

The impact of galactic feedback on the shapes of dark matter haloes

Kun Ting Eddie Chua     Mark Vogelsberger   Annalisa Pillepich   and Lars Hernquist⁴

¹*Institute of High Performance Computing, 1 Fusionopolis Way, 138632, Singapore*

²*Department of Physics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA*

³*Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*

⁴*Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA*

Accepted 2022 July 2. Received 2022 July 2; in original form 2021 August 26

ABSTRACT

We quantify the impact of galaxy formation on dark matter halo shapes using cosmological simulations at redshift $z = 0$. Using magnetohydrodynamic simulations from the IllustrisTNG project, we focus on haloes of mass $10^{10\text{--}14} M_{\odot}$ from the 50 Mpc (TNG50) and 100 Mpc (TNG100) boxes and compare them to dark matter-only (DMO) analogues and other simulations, e.g. Numerical Investigation of a Hundred Astrophysical Objects (NIHAO) and Evolution and Assembly of GaLaxies and their Environments (EAGLE). We further quantify the prediction uncertainty by varying the feedback models using smaller 25 Mpc h^{-1} boxes. We find that (i) galaxy formation results in rounder haloes compared to DMO simulations, in qualitative agreement with past results. Haloes of mass $\approx 2 \times 10^{12} M_{\odot}$ are most spherical, with an average minor-to-major axial ratio of $\langle s \rangle \approx 0.75$ in the inner halo, an increase of 40 per cent compared to their DMO counterparts. No significant difference is present for low-mass $10^{10} M_{\odot}$ haloes; (ii) stronger feedback, e.g. increasing galactic wind speed, reduces the impact of baryons; (iii) the inner halo shape correlates with the stellar mass fraction, explaining the dependence of halo shapes on feedback models; and (iv) the fiducial and weaker feedback models are most consistent with observational estimates of the Milky Way halo shape. At fixed halo mass, very diverse and possibly unrealistic feedback models all predict inner shapes closer to one another than to the DMO results. Because of the large halo-to-halo variation in halo shape, a larger observational sample is required to statistically distinguish different baryonic prescriptions.

Key words: methods: numerical – methods: statistical – galaxies: formation – galaxies: haloes – dark matter.

1 INTRODUCTION

The rise in computational power has led to an increase in the scope of hydrodynamic cosmological simulations. By implementing galaxy formation physics, these simulations aim to model the formation of realistic galaxies in a cosmological context (for a review, see Vogelsberger et al. 2020a). In addition to their effects on the stellar and gaseous content of galaxies, baryonic processes such as radiative cooling, star formation, and stellar and active galactic nuclei (AGN) feedback can also have significant effects on the structure of their host dark matter (DM) haloes and subhaloes.

In particular, theoretical predictions for halo shapes have been studied extensively using cosmological simulations. Hierarchical structure formation predicts anisotropic halo growth, since the accretion of matter on to DM haloes during their growth is preferentially aligned with DM sheets and filaments. This prediction is well supported by a wealth of N -body simulations (Dubinski & Carlberg 1991; Warren et al. 1992; Bullock 2002; Jing & Suto 2002; Bailin & Steinmetz 2005; Allgood et al. 2006; Macciò, Dutton & van den Bosch 2008; Vera-Ciro et al. 2011), which have shown that dark matter-only (DMO) haloes are triaxial and prolate ($c/b > b/a$).¹

In Milky Way (MW)-sized haloes ($\sim 10^{12} M_{\odot}$), DMO simulations predict a sphericity $c/a \approx 0.5\text{--}0.6$ near the galactic centre (within few tens of kpc). These results are, however, in contrast to the MW halo shape estimated using stellar kinematics and stellar streams, which suggest a more spherical inner halo e.g. $c/a \geq 0.8$ (Ibata et al. 2001), and $c/a = 0.72$ (Law & Majewski 2010).

The condensation of baryons to halo centres can have an impact on halo shapes since the formation of a central baryonic mass can scatter and modify the orbits of approaching DM particles. Such an effect was derived in Debattista et al. (2008), which simulated an isolated halo and found the growth of a central component can deform box orbits into rounder trajectories. Consequently, haloes from galaxy formation simulations have been found to be transformed into more spherical configurations than their analogues from DMO simulations (e.g. Katz & Gunn 1991; Katz & White 1993; Dubinski 1994; Abadi et al. 2010; Kazantzidis, Abadi & Navarro 2010; Tissera et al. 2010; Bryan et al. 2013; Butsky et al. 2016; Chisari et al. 2017; Chua et al. 2019). This has also led to better agreement between halo shapes in baryonic simulations and the observationally inferred DM halo shape of our Galaxy (e.g. Chua et al. 2019; Prada et al. 2019). More recently, Cataldi et al. (2021) further concluded using the Fenix and Evolution and Assembly of GaLaxies and their Environments (EAGLE) simulations that DM halo shapes are related to the morphology of the central galaxies, pointing to a fundamental relation between the growth of the central galaxy and its parent halo.

* E-mail: eddie@cheleb.com

¹We denote the major, intermediate, and minor axes semilengths as a , b , and c , respectively, with $a > b > c$.

Despite the role of galaxy formation models in resolving disagreements between N -body predictions and observations, as well as in reproducing stellar galactic properties, significant uncertainty regarding feedback implementations in simulations remains. Such variations arise because many baryonic processes such as star formation and black hole accretion are not directly resolved in most galaxy formation simulations. Instead, these effects are described through subgrid models, where multiple parametrizations often exist. In many cases, large galaxy formation simulations rely on a single fiducial galaxy formation model, where parameters and parametrizations are tuned or chosen to reproduce a set of observations. As such, many numerical studies relying on these simulations do not take into account inherent uncertainty in baryonic feedback, which can have important consequences on the predicted properties of the galaxy and its surrounding medium (e.g. Suresh et al. 2015).

In the absence of outflows and star formation, the gaseous disc that forms at the centres of simulated haloes is unrealistically small and massive, but provides an idea of the maximal effect of galaxy assembly on the DM haloes. For example, such a configuration was investigated by Abadi et al. (2010), who estimated such haloes to be completely oblate with axial ratios of the isopotential contours to be $b/a \sim 1$ and $c/a \sim 0.85$. Newer cosmological simulations allow realistic galaxies to form in a cosmological context and enable a more accurate characterization of halo shapes (e.g. Butsky et al. 2016; Chua et al. 2019; Prada et al. 2019; Emami et al. 2021). Although many of these galaxy simulations qualitatively agree on the sphericalization effects of baryons, differing methods in measuring halo shapes often render quantitative comparisons difficult (Zemp et al. 2011).

Furthermore, due to the high computational costs of galaxy formation simulations, only one ‘fiducial’ galaxy formation model is often considered. As such, few studies exist looking into the uncertainty in halo shape predictions arising from the strength of feedback prescriptions. Earlier work by Bryan et al. (2013) using the OverWhelmingly Large Simulations (OWLS; Schaye et al. 2010) showed that varying the stellar and AGN feedback can significantly change the inner region shape of DM haloes. However, the non-iterative method used to measure halo shape is considered less accurate compared to other iterative methods (Zemp et al. 2011).

In this paper, we study the result of realistic galactic feedback on DM halo shapes, and further investigate halo shapes when subjected to a variety of baryonic feedback prescriptions. Our fiducial results are obtained using haloes from IllustrisTNG² (hereafter TNG; Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018b; Springel et al. 2018), a suite of cosmological magnetohydrodynamical simulations carried out in various box sizes. The IllustrisTNG simulations were conceived to improve upon the original Illustris simulation (Vogelsberger et al. 2013, 2014a,b; Genel et al. 2014), which was found to exhibit tensions with observations in terms of (1) a cosmic star formation rate density at $z < 1$ that was too high; (2) a stellar mass function and stellar mass fraction at $z = 0$ that were too high for both high-mass and low-mass galaxies; (3) $z = 0$ galaxy sizes (stellar half-mass radii) that were too large; and (4) halo X-ray emission from high-mass galaxies severely underestimating that seen in observations. The updated TNG galaxy formation model alleviates the above-mentioned tensions while continuing to demonstrate good agreement with observational constraints (see Nelson et al. 2019a, for a partial list). Through radiative transfer post-processing, the suite of IllustrisTNG simulations has also provided

detailed predictions of the high-redshift galaxy populations expected to be observed by the *James Webb Space Telescope* (Shen et al. 2020, 2022; Vogelsberger et al. 2020b).

With a higher resolution compared to TNG100, TNG50 resolves shapes of smaller haloes down to $10^{10} M_{\odot}$. As such, we are able to statistically characterize halo shapes of diverse masses, ranging from $10^{10} M_{\odot}$ dwarf haloes to $5 \times 10^{14} M_{\odot}$ clusters. The similar box size, element count, and initial condition of TNG100 and the older Illustris simulation (previously analysed in Chua et al. 2019) enable a direct comparison of both galaxy formation models.

In addition to the flagship IllustrisTNG runs, we analyse an additional set of simulations that include variations of the TNG feedback model. While these simulations are carried out in smaller boxes of side length $25 h^{-1}$ Mpc, they provide important insights into the relationship between the feedback strength and halo shapes. By making specific variations to feedback parameters of the galactic wind and AGN model, we can further and systematically quantify the impact of baryonic physics on DM halo shapes.

This paper is structured as follows. We describe our simulation methods and definitions in Section 2. We present the main results on halo shapes from TNG100 and TNG50 in Section 3, and compare them to those from other hydrodynamic simulations. In Section 4, we analyse the impact of feedback variations using the smaller boxes and present a comparison to observational estimates of the MW halo shape. Finally, our summary and conclusions are presented in Section 5.

2 METHODS AND DEFINITIONS

2.1 IllustrisTNG simulations

The haloes we study in this work are drawn from *The Next Generation Illustris Simulations* (IllustrisTNG), a suite of cosmological magnetohydrodynamical simulations (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018b; Springel et al. 2018). The simulations are performed using the simulation code AREPO (Springel 2010), which calculates the gravitational forces using a Tree-Particle-Mesh method and solves the ideal magnetohydrodynamic (MHD) equations using a finite volume method on an adaptive mesh. The cosmology utilized in IllustrisTNG is consistent with that of *Planck*, given by $\Omega_m = 0.3089$, $\Omega_{\Lambda} = 0.6911$, $\Omega_b = 0.0486$, $\sigma_8 = 0.8159$, $n_s = 0.9667$, and $h = 0.6774$ (Planck Collaboration XVI 2014; Spergel, Flauger & Hložek 2015).

The galaxy formation model adopted in TNG accounts for (i) primordial and metal-line gas cooling; (ii) a spatially uniform and time-dependent ultraviolet (UV) background; (iii) stellar formation and feedback; and (iv) kinetic and thermal feedback from black holes. The TNG galaxy formation model builds upon the physics model introduced previously in Illustris to address the shortcomings that were identified. Compared to Illustris, these improvements include (i) an updated kinetic AGN feedback model at low accretion rates; (ii) improved isotropic galactic winds; and (iii) ideal MHD. For specific details, we refer the reader to Weinberger et al. (2017) and Pillepich et al. (2018a). The model parameters of IllustrisTNG are that of the default (*fiducial*) model described in Pillepich et al. (2018a), selected to produce good agreements between simulated galaxies and observations of the cosmic star formation rate density and the stellar content of the galaxy population at $z = 0$.

