches values of 4 per cent to 6 per cent for scales close to the Nyquist frequency at low redshift, and up to 8 per cent to 10 per cent at $z = 1$ –2. Additionally, the model is less accurate for simulations with extreme feedback prescriptions (low baryon fractions), for scales $k \sim 1 h \text{ Mpc}^{-1}$. While a wide range of feedback prescriptions were used to build and calibrate our model, it is important to highlight that for simulations with mean group-scale baryon fractions roughly consistent with observations, the accuracy of our model is at its best, with mean error of ≈ 1 per cent, and a maximum mean error of ≈ 3 per cent for the whole scale and redshift range. As an example, Fig. 9 shows such accuracy around the fiducial BAHAMAS simulation (orange line), which was specifically calibrated to reproduce the galaxy stellar mass function and the observed gas fractions of galaxy groups and clusters at redshift $z \approx 0$ (see Section 4),

3.5 How the baryon fraction shapes the suppression of the total matter power spectrum

We now use our model to systematically explore the effect of the f_b – M_{halo} relation on the shape of the suppression of the total matter power spectrum. For the mass range that can be relatively well probed in current X-ray and Sunyaev–Zel’dovich effect observations (roughly $10^{13} \lesssim M_{200c} [M_\odot] \lesssim 10^{15}$), the total baryon fraction of haloes can be roughly approximated by a power law with constant slope (e.g. Mulroy et al. 2019; Akino et al. 2022). Hence, we use the

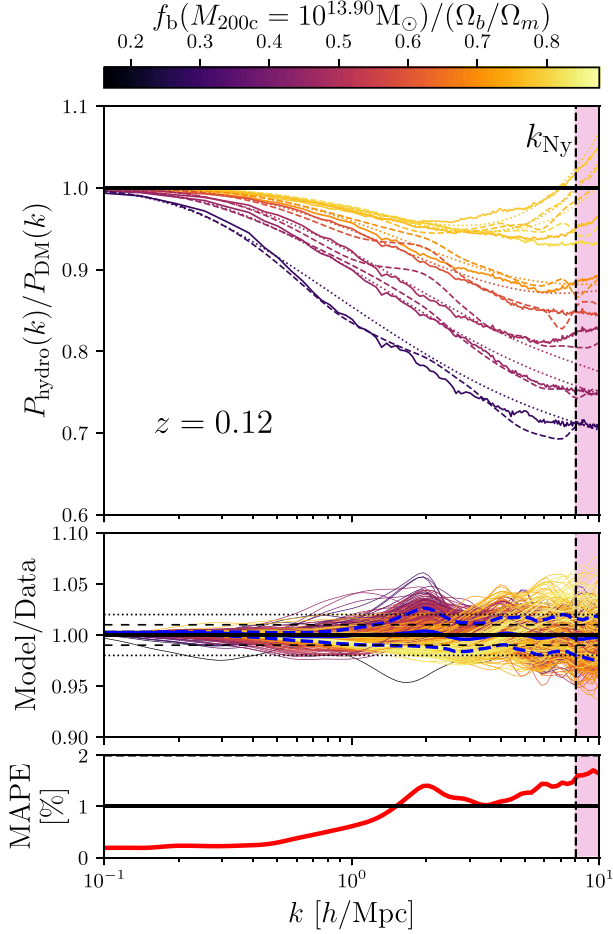


Figure 8. Top panel: Fractional impact of baryons on the total matter power spectrum for 8 randomly selected samples from the 400 models. Solid lines use power spectra computed from the simulations. Dashed lines represent the best-fitting model obtained with $SP(k)$ using binned data for the f_b – M_{halo} relation. Dotted lines use a simple power-law fit to the total baryon fraction instead. Colour-coding represents the total baryon fraction of haloes of mass $M_{200c} = 10^{13.90} M_{\odot}$, which corresponds to the optimal mass at $k = 1.0 h \text{ Mpc}^{-1}$, i.e. $\hat{M}_{k,200c}(k = 1.0 h \text{ Mpc}^{-1}) = 10^{13.90} M_{\odot}$. The vertical dashed line indicates the Nyquist frequency of the simulations k_{Ny} . The pink lightly shaded region indicates the scales where our model is not a good indicator of the uncertainty as $k > k_{\text{Ny}}$. Middle panel: Ratio between the measurements of the suppression in the power spectrum induced by baryons as measured in the hydrodynamical simulations to our best-fitting model for all 400 models. Dashed black lines indicate 1 per cent accuracy, while dotted lines indicate 2 per cent accuracy. The solid blue line shows the median ratio, while the blue dashed lines enclose the 15.9 and 84.1 percentile. Bottom panel: Mean absolute percentage error for our best-fitting model for all the 400 simulation models.

following parametrization,

$$f_b/(\Omega_b/\Omega_m) = a \left(\frac{M_{200c}}{10^{13.5} M_{\odot}} \right)^b, \quad (11)$$

where a is the normalization of the f_b – M_{200c} relation at $M_{200c} = 10^{13.5} M_{\odot}$ and b is the power-law slope.