In this paper, we analyse halo shapes from both TNG50 (Nelson et al. 2019b; Pillepich et al. 2019) and TNG100, focusing primarily on the highest resolution simulation of each box size. TNG50 has a volume of $(50 \text{ Mpc})^3$, with a DM mass resolution of 4.5×10^5

²www.tng-project.org

Table 1. Summary of the main simulations and their resolution parameters examined in this work: (1) simulation name; (2) the type of simulation; (3) length of simulation box; (4) number of cells and particles in the simulation; (5) mass per dark matter (DM) particle; (6) target mass of baryonic particles; and (7) Plummer-equivalent gravitational softening lengths at redshift $z = 0$. In the TNG runs, the softening lengths for all particle types are comoving kpc for $z > 1$, after which they are fixed to their $z = 1$ values in physical space. Note that in Illustris, this procedure is not applied to the DM particles, thus DM particles have twice the softening lengths (first value) as the stellar particles (second value). The MHD simulations are based on the fiducial TNG galaxy formation model described in Pillepich et al. (2018a). The nine model variations in the small boxes are further listed in Table 2.

| Name | Type | Box size (Mpc) | DM particles and cells | m_{DM} ($10^6 M_{\odot}$) | m_{baryon} ($10^6 M_{\odot}$) | $\epsilon_{\text{DM,stars}}^{z=0}$ (kpc) |
|--------------------------|--------------------------|-------------------------|------------------------|--------------------------------------|--|--|
| TNG50 | MHD | $35 h^{-1} \approx 52$ | 2×2160^3 | 0.454 | 0.085 | 0.288 |
| TNG50-DM | DMO | $35 h^{-1}$ | 2160^3 | 0.54 | – | 0.288 |
| TNG100 | MHD | $75 h^{-1} \approx 110$ | 2×1820^3 | 7.5 | 1.4 | 0.74 |
| TNG100-DM | DMO | $75 h^{-1}$ | 1820^3 | 8.9 | – | 0.74 |
| Illustris | Hydrodynamic | $75 h^{-1}$ | 2×1820^3 | 6.3 | 1.3 | 1.42/0.71 |
| Illustris-Dark | DMO | $75 h^{-1}$ | 1820^3 | 7.52 | – | 1.42/– |
| L25n512 (Small boxes) | Various (See Table 2) | $25 h^{-1} \approx 37$ | 2×512^3 | 12.4 | 2.4 | 0.74 |

and baryonic mass resolution of $m_b = 8.5 \times 10^4 M_{\odot}$. TNG100 has a volume of $(100 \text{ Mpc})^3$, with a DM mass resolution of 7.5×10^6 and baryonic mass resolution of $m_b = 1.4 \times 10^6 M_{\odot}$. The lower resolution run TNG100-2 uses 2×910^3 elements with two times worse spatial resolution, while TNG100-3 uses 2×455^3 elements with four times worse spatial resolution compared to TNG100. For each box size, DMO simulations are performed, providing baseline N -body results for comparison. The important parameters of the simulations are summarized in Table 1. In general, both the MHD and hydrodynamics simulations will be referred to as the *Full-Physics* runs.

2.2 TNG model and variations

Besides the primary 50 and 100 Mpc boxes, variations of the fiducial TNG model were also simulated in smaller boxes of size $L = 25 h^{-1} \sim 37 \text{ Mpc}$, with 2×512^3 resolution elements. As shown in Table 1, these L25n512 small boxes have similar softening lengths and resolutions (within a factor of 2) compared to TNG100. All simulations of model variations were carried out with the same initial conditions, which were chosen by conducting low-resolution DMO simulations from 10 different initial density fields and then choosing the realization with the DM halo mass function closest to the average (Pillepich et al. 2018a). This was done to reduce the effect of sample variance in this smaller box.

In this work, we focus on the effect of galactic winds and black holes. In the TNG model, galactic winds are driven by star formation, which launches wind particles with an initial speed v_w scaling with the local DM velocity dispersion σ_{DM} :

$$v_w = \max \left[\kappa_w \sigma_{\text{DM}} \left(\frac{H_0}{H(z)} \right)^{1/3}, v_{w,\min} \right], \quad (1)$$

where $H(z)$ is the redshift-dependent Hubble parameter, and $v_{w,\min}$ is a velocity floor. κ_w is thus a dimensionless factor that determines the speeds of launched winds, taken to be $\kappa_w = 7.4$ in the fiducial model. We investigate variations with doubled κ_w (*fast winds*) and with halved κ_w (*slow winds*).

For a given wind speed, mass loading of the wind depends on the energy available for wind generation, which is directly proportional to the free parameter \bar{e}_w in the TNG model. This quantity \bar{e}_w reflects the energy released per core-collapse supernova, taken to be $\bar{e}_w =$

Table 2. Variations of the galaxy formation model simulated in smaller boxes (L25n512). Each parameter or choice change is made with respect to the fiducial TNG model. Further description and usages of these model-variation runs can be found in Pillepich et al. (2018a).

| Variation type | Details |
|------------------------|---|
| Fiducial | $\bar{e}_w = 3.6, \kappa_w = 3.6$ |
| Strong winds | Doubled wind energy ($\bar{e}_w = 7.2$) |
| Weak winds | Halved wind energy ($\bar{e}_w = 1.8$) |
| Fast winds | Doubled wind speed ($\kappa_w = 14.8$) |
| Slow winds | Halved wind speed ($\kappa_w = 3.2$) |
| No winds | No galactic winds |
| No BH | No black holes |
| No BH kinetic mode | Only BH quasar mode |
| Dark matter-only (DMO) | No baryons |

3.6 in the fiducial case. We examine the effects of increasing the wind energy, with this factor doubled to $\bar{e}_w = 7.2$ (*strong winds*). Additionally, a case without galactic winds (*no winds*) is also carried out.

In the TNG model, black hole feedback is driven by a combination of thermal feedback heating the surrounding gas at high accretion rates, and a kinetic AGN feedback model driving black hole-driven winds at low accretion rates (Weinberger et al. 2017). We investigate the effect of suppressing kinetic AGN feedback (*no BH kinetic winds*), which has been found to be important in quenching star formation in high-mass haloes. Finally, a case neglecting black hole feedback (*no BHs*) is also examined.

The nine model variations (including the fiducial and DMO models) are summarized in Table 2.

2.3 Identifying and matching haloes and subhaloes

In the simulations, haloes are identified using a Friends-of-Friends (FOF) group finder algorithm with a linking length of 0.2 (Davis et al. 1985). The SUBFIND algorithm subsequently identifies gravitationally self-bound subhaloes (Springel et al. 2001; Dolag et al. 2009). The subhalo in each FOF group with the lowest potential resolution element is classified as *central*, and it is typically the most massive: the remaining subhaloes are called *satellites*, whether they include

Table 3. Number of haloes at redshift $z = 0$ adopted in this paper and extracted from TNG100, TNG50, Illustris, and their DMO counterparts, according to mass. In the 50 Mpc boxes (TNG50 and TNG50-DM), we focus on haloes of mass 10^{10} – $10^{13} M_{\odot}$. In the other simulations, which have lower resolutions, the minimum halo mass we consider is $10^{11} M_{\odot}$, hence the number of haloes in the boxes is given in parentheses. For the smaller boxes (L25n512), the number of haloes corresponds to that of the fiducial model run.

| Simulation | Halo mass M_{200} (M_{\odot}) | | | | |
|--------------------------|-------------------------------------|--------------|--------------|--------------|--------------|
| | 10^{10-11} | 10^{11-12} | 10^{12-13} | 10^{13-14} | 10^{14-15} |
| TNG50 | 9179 | 1441 | 183 | (23) | (1) |
| TNG50-DM | 11 398 | 1537 | 193 | (24) | (1) |
| TNG100 | (87 025) | 12 963 | 1708 | 168 | 14 |
| TNG100-DM | (109 853) | 13 912 | 1743 | 198 | 13 |
| Illustris | (82 249) | 12 875 | 1317 | 109 | 10 |
| Illustris-Dark | (93 569) | 11 781 | 1440 | 160 | 11 |
| L25n512 (fiducial model) | (3274) | 496 | 77 | 9 | 0 |

stars or not. For each halo, the virial mass M_{200} and virial radius R_{200} are calculated.³

In this work, we examine haloes at the current redshift $z = 0$. Table 3 compares the number of haloes in four mass bins in Full-Physics and DMO runs. In TNG100, there are $\approx 14\,000$ haloes of mass greater than $10^{11} M_{\odot}$, including 14 cluster-sized ($\sim 10^{14}$) haloes. We focus on haloes of small to intermediate masses (10^{10} – $10^{13} M_{\odot}$) in TNG50, since its higher resolution allows the shapes of lower mass haloes to be resolved.

We also identify matching haloes between the Full-Physics and DMO runs to enable one-to-one halo shape comparisons between each halo pair. For any given subhalo in the Full-Physics run, the matching DMO subhalo is the one with the largest fraction of matching DM particles, identified using their unique IDs. The process is reversed, and only bidirectional matches are considered to be successful. At the masses considered in this work, almost all central subhaloes are successfully matched: this corresponds to 10 804 out of 10 827 haloes with $M_{200} \geq 10^{10} M_{\odot}$ in TNG50, and 14 778 out of 14 853 with $M_{200} \geq 10^{11} M_{\odot}$ in TNG100.

2.4 Defining the halo shape

The algorithm we used to quantify halo shapes is identical to that described in Chua et al. (2019), which we summarize as follows. Assuming that DM haloes are triaxial, their shapes are determined by the axis ratios of the isodensity surface, $q \equiv b/a$ and $s \equiv c/a$, where a , b , and c are the major, intermediate, and minor axes respectively. These parameters can be found using an iterative algorithm with the *unweighted shape tensor* (e.g. Bailin & Steinmetz 2005; Zemp et al. 2011),

$$S_{ij} = \frac{1}{\sum_k m_k} \sum_k m_k r_{k,i} r_{k,j}, \quad (2)$$

related to the second moment of the mass distribution. Here, m_k is the mass of the k th particle, and $r_{k,i}$ is the i th component of its position vector.

The iterative algorithm allows the shape of the integration volume to adapt to the shape of the halo. The first iteration begins with a

³ R_{Δ} is the radius within which the enclosed mass density is Δ times the critical value ρ_c , i.e. $\rho_{\text{halo}} = \Delta \rho_c$. M_{Δ} is the total mass of the halo enclosed within R_{Δ} . In this work, we use the value $\Delta = 200$.

spherical shell ($q = s = 1$), within which a set of particles is selected. We then calculate and diagonalize the shape tensor using equation (2). The eigenvectors denote the directions of the principal axes, while the eigenvalues are related to the square-roots of the principal axes lengths ($i \propto \sqrt{\lambda_i}$, $i = a, b, c$). In subsequent iterations, the new values of $q = b/a$ and $s = c/a$ are used to select a new set of particles, and the ellipsoidal shell is deformed keeping the semimajor length constant. The process continues until both q and s converge, i.e. when fractional differences in successive values differ by less than 1 per cent.

To obtain the shape profiles, we calculate the local shape $q(r)$ and $s(r)$ in ellipsoidal shells as a function of distance from the halo centre. The elliptical radius,

$$r_{\text{ell}}^2 = x^2 + \frac{y^2}{q^2} + \frac{z^2}{s^2}, \quad (3)$$

is used to determine if a given particle falls within an ellipsoidal shell with axis ratios q and s . Assuming that ellipsoids are oriented with x along the major axis and z along the minor axis, r_{ell} is also the semimajor length of the ellipsoid where the particle is residing. Throughout this paper, halocentric distances involving the halo shape will in general refer to this elliptical radius r_{ell} .

When using shell-enclosed particles, the density distribution and thus the shape tensor can be sensitive to the presence of large satellites (Zemp et al. 2011). We avoid substructure contamination by considering only particles bound to the central subhalo, as identified by SUBFIND. Using the unweighted shape tensor, we apply the iterative procedure to calculate the shape of each halo. To obtain a radial profile, 15 ellipsoidal shells logarithmically spaced in the range $0.01 \leq r/R_{200} \leq 1$ are used, each with a logarithmic width of $\Delta(r/R_{200}) = 0.1$ dex. Finally, the triaxiality parameter $T \equiv (1 - q^2)/(1 - s^2)$ measures how prolate ($T = 1$) or oblate ($T = 0$) the halo is. The axis ratios q and s , as well as the triaxiality T , are collectively termed the halo shape parameters.

3 IMPACT OF BARYONS IN TNG100 AND TNG50

We begin by focusing on the fiducial TNG galaxy formation model, which presents our best estimate of the impact of baryons on DM halo shapes. Here, we rely on the flagship runs TNG50 and TNG100 to take advantage of their large box sizes and high resolutions.

We visualize the shapes of some representative haloes of TNG100 and TNG50 in Figs 1 and 2, respectively. Each plot shows the DM density in a 30 kpc thick slice within half of the virial radius, for matching haloes in the MHD (left-hand column) and DMO (right-hand column) simulations. To better highlight the halo shapes, contour lines corresponding to three different densities have also been included. For each simulation, the haloes are ordered top to bottom, from less to more massive. Qualitatively, it is apparent that the MHD haloes are more spherical than their DMO counterparts. For example, the MHD haloes in the third and fourth rows of Fig. 2 (TNG50) appear almost completely spherical at these radii. However, we caution making conclusions based on projections of the DM density since the projected shape can depend on the projection angle, and the small sample size can be unrepresentative of the halo population at large.

Although not a focus of this paper, we point out that apart from transforming halo shapes through the axis ratios s and q or the triaxiality T , the tilt of the halo can also be changed by baryonic physics. This is most evident in the second row of Fig. 2, where the major axis is aligned closer to the vertical in TNG50 compared to

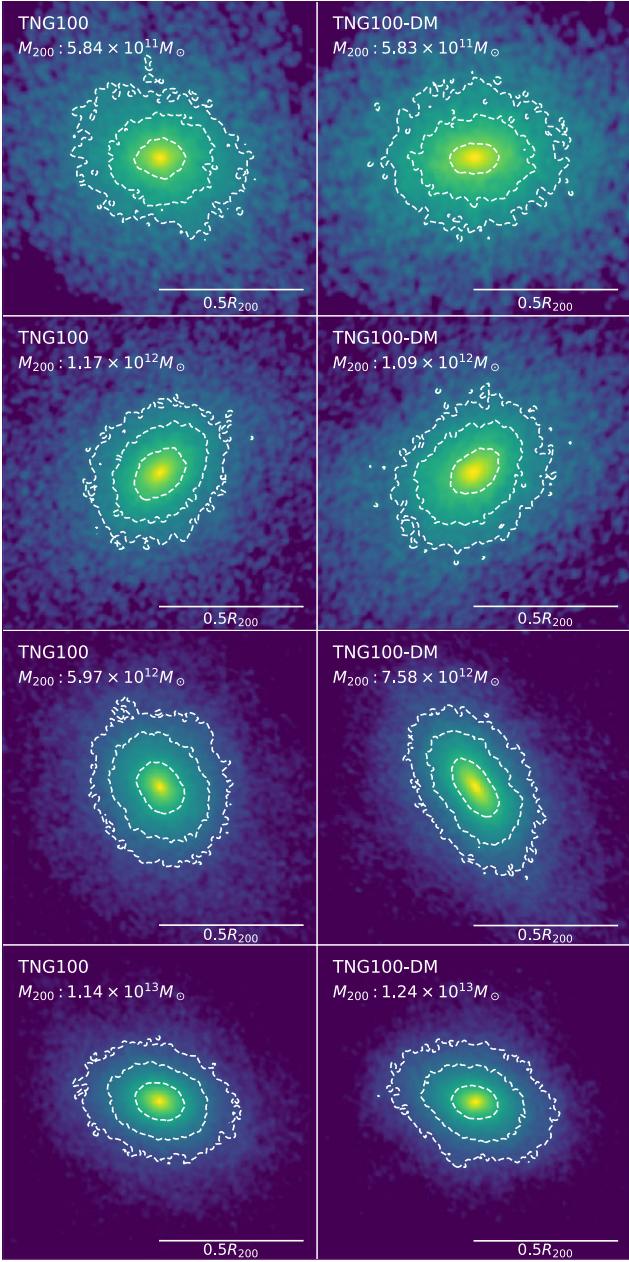


Figure 1. DM density on a logarithmic colour scale within a slice of thickness 30 kpc, for TNG100 and TNG100-DM. Four representative haloes in TNG100 (left) with their matched DMO counterparts in TNG100-DM (right) are shown, and the haloes are randomly projected on a 2D plane. Dashed lines indicate isodensity contours. These examples highlight the sphericalization of haloes when including galaxy formation physics in the simulations. Additionally, we see that halo shapes are not uniform, but vary with distance to the halo centre.

TNG50-DM, and can further impact the alignment of the galaxy with respect to the halo (e.g. Tenneti et al. 2015; Velliscig et al. 2015).

3.1 Determining resolved regions

Because of limitations in the mass and spatial resolutions, not all regions of a halo can be reliably resolved in the simulations. We determine the convergence radius r_{conv} , the smallest radius for which a halo shape is resolved, by comparing the median DMO profiles

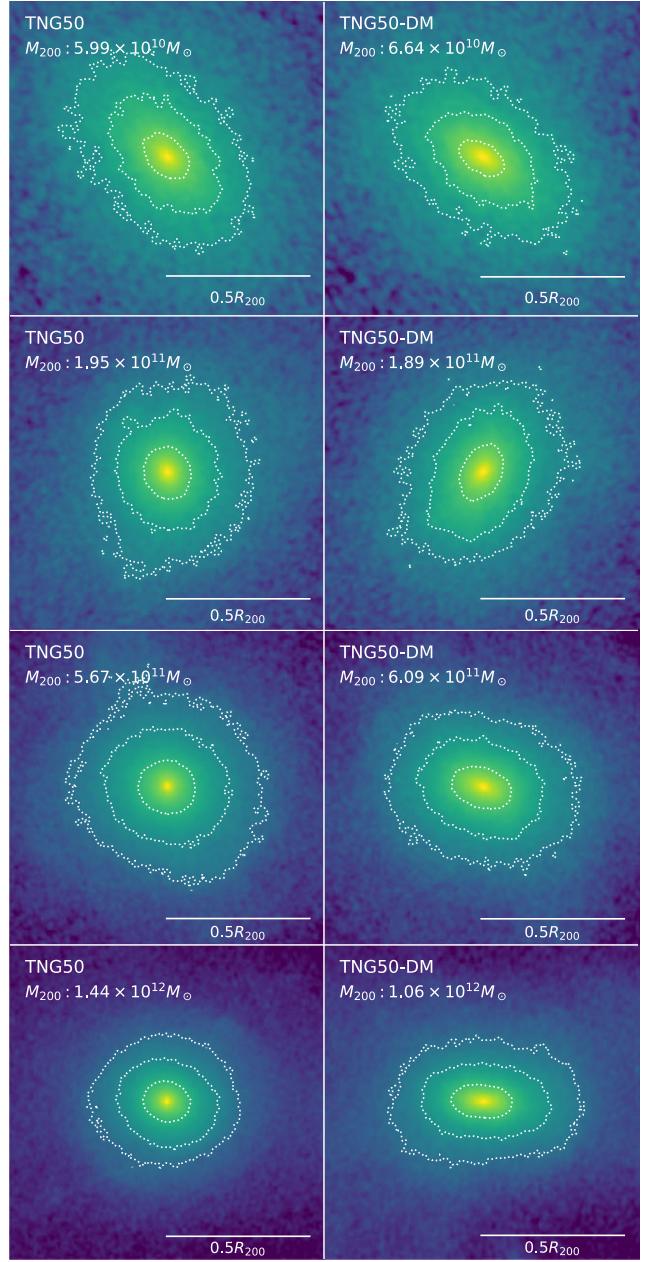


Figure 2. DM density on a logarithmic colour scale within a slice of thickness 30 kpc, for TNG50 and TNG50-DM. Compared to TNG100 (Fig. 1), the density and contours are smoother due to the higher mass resolution in TNG50.

at different resolutions (i.e. TNG100-DM with TNG100-DM-2 and TNG100-DM-3); see Appendix A. There we find that $\kappa(r_{\text{conv}}) = 7$ describes the convergence radius well, where

$$\kappa(r) \equiv \frac{\sqrt{200}}{8} \frac{N(r)}{\ln N(r)} \left[\frac{\bar{\rho}(r)}{\rho_{\text{crit}}} \right]^{-1/2}. \quad (4)$$

$N(r)$ is the number of DM particles enclosed within a radius r , $\bar{\rho}(r)$ is the mean density within r , and ρ_{crit} is the critical density of the universe. The expression for κ is derived from the ratio of the two-body relaxation time-scale to the circular orbit time-scale at the virial radius (Power et al. 2003; Vera-Ciro et al. 2011).

In TNG100-DM, the convergence radius is $r_{\text{conv}} \approx 10$ kpc, which corresponds to ≈ 15 per cent of the median virial radius for $10^{11} M_{\odot}$