We use our best-fitting parameters for M_{200c} in Tables 3 and 4 and illustrate the effect of changing the power-law fitting parameters a and b at redshift $z = 0.125$ in Fig. 10. The top panel shows variations in the power-law normalization while keeping $b = 0.3$. The bottom panel shows variations in the power-law slope while keeping $a =$

0.4. The parameters $a = 0.4$ and $b = 0.3$ were selected as they are roughly consistent with the baryon fraction–halo mass relation for the fiducial BAHAMAS simulation at the corresponding redshift.

Consistent with van Daalen et al. (2020), the top panel shows the large effect of the normalization of the baryon fraction–halo mass relation on the suppression of the matter power spectrum. As mentioned above, the effect of baryon physics on the matter power spectrum is fundamentally a gravitational process. Hence, the larger the amount of mass removed from haloes (lower f_b), the stronger the suppression on $P(k)$ relative to a baryon complete (DM-only) model. Furthermore, very low baryon fractions (associated with stronger AGN feedback) can have a non-negligible impact the power spectrum suppression on very large scales ($k \sim 0.1 h \text{ Mpc}^{-1}$, dark-blue line).

In the bottom panel, we show that the slope of the baryon fraction–halo mass relation has a large effect on the scale of the minimum of the ‘spoon-like’ shape of the power spectrum suppression.

3.6 Alternative approaches to fitting the suppression

Our approach above was to apply simple parametric forms to describe the correlation between the power spectrum suppression and the baryon fraction of haloes that contribute most significantly at a given scale. We have shown that, in spite of its simplicity, the model provides a very accurate description of our large simulation data set. Nevertheless, our approach to modelling the simulation data is not unique and there may be alternative (e.g. non-parametric) approaches that could lead to even more accurate results. For example, principle component analysis applied to simulation power spectra has been used as a way of characterizing the impact of baryons and allowing for their marginalization in cosmological pipelines (e.g. Eifler et al. 2015). A related strategy is to model the principle components via Gaussian Processes (e.g. Heitmann et al. 2014) or sparse polynomial chaos expansion (e.g. Euclid Collaboration 2019), although to our knowledge these methods have so far been applied to only the non-linear correction to the matter power spectrum and not on the impact of baryons. In addition, neural networks have recently been used to construct a baryon emulator by Aricò et al. (2021). The neural net was trained on a large sample of power spectra generated by the baryonification approach.

While there are various advantages and disadvantages to the alternative approaches described above, we have elected to use the approach described above because it is more intuitive in terms of visualizing the correlations and it allows us to more easily enforce physically motivated boundary conditions, such as requiring that suppression of the power spectrum goes to unity as $k \rightarrow 0$ (very large scales). That being said, we have also experimented with both a Gaussian Process-based emulator and a simple neural network using the baryon fractions as the input quantities. We find that these non-parametric approaches do not significantly improve on the accuracy of our fiducial analytical method and we therefore elected to retain the latter as our main model.

4 TESTING AGAINST BAHAMAS

We employ the BAHAMAS suite of cosmological hydrodynamical simulations (McCarthy et al. 2017, 2018) to test the accuracy of our model using simulations outside our calibration set. The reference BAHAMAS simulation uses the same cosmology and particle mass resolution as the simulations used in this study, but in a much larger box. The fiducial BAHAMAS run consist of a comoving volume with $400 c\text{Mpc } h^{-1}$ on a side and 2×10^{24} particles of $m_{\text{DM}} = 3.85 \times 10^9 M_{\odot} h^{-1}$ dark matter particle mass, and $m_g =$