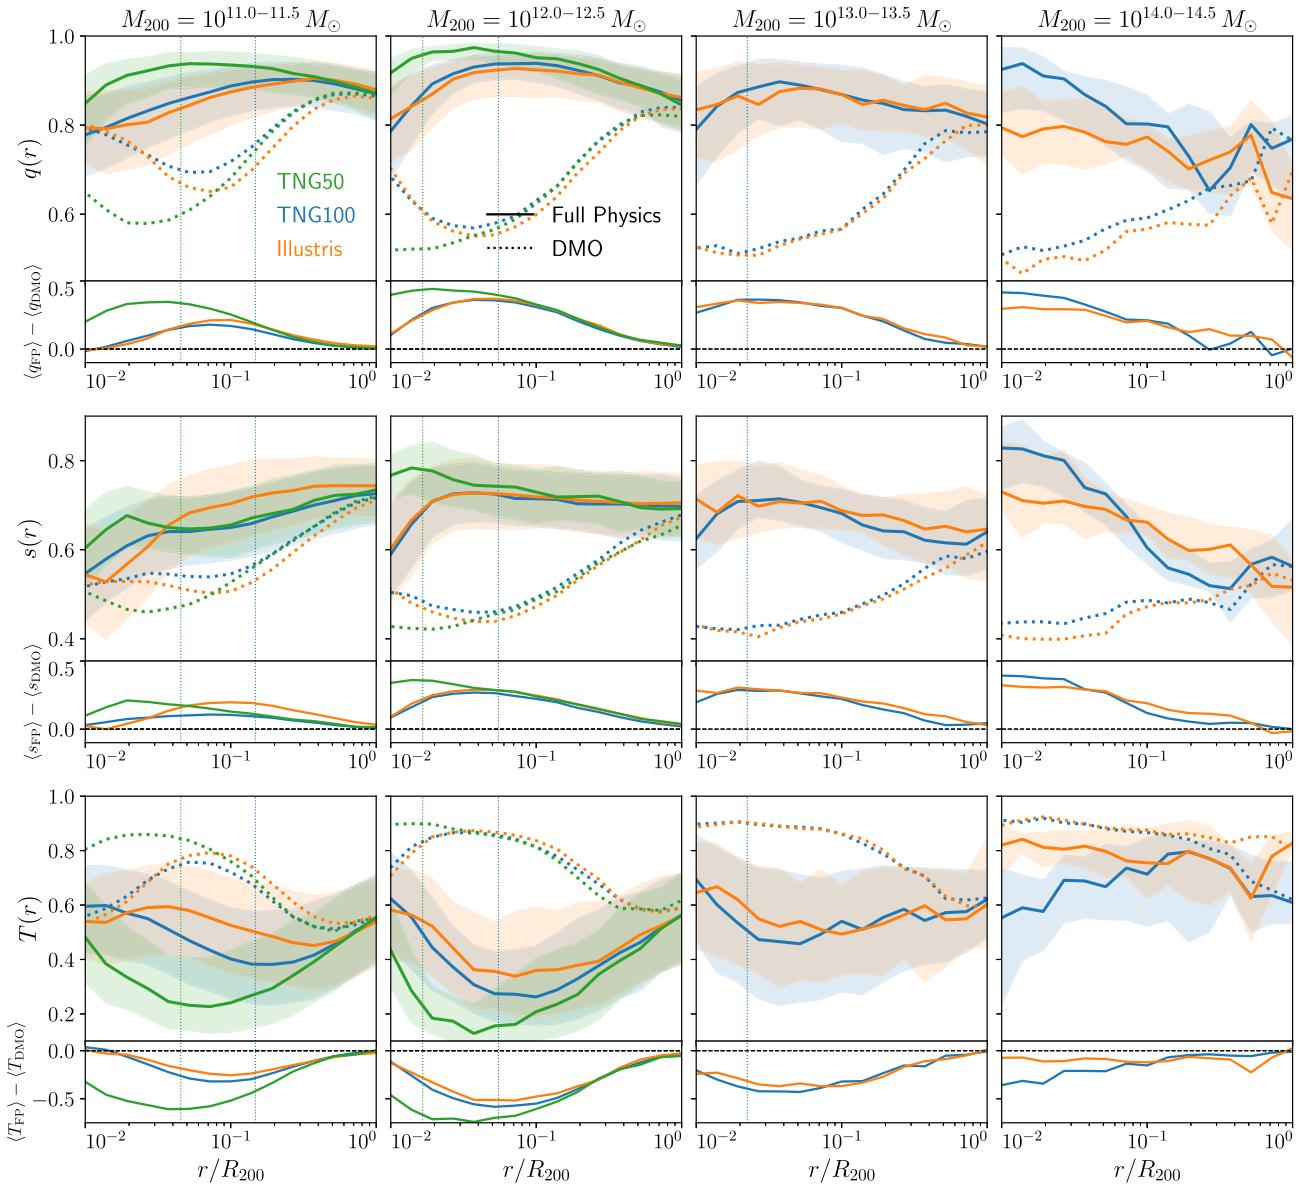


Figure 3. Shape profiles of DM haloes in TNG50 (green), TNG100 (blue), and Illustris (orange). Haloes of mass between 10^{11} and $10^{14} M_{\odot}$ are selected in four mass intervals. Solid curves denote the median profiles in the Full-Physics (MHD and hydrodynamic) simulations, while dotted curves denote median profiles in the DMO simulations. The median convergence radius r_{conv} in TNG50 and TNG100 is indicated by the vertical green and blue lines, respectively. Halo-to-halo variations for the Full-Physics runs are represented by shaded regions that enclose the 25–75th percentiles. The lower attached panels indicate the change in the shape parameters between the Full-Physics and DMO runs. In the Full-Physics simulations, the median halo becomes substantially more spherical (larger q and s) and more oblate (smaller T) for all resolved radii, compared to the DMO results. From the lower attached panels, it is clear that the sphericalization is the largest near the halo centre, and diminishes towards the virial radius. In the DMO runs, the DMO haloes from TNG50-DM and TNG100-DM are also slightly more spherical and oblate on halo scales compared to Illustris-Dark: this is due to the larger σ_8 and Ω_m used in the TNG cosmology.

haloes, and 5 per cent for $10^{12} M_{\odot}$ haloes. Because of the higher mass and spatial resolutions in TNG50-DM, the convergence radius $r_{\text{conv}} \approx 4$ kpc, corresponding to around 10 per cent of the median virial radius in $10^{10} M_{\odot}$ haloes. Further details of the convergence tests are presented in Appendix A. We note that $\kappa(r_{\text{conv}}) = 7$ is consistent with high-resolution results from the Aquarius simulation, and allows the circular velocity to converge to better than 2.5 per cent (Navarro et al. 2010; Vera-Ciro et al. 2011). It is not obvious how equation (4) should be applied or modified in the case of the Full-Physics runs, to account for their larger number and greater variety of resolution elements within the central regions of haloes. Although it is possible that the convergence radii of the halo shapes in e.g. TNG100 and

TNG50 may be in principle smaller than those in TNG100-DM and TNG50-DM, we take the DMO estimates as reference throughout.

3.2 Radial profile of halo shapes in TNG

The radial profiles of the shape parameters q , s , and T of TNG50, TNG100, and Illustris are shown in Fig. 3 for haloes in different mass bins. Solid lines show the median in each of the three Full-Physics (both MHD and hydrodynamic) simulations, while dotted lines correspond to the median results in the DMO counterparts (dotted lines). The shaded regions represent the 25–75th percentile of the distributions in the Full-Physics cases and illustrate the

halo-to-halo variation of the halo shape. The vertical dotted lines denote the convergence radii r_{conv}^{100} in TNG100-DM (blue) and r_{conv}^{50} in TNG50-DM (green).

3.2.1 DMO results

In the DMO simulations (dotted lines), DM haloes become more spherical and oblate towards the virial radius, a well-known result produced by a number of previous N -body studies (e.g. Allgood et al. 2006; Hayashi, Navarro & Springel 2007; Vera-Ciro et al. 2011). Although the shape profiles from all three DMO runs are in good agreement, we note that the average DMO halo from both TNG runs is slightly more spherical and more oblate compared to Illustris-Dark. This effect is most evident for the lower mass 10^{11} and 10^{12} M_\odot haloes and is a result of the larger σ_8 and Ω_m parameters in the TNG cosmology ($\sigma_8 = 0.8159$, $\Omega_m = 0.31$), compared to the Illustris cosmology ($\sigma_8 = 0.81$, $\Omega_m = 0.27$). Thus, Illustris-Dark haloes form later and are less spherical on average (Allgood et al. 2006; Macciò et al. 2008).

3.2.2 Impact of baryons

For both the TNG and Illustris galaxy formation models, Fig. 3 shows that haloes in Full-Physics calculations are both rounder (larger q and s) and more oblate (smaller T) at all radii compared to their DMO counterparts, in qualitative agreement with past hydrodynamic simulations (e.g. Abadi et al. 2010; Butsky et al. 2016; Cataldi et al. 2021). This can be clearly seen from the lower attached panels, which plot the difference between the median shapes in the Full-Physics and DMO simulations. The amount of sphericalization depends on the distance to the halo centre: the effect of baryons is largest near the halo centre but reduces towards the virial radius, where differences become negligible. Taking 10^{12} M_\odot haloes in TNG100, for example, the maximal increase is $\Delta q \approx 0.35$ and $\Delta s \approx 0.25$ near the convergence radius ($r_{\text{conv}}^{100} = 0.05R_{200}$). The effect of baryons extends far beyond the stellar half-mass radius of the central galaxy, which is typically smaller than 20 per cent of the virial radius. Even at half the virial radius, there is non-negligible sphericalization of $\Delta q \approx 0.05$ and $\Delta s \approx 0.05$ in all Full-Physics runs.

In general, there is good agreement between the TNG50 and TNG100 results for 10^{11} and 10^{12} M_\odot haloes, where the simulation results can be compared in a statistically sound manner. With a higher resolution, TNG50 further demonstrates that the TNG100 results also apply to smaller radii: the difference between the Full-Physics and DMO results continues to widen down to the TNG50 convergence radius. There is a tendency for TNG50 haloes to be more spherical and oblate than TNG100 haloes in the same mass bin, a result of the resolution dependency of the galaxy formation model. Briefly, changing the simulation resolution impacts the growth of galaxies (e.g. in terms of star formation and/or strength of baryonic feedback). For the TNG model, at fixed halo mass, higher resolution implies larger stellar masses (see e.g. appendix sections of Pillepich et al. 2018a,b, 2019). Consequently, the halo shape is also affected, due to the relation between galaxy properties and halo shape, which is discussed in Section 4. Further discussion of resolution effects on halo shapes in the Full-Physics runs can also be found in Appendix A.

3.2.3 Comparing TNG100 and Illustris

On average, halo shapes in TNG100 and the older Illustris simulation are similar, with the biggest discrepancies in low-mass and high-mass

haloes. At the low-mass end (10^{11} M_\odot), the TNG model decreases the sphericity s but not q , hence causing the triaxiality to decrease (more oblate) relative to Illustris. At cluster scales (10^{14} M_\odot), TNG100 haloes are in general more spherical and oblate compared to Illustris, especially at smaller radii ($r \lesssim 0.1R_{200}$). Both these results can be traced to the TNG galaxy formation model, which suppresses the stellar mass function at both the low-mass and high-mass ends (Weinberger et al. 2017; Pillepich et al. 2018a; Springel et al. 2018).

3.3 Dependence on halo mass

To examine the dependence of halo shape on halo mass, we focus on the local shape at two specific radii: (i) the virial radius R_{200} representing the outer halo, and (ii) $r_{15} := 0.15R_{200}$ representing the inner halo, where the changes on halo shapes induced by galaxy formation are the strongest. Because of resolution constraints, r_{15} has been determined by the smallest converged radius of 10^{11} M_\odot haloes in TNG100-DM, and 10^{10} M_\odot haloes in TNG50-DM. For simplicity, we denote shape parameters measured at r_{15} with a subscript; i.e. $s_{15} := s(r_{15})$, and similarly for q_{15} and T_{15} .

Fig. 4 plots the halo shape parameters as a function of halo mass at the virial radius (top panels) and at r_{15} (lower panels). Each curve represents the median halo shape, while the shaded regions denote the 25–75th percentile and measure the halo-to-halo variation in the baryonic runs. The lower attached panels show the difference in the median shapes between the baryonic and DMO counterparts, denoted as $\Delta \langle q \rangle = \langle q_{\text{FP}} \rangle - \langle q_{\text{DMO}} \rangle$. Compared to the previous section, we further consider in TNG50 haloes of mass 10^{10} – 10^{11} M_\odot .

3.3.1 Halo shape at the virial radius and in the inner halo

At the virial radius (upper panels), there is little difference between the Full-Physics and DMO runs. Hence, we conclude once again that the impact of baryons is negligible at the virial radius, consistent with the radial profiles shown in Fig. 3. In general, the parameters q and s exhibit a negative correlation with halo mass, while T exhibits a positive correlation with halo mass: more massive haloes tend to be less spherical and more prolate on average. This is in good agreement with results of previous N -body simulations (e.g. Springel, White & Hernquist 2004; Allgood et al. 2006; Macciò et al. 2008; Schneider, Frenk & Cole 2012).

In the inner halo ($r = r_{15}$), differences between the baryonic and DMO runs become significant. Although the axis ratios continue to correlate negatively (and positively for T) with halo mass, this dependence is no longer monotonic in the Full-Physics runs: there are clear peaks at $M_{200} \approx 2 \times 10^{12} \text{ M}_\odot$, where haloes tend to be most spherical and oblate. At this mass, we find that $\langle q_{15} \rangle \approx 0.95$, $\langle s_{15} \rangle \approx 0.75$, and $\langle T_{15} \rangle \approx 0.30$. From the lower panels, these correspond to a difference of $\Delta \langle q_{15} \rangle \approx 0.25$, $\Delta \langle s_{15} \rangle \approx 0.2$, and $\Delta \langle T_{15} \rangle \approx -0.5$ relative to the DMO runs.

Away from the maximum, the impact of baryons in the inner halo decreases. In fact, the TNG50 results show that the differences between haloes in the Full-Physics and DMO runs are negligible for 10^{10} M_\odot haloes, where galaxy formation does not appear to have had an impact on halo shapes.