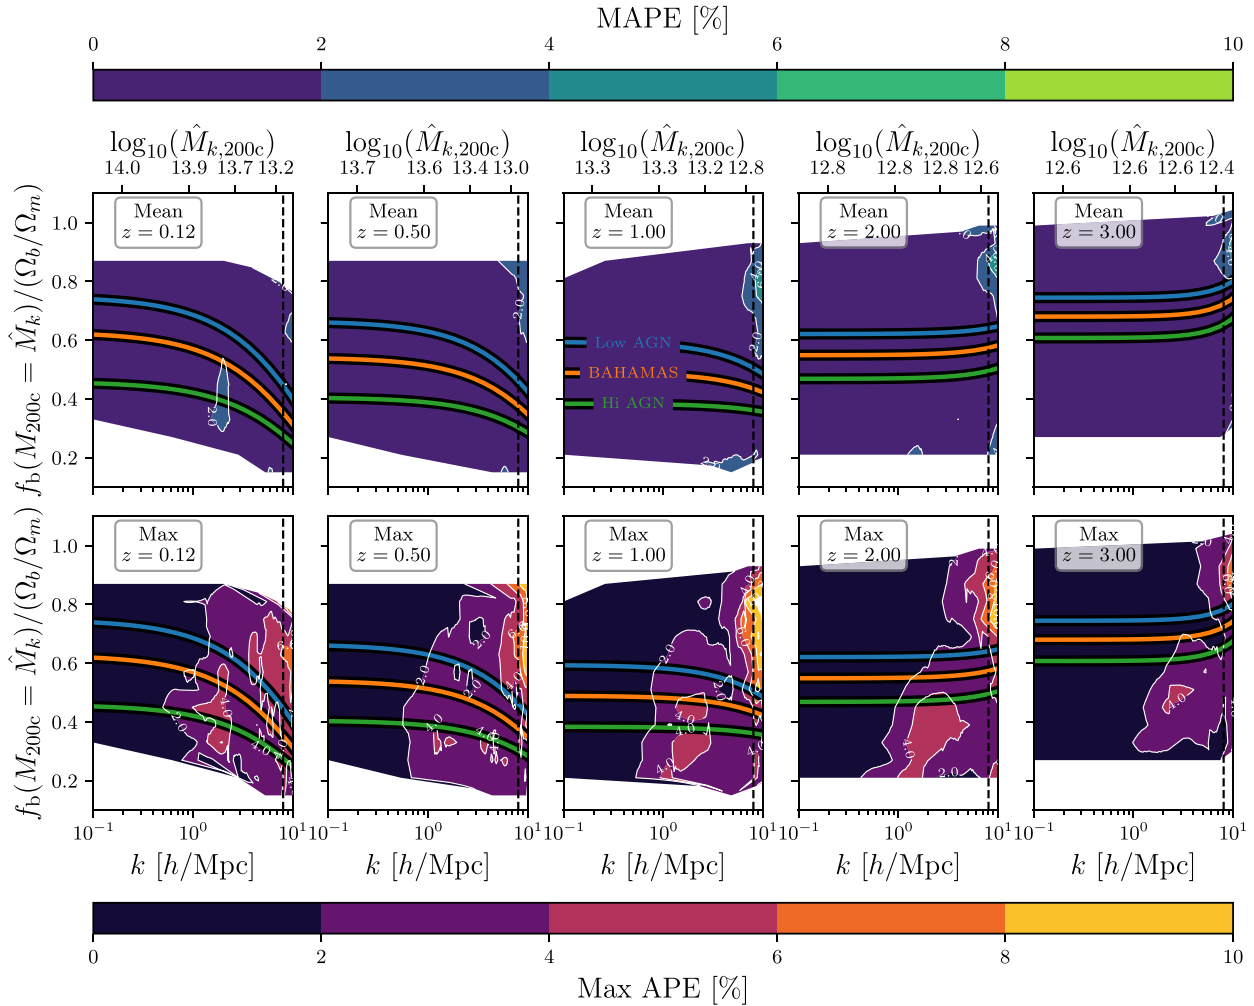


Figure 9. 2D mean (top) and maximum (bottom) absolute percentage error distribution for our best-fitting model at $z = 0.125, 0.5, 1.0, 2.0$, and 3.0 . The baryon fraction is computed using the optimal mass of haloes $\hat{M}_{k,200c}$. The vertical dashed line indicates the Nyquist frequency of the simulations k_{Ny} . $\hat{M}_{k,200c}$ as function of k and z is calculated using equation (2) and given along the top axis. The three coloured lines indicate the total baryon fraction using the optimal mass for the three BAHAMAS models (Low AGN, fiducial BAHAMAS, and High AGN; McCarthy et al. 2017, 2018).

$7.66 \times 10^8 M_{\odot} h^{-1}$ initial gas particle mass. A full discussion of the subgrid implementation, including the prescriptions for star formation, gas heating and cooling, BH formation, and SNe and AGN feedback models can be found in McCarthy et al. (2017, see also Schaye et al. 2010). As the BAHAMAS suite of simulations aim to study the impact of baryonic process on LSS, it was of paramount importance to ensure that group size and more massive haloes had the correct baryon fractions, as these haloes contribute the most to the matter power spectrum. Therefore, the BAHAMAS approach was to calibrate the efficiencies of the stellar and AGN feedback to reproduce galaxy stellar mass function and the observed gas fractions of galaxy groups and clusters at redshift $z \approx 0$. We utilize three different variations of the BAHAMAS simulations that differ only in the strength of their AGN feedback prescription. This strength is determined by the AGN subgrid heating temperature, which is the temperature increase given to gas particles associated with each AGN feedback event.

Additionally, as a ‘proof of concept’, we test our model against two additional simulations, one that varies the mass resolution, but is otherwise identical to the fiducial BAHAMAS model, and one that changes the cosmology by including the effects of massive neutrinos,

both on the background expansion rate and the growth of density fluctuations (see further discussion in Section 5). For the high-resolution study, we used the ‘high res’ model presented in McCarthy et al. (2017), which left the subgrid feedback parameters unchanged with respect to the BAHAMAS fiducial resolution model, but it has 8 times better mass resolution (comoving volume of $100 \text{ cMpc } h^{-1}$ on a side and 2×512^3 particles). This represents a ‘strong’ convergence test in the terminology of Schaye et al. (2015). We note that only the particle data at redshifts $z = 0$ and $z = 0.125$ for these runs is still available. Hence, we limited our comparison for the ‘high res’ run to redshift $z = 0.125$.

We have also used our model to predict the baryonic suppression for the most extreme massive neutrino variation of the WMAP9 cosmology in the BAHAMAS suite (Mummery et al. 2017; McCarthy et al. 2018). This model left the subgrid feedback parameters unchanged with respect to the fiducial BAHAMAS model, but uses a total summed neutrino mass of $M_{\nu} = 0.48 \text{ eV}$. For this model, the amplitude of the density fluctuations at the epoch of recombination A_s , as inferred by WMAP9 data assuming massless neutrinos, was held fixed in order to retain agreement with the observed CMB angular power spectrum, and the value of σ_8 was

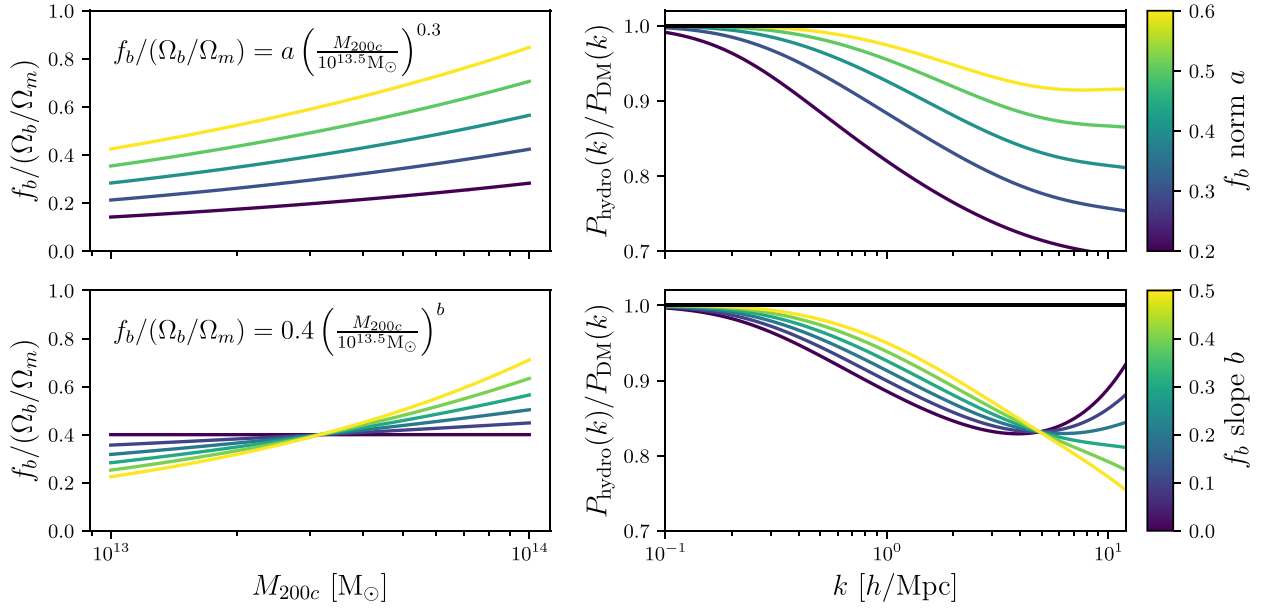


Figure 10. Top panel: Effect of variations in the power-law normalization (a) of the f_b – M_{halo} relation (left), as parametrized by equation (11), on the shape of the suppression of the total matter power spectrum (right) at redshift $z = 0.125$. The larger the amount of mass removed from haloes (lower f_b), the stronger the suppression on $P(k)$ relative to a baryon complete (DM-only) model. Bottom panel: Effect of variations in the power-law slope (b).

adjusted accordingly. Additionally, the matter density parameter of cold dark matter was adjusted to maintain a flat universe model. All other cosmological parameters are held fixed.