3.3.2 Comparison to other hydrodynamic simulations

We compare our TNG results to those from the Numerical Investigation of a Hundred Astrophysical Objects (NIHAO) and EAGLE

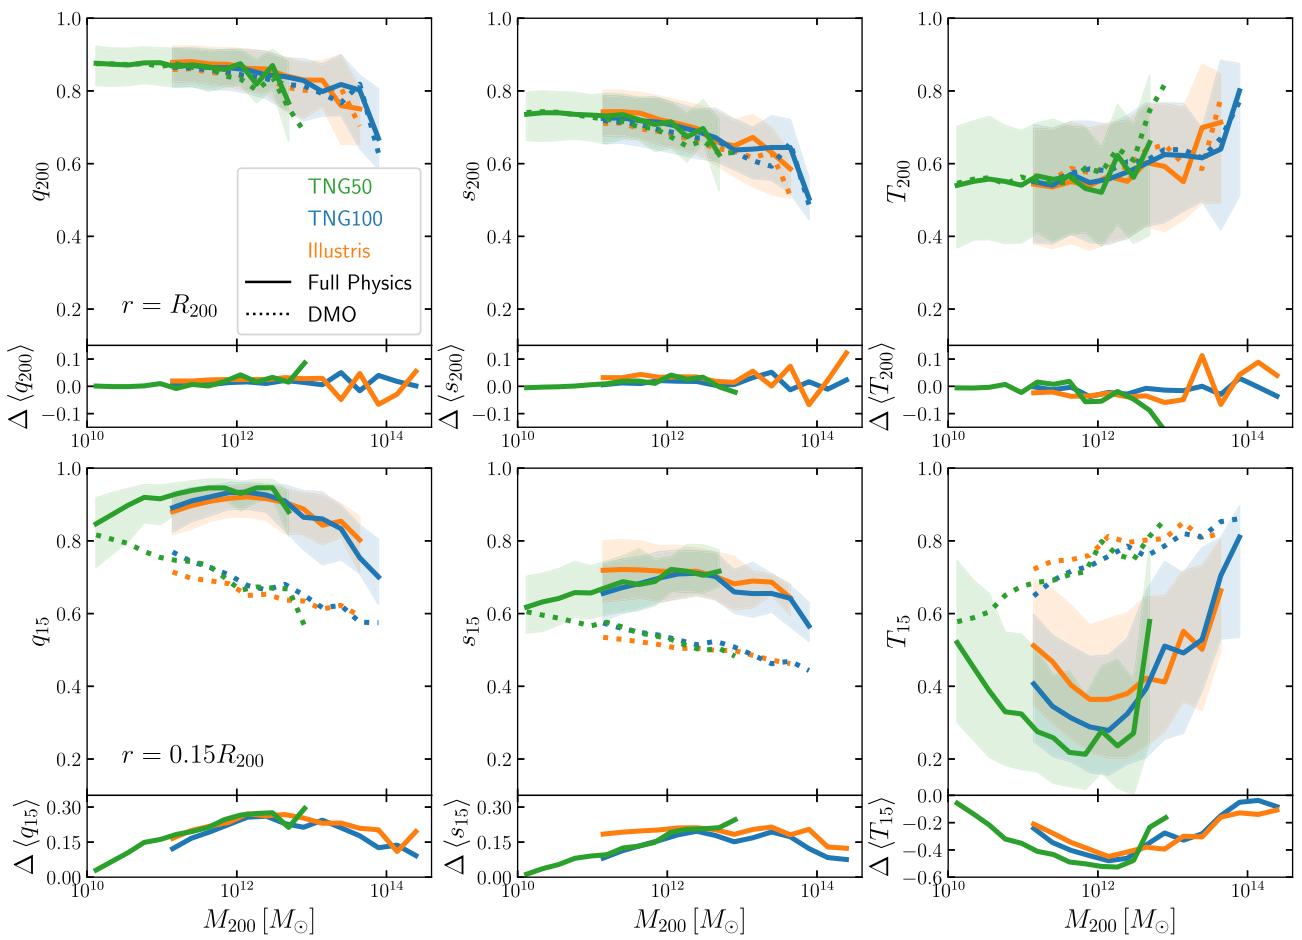


Figure 4. Dependence of DM halo shapes on halo mass at the virial radius (top panels) and in the inner halo (bottom panels, $r_{15} \equiv 0.15R_{200}$). Solid lines denote results from the Full-Physics simulations (TNG and Illustris), while dotted lines denote results from the DMO simulations. The lower attached panels plot the difference in the median parameters between the Full-Physics and DMO simulations, i.e. $\Delta q = \langle q_{FP} \rangle - \langle q_{DMO} \rangle$. In TNG100, baryonic physics and the associated galaxy formation result in non-monotonic behaviour of the halo shapes, peaking at halo masses of $M_{\odot} \approx 2 \times 10^{12} M_{\odot}$.

simulations. NIHAO is a suite of high-resolution simulations of dwarf to MW-size galaxies (Wang et al. 2015), and we make use of the relation between halo shape and mass presented in equation (1) of Butsky et al. (2016). Since the impact of baryons in NIHAO was presented on a halo-to-halo basis, we extract similar results by matching subhaloes between the Full-Physics and DMO runs for TNG50, TNG100, and EAGLE, as described in Section 2. To be consistent with NIHAO, we make the following adjustments for the shape calculation in this section: (i) the inner halo shapes are evaluated interior to $0.12R_{200}$; and (ii) halo shapes are calculated using the *reduced shape tensor*, where the contribution of each particle is weighted by r_{ell}^{-2} . Correspondingly, the minimum halo mass in TNG100(-DM) is increased to $3 \times 10^{11} M_{\odot}$ to ensure that the halo shape at $0.12R_{200}$ is resolved.

Like TNG, the EAGLE project consists of cosmological hydro-dynamical simulations of galaxy formation (Crain et al. 2015). Although the shapes of DM haloes in EAGLE have been analysed and presented in previous work (e.g. Cataldi et al. 2021), the results presented here are based on our own analysis: the haloes and subhaloes are identified identically to TNG (i.e. with the same FOF group finder and SUBFIND), and halo shapes are calculated with the same code used for Illustris and TNG. This removes differences resulting from variations of these procedures and ensures that the TNG and EAGLE halo shapes can be compared consistently.

Fig. 5 shows the scatter plots of the relative halo shapes in the Full-Physics and DMO simulations as a function of halo mass. On the x -axis, masses refer to the halo mass of the Full-Physics counterpart. The median relations in TNG50 and TNG100 are presented by green and blue curves. In general, we find that the curves are qualitatively similar to those in the lower attached panels of Fig. 4. Combining haloes from both TNG50 and TNG100, we decided to capture the shape of the ratio s_{FP}/s_{DMO} by fitting a cubic function:

$$f(M) = \frac{s_{FP}}{s_{DMO}} = \alpha + \gamma(\log_{10} M - \beta)^2 + \delta(\log_{10} M - \beta)^3, \quad (5)$$

where $M \equiv M_{200}$ is the virial mass of the halo. By nature of the halo mass function, small haloes strongly outnumber more massive ones, hence we weight each halo inversely proportional to the halo mass function. The Levenberg–Marquardt algorithm is applied to solve the least-squares problem, with the best-fitting parameters shown in Table 4. The best-fitting curve to the combined TNG50 and TNG100 data is shown by the black dashed line in Fig. 5, with a maximum average ratio of 1.5 for haloes close to $10^{13} M_{\odot}$.

The EAGLE results based on our own analysis of the data are plotted in red, which shows good agreement between EAGLE and TNG100. In particular, the impact of baryons, as captured by the Full-Physics to DMO ratios, exhibits similar dependencies on the halo mass for all three parameters, with maxima/minima located

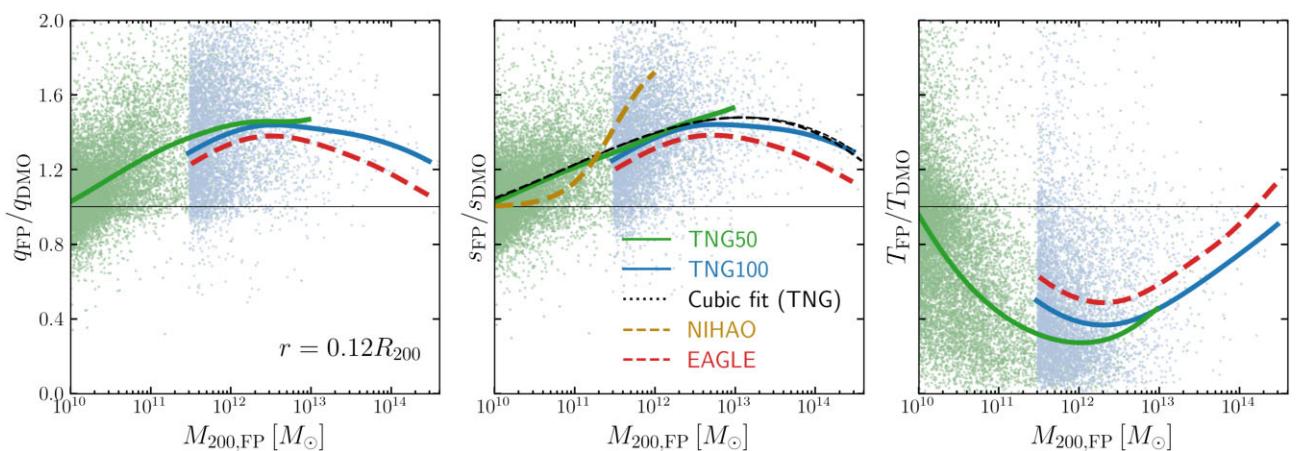


Figure 5. Comparison with NIHAO (golden dashed lines; Butsky et al. 2016) and EAGLE (red dashed lines; Crain et al. 2015). We plot the ratio between the inner halo shape of matched haloes between the Full-Physics and DMO simulations, as a function of the halo mass of the Full-Physics counterpart. The scatter plots indicate individual matched halo pairs in TNG50 and TNG100, while the median is represented by solid lines. For $s_{\text{FP}}/s_{\text{DMO}}$ (middle plot), a least-squares fit is performed to a cubic equation (equation 5), and the resulting fit is shown as the black dashed line. To facilitate a consistent comparison with NIHAO, here we show results for halo shapes evaluated interior to $r = 0.12R_{200}$ using the reduced shape tensor.

Table 4. Best-fitting parameters to equation (5) for $s_{\text{FP}}/s_{\text{DMO}}$, the ratio between the shape parameter s in the baryonic and DMO simulations (middle panel of Fig. 5).

| α | β | γ | δ |
|----------|---------|----------|----------|
| 1.48 | 13.1 | -0.081 | -0.011 |

approximately at $3 \times 10^{12} M_{\odot}$. There is a small overall vertical shift implying a smaller impact of baryons in EAGLE compared to TNG100, especially at the high-mass end. However, this vertical shift is small compared to the halo-to-halo variation, and to the difference between TNG50 and TNG100 in the 10^{11} – $10^{12} M_{\odot}$ halo mass range.

The NIHAO relation presented in Butsky et al. (2016) is shown as the brown dashed curve in the middle panel, which is an increasing S-shaped curve. At a halo mass of $10^{10} M_{\odot}$, both the TNG50 and NIHAO results agree that the influence of baryons is negligible. Although both sets of simulations show that the ratio $s_{\text{FP}}/s_{\text{DMO}}$ increases with halo mass up to $\approx 10^{12} M_{\odot}$, NIHAO exhibits a much stronger dependence and the curve increases sharply between 10^{11} and $10^{12} M_{\odot}$. At $10^{12} M_{\odot}$, the value of the ratio in NIHAO is ≈ 1.7 , substantially larger than the TNG and EAGLE curves. Since the NIHAO sample only contains haloes hosting dwarf galaxies to MW-size galaxies, it is unclear if the effect of baryons decreases for more massive haloes. The higher sphericities in NIHAO can be explained by the larger central densities of NIHAO haloes. Although NIHAO is able to reproduce the stellar mass–halo mass relation and star formation rates across cosmic time, black hole feedback is not included in the model, which results in high central peaks in some of the circular velocity curves for $10^{12} M_{\odot}$ haloes (Wang et al. 2015). In contrast, the circular velocity profiles of similar-mass haloes in TNG100 are flatter, owing to the inclusion of black hole feedback (Lovell et al. 2018).

4 GALACTIC FEEDBACK VARIATIONS

In this section, we examine the effects of baryonic physics variations using the $25 \text{ Mpc } h^{-1}$ smaller boxes. A summary of the nine feedback variations examined has been presented in Table 2.

4.1 Stellar mass–halo mass relation

Observations have shown that the galaxy formation efficiency depends non-monotonically on the halo mass and is maximal for MW-sized haloes of $10^{12} M_{\odot}$ (e.g. Conroy & Wechsler 2009; Leauthaud et al. 2012). We parametrize the galaxy formation efficiency using the stellar mass fraction m_{*}/M_{200} , where m_{*} refers to the stellar mass content within twice the stellar half-mass radius.⁴

We first show the stellar mass fraction at $z = 0$ as a function of halo mass for the various feedback models in Fig. 6. For comparison, we include results from TNG100 (blue dashed), and the semi-empirical model of Behroozi et al. (2013) (solid grey). The fiducial small box (solid blue) matches well with both TNG100 and Behroozi et al. (2013), with the stellar mass fraction maximized for haloes close to $10^{12} M_{\odot}$. Note however that a quantitative comparison between our results and those from semi-empirical models can be difficult due to contrasting definitions for the stellar mass between simulations and observational data sets (see Pillepich et al. 2018a,b for a detailed discussion).

For lower mass haloes (10^{11} – $12 M_{\odot}$), the stellar mass fraction is primarily affected by changes in the galactic winds. Turning off galactic winds (*no winds*) results in the largest increase in the stellar mass fraction. Similarly, decreasing the strength and speed of galactic winds both lead to an increase in the stellar mass fraction, evident from the *slow winds* and *weak winds* cases. Conversely, increasing the wind speed and energy inhibits star formation, and decreases the stellar mass fraction compared to the fiducial model. The corresponding galaxy stellar mass functions at $z = 0$ can be found in Pillepich et al. (2018a): it is manifest that some of the TNG model-variation runs shown here are known to be ruled out by observational constraints.

Although turning off BH feedback also increases star formation, the resulting increase is less significant compared with changes to the galactic wind, especially at lower halo masses. For more massive haloes ($M_{200} > 10^{12} M_{\odot}$), BH feedback plays a larger role in determining the stellar mass fraction. For example, turning off the

⁴The galaxy formation efficiency can also be defined as $(m_{*}/M_{\text{halo}})/(\Omega_b/\Omega_m)$. Also, see Pillepich et al. (2018b) regarding alternative definitions of the stellar mass content.

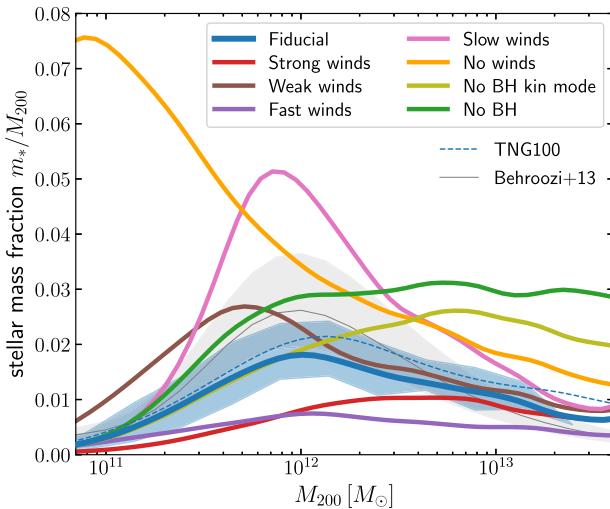


Figure 6. Stellar mass fraction as a function of halo mass in the small L25n512 boxes. Galaxy stellar masses are evaluated within twice their stellar half-mass radius. The blue shaded region denotes the 25–75th percentiles for the fiducial run. Results from TNG100 (dashed blue) and the semi-empirical model of Behroozi, Wechsler & Wu (2013) (solid grey) are also shown for comparison. In general, reducing the strength of baryonic feedback (winds and BH) leads to an increase in the stellar mass fraction; increasing the strength and speed of galactic winds (*fast winds* and *strong winds*) leads to a decrease in the stellar mass fraction.

BH kinetic feedback or turning off BH feedback completely both result in the largest increase to the stellar mass fraction. Unlike in the smaller haloes, increasing wind speeds appears to be more effective than increasing the wind energy in suppressing star formation.

4.2 Radial profile of halo shapes with physics variation

The effects of the various feedback models on the radial profiles of the shape parameters are shown in Fig. 7. Here, we note that the convergence radius (dotted vertical line), as determined using the DMO run, is similar to TNG100-DM due to their similar resolutions. Hence the radius r_{15} remains resolved in the small boxes.

For $10^{11-11.5} M_\odot$ haloes (top row), the largest impact results from turning off galactic winds (*no winds*), which causes haloes to become significantly rounder compared to the fiducial model. At $r = 0.15R_{200}$, the median sphericity is increased by $\Delta \langle s \rangle \approx 0.1$ compared to the fiducial model. However, such a feedback model is clearly unrealistic and leads to an overproduction of stars in low-mass haloes (Fig. 6). In the *strong winds* and *fast winds* cases, both q and s are slightly decreased, while the triaxiality T is increased compared to the fiducial run. There is little impact from changing the black hole feedback on the halo shapes of low-mass haloes. In the inner halo (at $r = 0.15R_{200}$), the median $\langle q \rangle$ varies by ≈ 0.6 across the feedback variations, while the median $\langle s \rangle$ varies by ≈ 0.75 when we ignore the unphysical *no winds* case.

For $10^{12-12.5} M_\odot$ haloes (bottom row of Fig. 7), the feedback variations result in a larger diversity of the median halo shapes. In the inner halo, decreasing the speed of galactic winds (*slow winds*) or turning off black hole feedback (*no BH*) results in similar increases to the sphericity s : the median s is increased by ≈ 0.05 . On the other hand, decreasing the wind energy (*weak winds*) does not have a noticeable impact on the halo shape profile. Interestingly, the effect of turning off the black hole kinetic mode (*no BH kin mode*) is similar to turning off BH feedback completely for the parameter q , but has

no effect on s in the inner halo. Both increasing the wind speed (*fast winds*) and the wind energy (*strong winds*) results in less spherical haloes, with the largest decrease observed for the *fast winds* case.

For $10^{12} M_\odot$ haloes, the median $\langle q \rangle$ and $\langle s \rangle$ varies by ≈ 0.1 in the inner halo across all the runs. This value is comparable to the halo-to-halo variation in the fiducial model, indicated by the blue shaded regions in Fig. 7, which suggests that ruling out different feedback models requires sufficiently large statistics. Furthermore, we note that the inner halo shape differences of ≈ 0.1 correspond to a change of almost 20 per cent with respect to the DMO halo shape ($s_{\text{DMO}} \approx 0.55$): different feedback model predictions are closer to one another than to the DMO ones, also in the case of models whose resulting galaxy populations are known to be unrealistic (e.g. the *no BH* and *no winds* cases).

4.3 Dependence on galaxy formation efficiency

The differences in halo shape profiles induced by the model variations suggest that the galaxy formation efficiency is closely linked to the resulting halo shapes in baryonic simulations, since the halo shape reflects the competition between the competing processes of star formation and feedback. The link between halo shape and the growth of the baryonic component is further supported by controlled N -body experiments that study the effect of baryon condensation by simulating the growth of baryonic discs within triaxial haloes. For example, Debattista et al. (2008) found that the final halo shape depends on the mass of the baryonic disc, and Kazantzidis et al. (2010) concluded that the halo response depends most strongly on the overall gravitational importance of the disc.

To verify these conclusions, we plot the shape parameters q and s in the inner halo ($r = 0.15R_{200}$) as a function of the stellar mass fraction in Fig. 8. The left-hand panels show the results in the large boxes (TNG50, TNG100, and Illustris), with curves denoting the median of the distributions for haloes of different masses. In general, the shape parameters correlate positively with the stellar mass fraction: haloes are rounder (larger q and s) with increasing stellar mass fraction, in agreement with results from previous simulations (e.g. Bryan et al. 2013; Butsky et al. 2016; Chua et al. 2019; Cataldi et al. 2021). For stellar mass fractions of $m_*/M_{200} \gtrsim 0.05$, the average sphericity is $\langle s \rangle \gtrsim 0.75$, significantly larger than the average DMO value (horizontal grey line and shaded region). As expected, halo shapes approach the DMO value as stellar mass fraction decreases, and baryonic effects become negligible for $m_*/M_{200} < 10^{-4}$. Note that most curves exhibit an upturn in the halo shapes at the lower end of the stellar mass fractions, which we believe is an indirect resolution effect. The strength of the correlations are further quantified using the Pearson r and Spearman ρ coefficients, which are shown in Table 5. In all cases, we observe moderate but significant correlations between the shape parameters and the stellar mass fraction.

The results here also suggest that even though the stellar mass fraction normalizes for the halo mass, there is residual mass dependence of the halo shape–stellar mass fraction relation: at a fixed stellar mass fraction, less massive haloes tend to be more spherical than their more massive counterparts. This is associated with the DMO halo shape–halo mass dependence previously touched upon in Section 3.3, which continues to hold in baryonic simulations after controlling for the galaxy stellar mass fraction.

The results of the physics variations shown in the right-hand panels of Fig. 8 paint a similar picture, for haloes of $10^{11-12} M_\odot$ only: the median curve for each variation displays a positive correlation between halo shape and stellar mass fraction. In particular, there is substantial overlap across all curves, and the differences across the

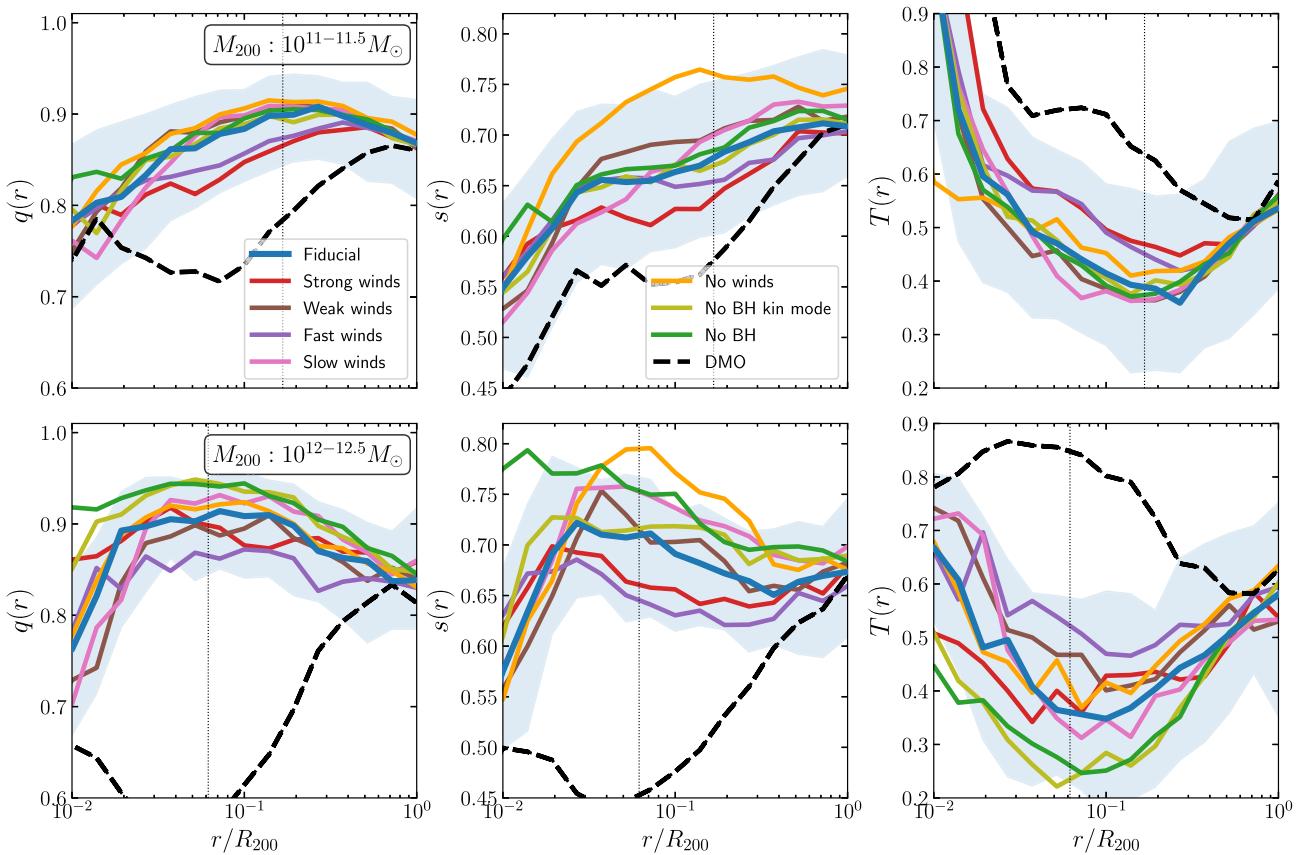


Figure 7. Effect of changing the feedback prescription on the radial profile of DM halo shapes. These simulations are carried out in the small L25n512 boxes. The top and bottom rows correspond to haloes of mass $10^{11-11.5} M_{\odot}$ and $10^{12-12.5} M_{\odot}$, respectively. The solid lines denote the median, while the blue shaded region represents the 25–75th percentile in the fiducial model. Changing the feedback prescription has a larger impact on s (middle panels) compared to q (left-hand panels). In general, faster and stronger galactic winds result in less spherical (smaller q and s) and more prolate (larger T) haloes. The massive haloes (bottom row) exhibit larger changes in response to variations in feedback prescription. An increase in the strength of galactic winds results in less spherical haloes, and conversely for a decrease. While black holes only have a minor effect on $10^{11} M_{\odot}$ haloes, $10^{12} M_{\odot}$ haloes become rounder when black hole feedback is turned off.

cases are small ($\lesssim 0.25$) at fixed stellar mass fractions. This provides a strong indication that the halo shape in baryonic simulations is primarily dependant on its galaxy formation efficiency, and independent of the particular galaxy formation model used. At halo masses of $10^{11-12} M_{\odot}$, the outcome of the model-variation runs also differs substantially from the DMO predictions, for most implementations and stellar mass fractions larger than 1–2 per cent.