In Fig. 11, we show the impact of baryons on the total matter power spectrum for the three BAHAMAS models (Low AGN, fiducial, and High AGN), the high-resolution (HiRes) and massive neutrino ($M_\nu = 0.48$ eV) variations. In order to compute the baryon fraction from the simulations to use in equation (5), we used a piecewise cubic Hermite interpolating polynomial using narrow bins of M_{200c} with 0.1 dex width. This approach provides a high degree of flexibility, as simulations show evidence for a slightly mass-dependent slope (Farahi et al. 2018). Additionally, we show in dotted lines the results using power-law fits to the median baryon fraction in the halo mass range of interest at each redshift.

The figure shows that our model follows the simulation response typically to within 1 percent, with a maximum error of a ~ 3 per cent across the entire range of scales and redshifts shown. The errors at all scales and redshifts are consistent with the expected errors from the mean error contours in Fig. 9. The largest errors are at $2 \lesssim k [h \text{ Mpc}^{-1}] \lesssim 5$ at $z = 1$ for most models. In particular, for the lowest baryon fraction model, as expected from the maximum error contours in Fig. 9. For the massive neutrino cosmology, the largest errors are at $2 \lesssim k [h \text{ Mpc}^{-1}] \lesssim 5$ at $z = 0.5$ and $z = 0.1$. By construction, the model recovers the suppression (or rather, lack thereof) at the large scales. Hence, the errors are very small at the largest of scales. Finally, we do not find a significant decrease of the performance of the model when using a simple power law, compared to a more precise interpolation of the baryon fraction as a function of halo mass.

5 DISCUSSION

It is important to note that our model was built from a suite of simulations that widely varied the feedback parameters but only within a single cosmology (i.e. the cosmological parameters were

fixed). How worried should we be about possible degeneracies between astrophysics in cosmology? In other words, does the scaling derived here apply to other cosmologies? Our previous work based on the BAHAMAS simulations has shown that the effects of baryon physics on the matter distribution appear to be separable from (independent of) variations in the cosmological parameters, including extensions to massive neutrinos (Mummery et al. 2017), dynamical dark energy (Pfeifer et al. 2020), and a running of scalar spectral index (Stafford et al. 2020). In those studies, we demonstrated that the resulting cluster baryon fractions and the suppression of the matter power spectrum match those of our fiducial cosmology model to typically per cent level accuracy or better. This is consistent with the findings of previous work based on the baryonification formalism, which has also shown that the suppression of the power spectrum due to baryons is independent of variations in the parameters of the standard cosmological model to approximately the per cent level, with the exception of a clear dependence on the universal baryon fraction, Ω_b/Ω_m (Schneider et al. 2020; Aricò et al. 2021).

The dependence on the universal baryon fraction is straightforward to interpret: in the absence of feedback, groups and clusters are expected to contain a baryon fraction that reflects that of the universe as a whole (White et al. 1993), thus a higher (lower) universal fraction requires more (less) efficient feedback to obtain a particular (e.g. observed) baryon fraction in groups/clusters. In other words, if the baryon fractions of groups/clusters are used as an independent constraint on the feedback-induced suppression of the matter power spectrum, one must also specify the universal baryon fraction that is assumed. This is why our model takes as its independent variable the halo baryon fraction normalized by the universal baryon fraction. In practice, the quantity Ω_b/Ω_m is precisely pinned down by observations of the CMB, such that the uncertainties in baryon fractions of haloes are dominated by the statistical and systematic errors in the halo (total and baryon) mass measurements. Nevertheless, marginalizing over the uncertainties