Finally, we note here that for MW-sized haloes, the most spherical haloes in TNG have sphericities of $s_{15} \approx 0.85$. This is consistent with the simulations of Abadi et al. (2010), which considered the maximal effects of baryons by neglecting stellar formation and other forms of baryonic feedback. In their analysis, they found that this model resulted in a halo with sphericity $s \sim 0.85$ approximately constant across radii.

4.4 Comparison to observations of the Milky Way halo shape

In our Galaxy, the shape of the inner halo ($\lesssim 60$ kpc) has been inferred through observations of the motion of individual stars. Kinematics of halo stars combined with equilibrium modelling using Jean's equations can be used to infer force fields and thus the MW halo shape (e.g. Loebman et al. 2012; Bowden, Evans & Williams 2016). Stellar streams formed from the tidal stripping of satellite galaxies or globular clusters can also be used to constrain the halo potential.

Examples of stellar streams that have been used include that of the Sagittarius dwarf galaxy (Ibata et al. 2001; Law & Majewski 2010; Vera-Ciro & Helmi 2013), as well as the Pal 5 and GD-1 tidal streams (Bovy et al. 2016). In general, these observations indicate that the MW inner halo is substantially more spherical than DMO predictions: $s = 0.72$ (20–60 kpc; Law & Majewski 2010), $s = 0.8$ (Vera-Ciro & Helmi 2013), $s = 1.05$ (20 kpc; Bovy et al. 2016), and $s = 0.82$ (Malhan & Ibata 2019), at radii somewhat smaller than 0.2–0.3 times the virial radius of the Galaxy.

In Fig. 8, we have presented a quick comparison to the MW observations by indicating a lower bound for the stellar stream measurements of the sphericity ($s \gtrsim 0.75$). The horizontal error bar represents the uncertainty in the stellar mass fraction of our Galaxy, ranging approximately between 0.01 and 0.1, consistent with observational estimates (e.g. Zaritsky et al. 2020). The stellar mass fraction the Galaxy corresponds to is the higher end in the simulations, where the observational estimates and numerical predictions are in good agreement.

A more thorough comparison between observational and simulation results is presented in Fig. 9, which plots the sphericity s_{15} as a function of q_{15} for MW-size haloes (7×10^{11} – $3 \times 10^{12} M_{\odot}$). The left-hand panel shows the results in TNG50, TNG100, and EAGLE, while the right-hand panel shows the results for the model variations. Estimates of the MW inner halo shape from stellar streams are shown

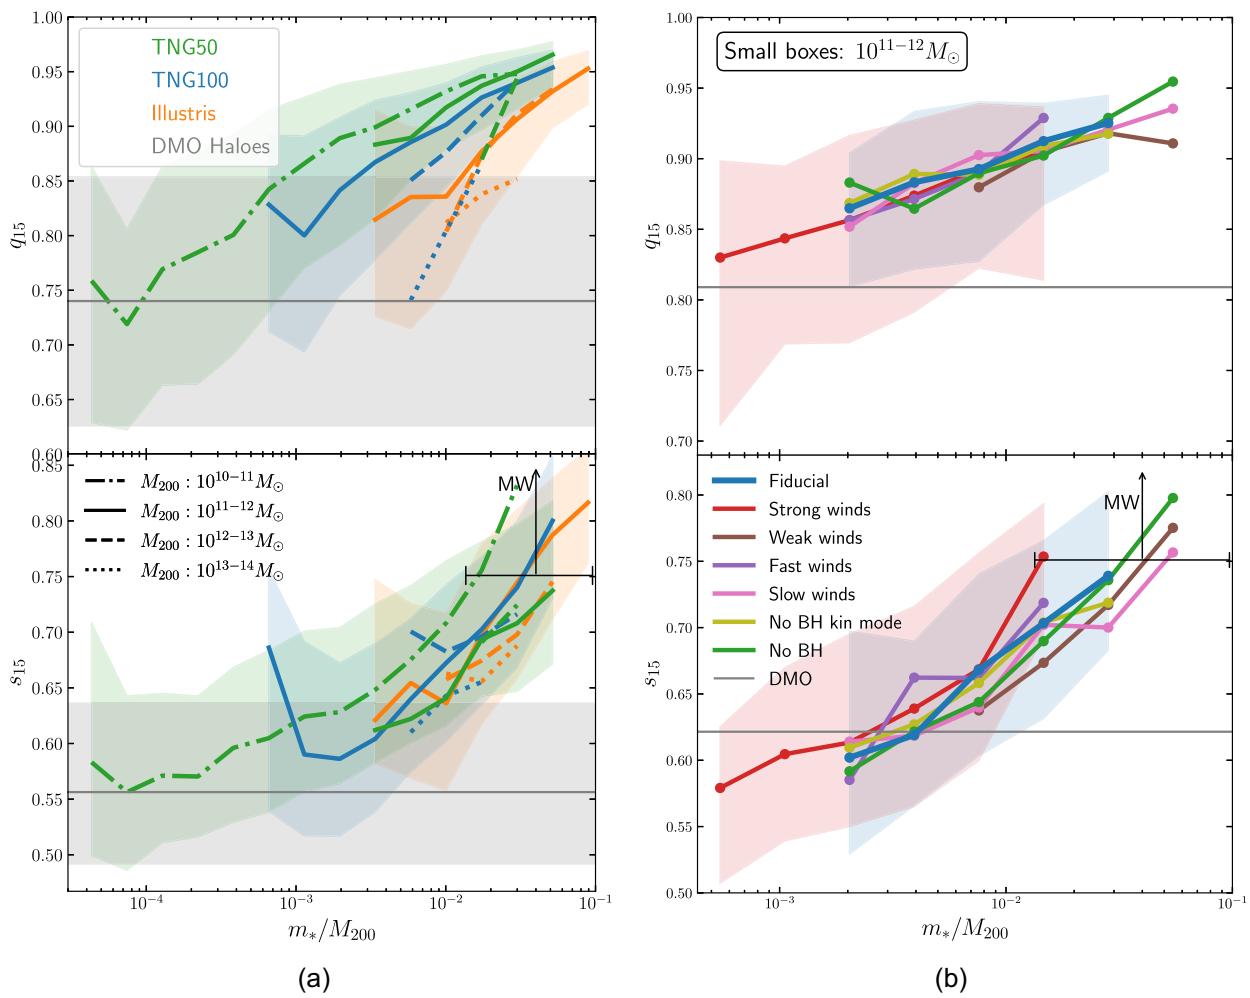


Figure 8. Relationship between the inner halo shape (at $r = 0.15R_{200}$) and the stellar mass fraction. (a) Panels on the left show the shape parameters s and q in the flagship boxes (TNG50, TNG100, and Illustris) for haloes in four mass intervals. The curves represent the median, while the shaded regions represent the 25–75th percentiles. The DMO results shown in grey for comparison consist of all resolved haloes with mass $> 10^{11} M_\odot$ in TNG50 and TNG100. The Full-Physics simulations predict a strong dependency of halo shapes on the stellar mass fraction, with maximal sphericalization for haloes with the highest galaxy formation efficiency. (b) The panels on the right compare the relationship between inner halo shape and stellar mass fraction for different galaxy formation implementations carried out in the small L25n512 boxes. Here, only haloes of mass $10^{11-12} M_\odot$ are considered. The different curves are in general agreement, which suggests that halo shape differences caused by varying baryonic feedback in the simulations can be explained by its cumulative effects on the stellar mass fraction. A lower bound for the sphericity s of the Galaxy (MW) from observations is shown for comparison, with the horizontal error bar representing the uncertainty in the stellar mass fraction.

for comparison, with error bars denoting the 1σ error (if provided by the reference). Three of the observations (Vera-Ciro & Helmi 2013; Bovy et al. 2016; Malhan & Ibata 2019) indicate that the MW halo is axisymmetric in the disc plane, hence $q \equiv b/a$ is almost exactly one.

In general, the shift to larger sphericities in the MHD; i.e. Full-Physics, runs means that TNG50 and TNG100 are in much better agreement with observations compared to their DMO counterparts. Since Bovy et al. (2016) estimate the MW halo to be almost completely spherical, we find a slight disagreement with their result at the 1σ level. Interestingly, this disagreement is not reduced even when we consider weaker feedback formulations, shown in the right-hand panel. Although the median s increases for models with weaker feedback (e.g. *weak winds* and *no BH* cases), the upper percentiles of the distributions (in s) do not shift upwards sufficiently to meet the Bovy et al. (2016) inferences: in fact, the galaxies produced by those models are unrealistic and so a possible agreement in MW halo shapes would still point towards a tension.

We illustrate these findings more clearly in Fig. 10, which plots the kernel density estimates of the distributions of q and s in the small boxes. The upper and lower panels show the halo shape distributions in the inner halo and at the virial radius, respectively. Although decreasing the baryonic feedback strength tends to shift the distribution of s_{15} (upper right) to larger values of s , the biggest changes are for haloes with $s_{15} \lesssim 0.85$. If accurate, the Bovy et al. (2016) results would rule out the DMO, *strong winds* and *fast winds* models, which have negligible fractions of haloes with $s_{15} > 0.9$. In the remaining models, the fractions of haloes with $s_{15} > 0.9$ are small but non-zero, thus it is possible for these models to produce haloes compatible with the Bovy et al. (2016) results. Since observations generally assume or find that the minor axis of the halo and of the stars (i.e. Galaxy) coincide, we note that it may be possible to improve the agreement with observations by projecting the halo shapes perpendicular to the stellar axis (Chua et al. 2019).

Table 5. Pearson and Spearman correlation coefficients and their associated p -values between the shape parameter s and the stellar mass fraction. The first seven rows are results for model variations in the small boxes, while the last two rows correspond to the larger boxes, TNG100 and TNG50. For TNG100 and the small boxes, all haloes of mass greater than $10^{11} M_{\odot}$ are considered. For TNG50, all haloes of mass greater than $10^{10} M_{\odot}$ are considered. All coefficients are positive and significant, pointing to a strong correlation between the sphericity of a halo and its galaxy formation efficiency.

| Simulation | Pearson r (p -value) | Spearman ρ (p -value) |
|------------------------|---------------------------------|---------------------------------|
| 1. Fiducial | 0.259 (4.9×10^{-17}) | 0.343 (1.7×10^{-17}) |
| 2. Strong winds | 0.338 (2.3×10^{-10}) | 0.259 (2.3×10^{-10}) |
| 3. Weak winds | 0.282 (2.3×10^{-12}) | 0.276 (5.9×10^{-12}) |
| 4. Fast winds | 0.223 (1.2×10^{-7}) | 0.214 (3.6×10^{-7}) |
| 5. Slow winds | 0.444 (2.8×10^{-30}) | 0.455 (4.7×10^{-32}) |
| 6. No BH kinetic model | 0.365 (1.0×10^{-19}) | 0.373 (1.3×10^{-20}) |
| 7. No BH | 0.435 (9.4×10^{-29}) | 0.440 (1.8×10^{-29}) |
| 8. TNG100 | 0.352 (0) | 0.368 (0) |
| 9. TNG50 | 0.301 (0) | 0.355 (0) |

In general, the results indicate that the inner DM halo shapes predicted by Full-Physics models differ substantially from those from DMO calculations and that, even for unrealistic feedback models, their predictions are closer to one another than to the DMO results. However, the distributions in Fig. 10 point to a difficulty in using the DM halo shape as a constraint on baryonic feedback models: varying the baryonic feedback only results in distributional differences that are small compared to the large halo-to-halo variation ($\sigma \approx 0.15$). Hence, a large statistic, i.e. a large halo sample, is necessary to statistically rule out the various feedback models. For hydrodynamic simulations, obtaining a large sample is challenging due to the high computational costs of high-resolution simulations with large box sizes such as TNG50 and TNG100. Observationally, current methods of inferring the three-dimensional DM halo shape rely strongly on stellar observations and are thus only feasible for nearby galaxies.