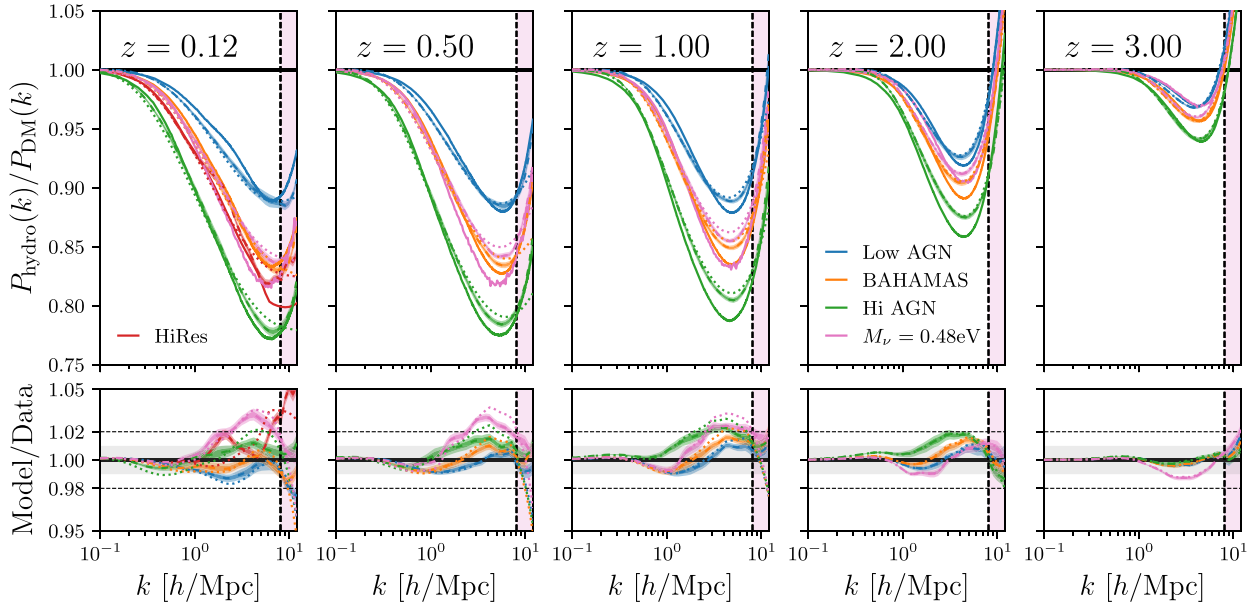


Figure 11. Fractional impact of baryons on the total matter power spectrum for three BAHAMAS models (Low AGN, fiducial, and High AGN; McCarthy et al. 2017, 2018), the high-resolution (HiRes McCarthy et al. 2017), and massive neutrino ($M_\nu = 0.48$ eV; Mummery et al. 2017; McCarthy et al. 2018) variations. Solid lines are the power spectra computed the hydrodynamical simulations. Dashed lines represent the best-fitting model obtained with $\mathcal{SP}(k)$ using binned data for the f_b – M_{halo} relation. Dotted lines use a simple power-law fit to the total baryon fraction. The light shaded regions enclose the 68 per cent and 95 per cent confidence interval from statistical errors (only shown for binned data). The vertical dashed line indicates the Nyquist frequency of the ANTILLES simulations k_{Ny} . The pink lightly shaded region indicates the scales where our model is not a good indicator of the uncertainty as $k > k_{\text{Ny}}$. Bottom panel: Ratio between the measurements of the suppression in the power spectrum induced by baryons as measured in the hydrodynamical simulations to our best-fitting model. The grey band highlights 1 per cent accuracy, while dashed lines indicate 2 per cent accuracy.

in both the halo and universal baryon fractions is recommended. Going forward, it will be important to continue to check for potential degeneracies between baryon and cosmological physics as we extend beyond the standard model of cosmology and also explore other cosmological observables (e.g. galaxy clustering, cluster counts, and so on).

Another caveat of our study is that, while we have explored a wide range of the feedback parameter volume in this study, this has only been done within the context of a given framework (or parametrization) for feedback. It will be important for other groups to independently test and perhaps refine the predictions for the impact of baryon physics on the matter power spectrum as new simulations become available. The study of van Daalen et al. (2020), which compared a limited number of simulations in the literature but which spanned a wide range of resolutions, hydro solvers, and feedback frameworks, is very suggestive that the physics of $P(k)$ is mainly driven by the baryon fractions. However, these simulations may not explore the full range of possibilities (see e.g. Debackere et al. 2020; Amon & Efstathiou 2022), so it is important to continue exploring the effects of baryons using a wide variety of methods, including the halo model, the baryonification formalism, and different hydrodynamical simulations.

As an initial ‘proof of concept’, we test the separability of feedback parameters (and resolution effects) from the cosmological model in Fig. 11 using a higher resolution simulation and one that changes the cosmology by including the effects of massive neutrinos. While this is a limited test, we find that our model based solely on the baryon fraction recovers the suppression of the matter power spectrum to a typically better than ~ 2 per cent accuracy. Further investigation including different feedback prescriptions, cosmological models, and resolution effects will be thoroughly explored in a follow-up paper.

6 SUMMARY AND CONCLUSIONS

We have presented ANTILLES, a new suite of 400 hydrodynamical simulations that explores the ‘feedback landscape’ associated with baryon physics. We have shown that a relatively simple yet remarkably accurate model can be constructed for the suppression of the matter power spectrum using only the mean baryon fraction of haloes (specifically the baryon fraction for haloes that dominate the power spectrum at a given scale, or wavenumber). Our work follows on from the recent study of van Daalen et al. (2020), by expanding greatly the number of simulations used to map this relationship and by pushing to a wider range of scales and redshifts, which are requirements for current and upcoming LSS surveys including cosmic shear.

The main specific findings of our study may be summarized as follows:

- (i) In agreement with van Daalen et al. (2020), we find that the fractional impact of baryon physics on the present-day non-linear matter power spectrum up to $k \sim 1 \, h \, \text{Mpc}^{-1}$ correlates very strongly with the baryon fraction of haloes with total masses of $\sim 10^{14} \, M_\odot$ (see Fig. 2). At smaller scales, however, the relation weakens somewhat, as also found by van Daalen et al. (2020).
- (ii) The weakening in the relation between the power spectrum suppression and the baryon fraction at a mass scale of $\sim 10^{14} \, M_\odot$ is a result of lower mass haloes contributing more significantly at smaller scales (see also Daalen & Schaye 2015; Mead et al. 2020). Using the simulations, we have empirically determined the halo mass scale whose baryon fraction correlates most strongly with the suppression of the power spectrum as a function of both scale and redshift (see Figs 4 and 5). We refer to this mass scale as the ‘optimal’ mass.
- (iii) Our $P(k)$ suppression model, called $\mathcal{SP}(k)$, is constructed by fitting a simple parametric form to the power spectrum suppression

as a function of scale (k), redshift (z), and the normalized baryon fraction at the optimal scale. See equation (4).

(iv) We characterize the error in our best-fitting model relative to our simulation training set in Figs 8 and 9, showing our model to be accurate to typically better than ≈ 2 per cent per cent accuracy over the range $0.1 \leq k [h \text{ Mpc}^{-1}] \leq 10$ and from $z = 0.1-3$.

(v) We tested our model against an independent set of BAHAMAS simulations (McCarthy et al. 2017, 2018), including a high-resolution box (McCarthy et al. 2017) and massive neutrino cosmology ($M_\nu = 0.48$ eV, Mummery et al. 2017; McCarthy et al. 2018), showing a similar level of accuracy (see Fig. 11).

(vi) We make a PYTHON implementation of our model publicly available at <https://github.com/jemme07/pyspk>.

With a fast, accurate model for characterizing the effects of baryon physics on the non-linear matter power spectrum, it is possible to straightforwardly incorporate these effects in existing theoretical pipelines based on gravity-only calculations (e.g. emulators for the absolute non-linear power spectrum or the halo model). This will allow one to consistently propagate the uncertainties in astrophysics (due to our uncertainties in the baryon fractions of groups/clusters as a function of mass and redshift) through to the cosmological constraints. An advantage that our model has over existing methods based on the halo model or the ‘baryonification’ formalism is that it effectively depends only on a single physically meaningful parameter (the baryon fraction). The benefit of this is that observational constraints on the baryon fraction could be used to inform the priors used in cosmological analysis, which in turn should help minimize the degradation of cosmological constraints that results from marginalizing over baryon effects (e.g. Amon & Efstathiou 2022). Furthermore, by providing a set of analytical equations, the model can be easily ‘inverted’ and allows for rapid experiments to be conducted, providing a powerful tool to explore the differential effects of baryonic physics.

Finally, as our model is expressed in terms of the baryon fraction of groups and clusters, we require observational constraints to either inform the priors that will be used in cosmological pipelines or to be used directly in a joint analysis. While the baryon fractions are pinned down reasonably precisely in the low-redshift universe (perhaps $z \lesssim 0.2$), there are currently very few constraints at higher redshifts for haloes with $M \lesssim 10^{14} M_\odot$. The situation may soon change, though, as eROSITA (X-ray), Advanced ACT (tSZ, kSZ), and Simons Observatory (tSZ, kSZ) data of large numbers of high- z groups starts to become available. An important consideration in these analyses will be carefully modelling the selection function to enable an unbiased measurement of the group baryon fractions. Analyses of mock X-ray and SZ observations of cosmological hydrodynamical simulations will be important in this endeavour.

ACKNOWLEDGEMENTS

The authors would like to thank Dr Alexandra Amon for her thorough report and valuable comments as the referee, which greatly contributed to the improvement of the paper. 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). This work was supported by the Science and Technology Facilities Council (STFC) grant ST/P000541/1. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1,

ST/P002293/1, and ST/R002371/1, 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

- Acuto A., McCarthy I. G., Kwan J., Salcido J., Stafford S. G., Font A. S., 2021, *MNRAS*, 508, 3519
- Akino D. et al., 2022, *PASJ*, 74, 175
- Amon A. et al., 2022, *Phys. Rev. D*, 105, 023514
- Amon A., Efstathiou G., 2022, *MNRAS*, 516, 5355
- Angulo R. E., Zennaro M., Contreras S., Aricò G., Pellejero-Ibañez M., Stücker J., 2021, *MNRAS*, 507, 5869
- Aricò G., Angulo R. E., Contreras S., Ondaro-Mallea L., Pellejero-Ibañez M., Zennaro M., 2021, *MNRAS*, 506, 4070
- Aricò G., Angulo R. E., Zennaro M., Contreras S., Chen A., Hernández-Monteagudo C., 2023, preprint ([arXiv:2303.05537](https://arxiv.org/abs/2303.05537))
- Asgari M. et al., 2021, *A&A*, 645, A104
- Baldry I. K. et al., 2012, *MNRAS*, 421, 621
- Balogh M. L., McCarthy I. G., Bower R. G., Eke V. R., 2008, *MNRAS*, 385, 1003
- Bernardi M., Meert A., Sheth R. K., Vikram V., Huertas-Company M., Mei S., Shankar F., 2013, *MNRAS*, 436, 697
- Bond J. R., Efstathiou G., Silk J., 1980, *Phys. Rev. Lett.*, 45, 1980
- Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
- Castro T., Borgani S., Dolag K., Marra V., Quartin M., Saro A., Sefusatti E., 2021, *MNRAS*, 500, 2316
- Chen A. et al., 2023, *MNRAS*, 518, 5340
- Chisari N. E. et al., 2018, *MNRAS*, 480, 3962
- Chisari N. E. et al., 2019, *Open J. Astrophys.*, 2, 4
- Cooray A., Sheth R., 2002, *Phys. Rep.*, 372, 1
- Cullen L., Dehnen W., 2010, *MNRAS*, 408, 669
- Dalla Vecchia C., Schaye J., 2008, *MNRAS*, 387, 1431
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- Debackere S. N. B., Schaye J., Hoekstra H., 2020, *MNRAS*, 492, 2285
- Debackere S. N. B., Schaye J., Hoekstra H., 2021, *MNRAS*, 505, 593
- DeRose J. et al., 2019, *ApJ*, 875, 69
- Deutsch J. L., Deutsch C. V., 2012, *J. Stat. Plan. Inference*, 142, 763
- Driver S. P. et al., 2022, *MNRAS*, 513, 439
- Durier F., Dalla Vecchia C., 2012, *MNRAS*, 419, 465
- Eifler T., Krause E., Dodelson S., Zentner A. R., Hearin A. P., Gnedin N. Y., 2015, *MNRAS*, 454, 2451
- Euclid Collaboration, 2019, *MNRAS*, 484, 5509
- Farahi A., Evrard A. E., McCarthy I., Barnes D. J., Kay S. T., 2018, *MNRAS*, 478, 2618
- Giri Sambit K., Schneider Aurel, 2021, *Journal of Cosmology and Astroparticle Physics*, 12, 046
- Gonzalez A. H., Sivanandam S., Zabludoff A. I., Zaritsky D., 2013, *ApJ*, 778, 14
- Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V. eds, *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*. astro-ph/0106018
- Heitmann K. et al., 2016, *ApJ*, 820, 108
- Heitmann K., Lawrence E., Kwan J., Habib S., Higdon D., 2014, *ApJ*, 780, 111
- Henden N. A., Puchwein E., Shen S., Sijacki D., 2018, *MNRAS*, 479, 5385
- Henden N. A., Puchwein E., Sijacki D., 2020, *MNRAS*, 498, 2114
- Heymans C. et al., 2021, *A&A*, 646, A140
- Hinshaw G. et al., 2013, *ApJ*, 208, 19
- Hopkins P. F., 2013, *MNRAS*, 428, 2840
- Kaiser N., 1987, *MNRAS*, 227, 1
- Krause E. et al., 2021, preprint ([arXiv:2105.13548](https://arxiv.org/abs/2105.13548))

- This paper has been typeset from a T_EX/L^AT_EX file prepared by the author.