Near the virial radius, no feedback variation is able to significantly affect the halo shape. This is illustrated in the lower panels of Fig. 10, which show that the distributions across all feedback variations are similar to the DMO case.

5 SUMMARY

We have used a suite of cosmological simulations to investigate the impact of galaxy assembly on the shape of DM haloes at redshift $z = 0$. To elucidate the effects of the fiducial galaxy formation model, we have examined haloes from TNG50 and TNG100, which are both part of the suite of MHD cosmological simulations IllustrisTNG. Because of the high resolution of TNG50, we were able to reliably resolve halo shapes down to halo masses of $10^{10} M_{\odot}$. In total, we have analysed statistically significant samples of ≈ 10000 and ≈ 14000 haloes in TNG50 and TNG100, respectively, spanning a halo mass range of $10^{10}–10^{14} M_{\odot}$. The convergence radii in TNG100-DM and TNG50-DM are determined to be $r_{\text{conv}} = 10$ and 4 kpc, respectively. These values correspond to 15 per cent of the virial radius for $10^{11} M_{\odot}$ haloes in TNG100, and 10 per cent of the virial radius for $10^{10} M_{\odot}$ haloes in TNG50. We have also investigated a set of $25 h^{-1}$ Mpc smaller boxes, which have numerical resolutions comparable to that of TNG100. A total of nine model variations have been analysed, including five galactic wind feedback variations and two black hole feedback variations: these have allowed us to bracket the effects on simulated halo shapes due to lingering uncertainties in the galaxy formation models, also including unrealistic feedback implementations.

Using an iterative algorithm for the unweighted shape tensor, we have quantified halo shapes in ellipsoidal shells as a function of radius, which are summarized with the parameters $q(r) \equiv b/a$, $s(r) \equiv c/a$, and the triaxiality parameter $T(r) \equiv (1 - q^2)/(1 - s^2)$. When focusing on the inner halo, we have chosen to represent the inner halo shape by the measurement at $r_{15} \equiv 0.15R_{200}$, where haloes across the examined mass range are determined to be resolved.

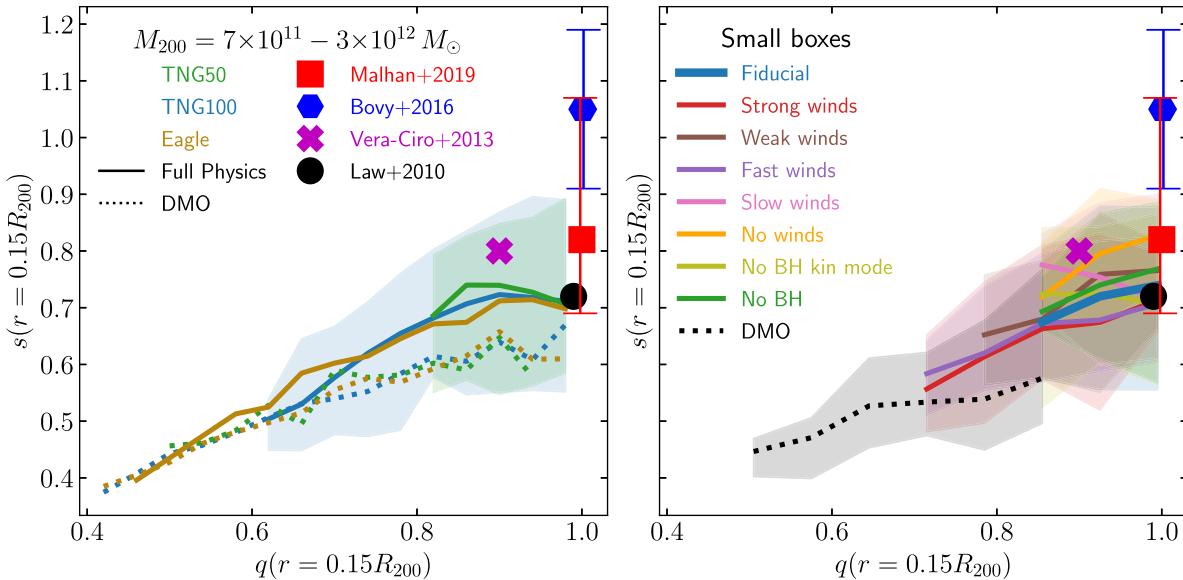


Figure 9. Distribution of the sphericity s as a function of the parameter q in the inner halo, for haloes of mass $M_{200} = 7 \times 10^{11}–3 \times 10^{12} M_{\odot}$. The left-hand panel shows the results from TNG50, TNG100, and EAGLE, while the right-hand panel shows the results for the different baryonic feedback models in the smaller boxes. Solid lines represent results from MHD simulations, while dashed lines represent results from the DMO simulation. Here, the shaded region represents the 1σ interval in the MHD simulations. The red, blue, magenta, and black symbols denote observational estimates of the MW halo shape by Malhan & Ibata (2019), Bovy et al. (2016), Vera-Ciro & Helmi (2013), and Law & Majewski (2010), with error bars denoting the 1σ uncertainty (if available).

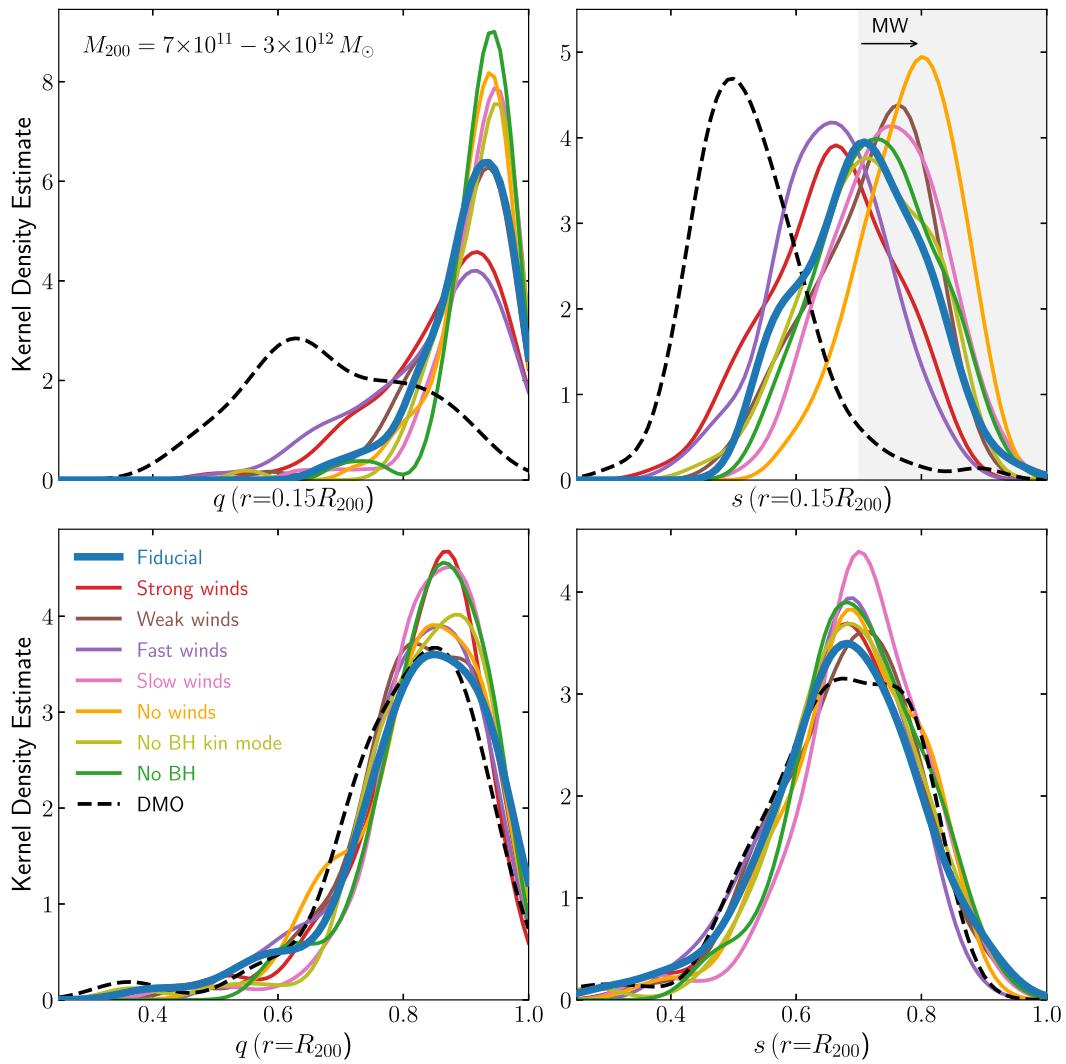


Figure 10. Effect of physics variations on the distributions of shapes of MW-sized DM haloes ($M_{200} = 7 \times 10^{11} - 3 \times 10^{12} M_{\odot}$) in the smaller boxes. The distribution for each feedback model is represented by a kernel density estimate. The grey shaded region denotes the range of observational estimates of the MW sphericity s in the inner halo. In the inner halo ($r = 0.15R_{200}$, top row), all the feedback models result in haloes that are substantially more spherical than those of the DMO simulation. Substantial overlap remains across the feedback variants, rendering it difficult to distinguish between these models based solely on measurements of the MW halo shape. At the virial radius (bottom row), regardless of the feedback model, the effect of baryons becomes negligible.

We compared the TNG haloes to their counterparts in the DMO analogues TNG50-DM and TNG100-DM, as well as the previous generation Illustris. Furthermore, we have also compared our results to previous simulations and observational estimates of the MW halo shape. We summarize our results as follows.

(i) In both TNG50 and TNG100, haloes are more spherical and oblate on average than their DMO counterparts (Figs 3 and 4), consistent with previous hydrodynamic simulations with realistic galaxy formation models (e.g. Butsky et al. 2016; Chisari et al. 2017; Prada et al. 2019; Cataldi et al. 2021). The effect of baryons is negligible in small $10^{10} M_{\odot}$ haloes, and maximal for $10^{12} M_{\odot}$ haloes, where the median $\langle q \rangle$ and $\langle s \rangle$ exceeds that of the DMO runs by almost 0.5 and 0.3 at $r = 0.02R_{200}$. Minor discrepancies between TNG50 and TNG100 can be noticed, due to the dependency of the baryonic physics models on simulation resolution.

(ii) Based on our own post-processing and analysis of the EAGLE data, we find good agreement between TNG100, EAGLE, and Illustris (Figs 5 and 9). We highlight that the impact of baryons, as

captured by the Full-Physics to DMO ratio, has a similar dependence on halo mass, which points to a quantitative agreement in numerical halo shape predictions across different hydrodynamic simulations.

(iii) Through the feedback variation runs in smaller simulation boxes (Table 2), we show, on the wake of the results of Weinberg et al. (2008) and Pillepich et al. (2018a), that the galaxy formation efficiency as parametrized by the stellar mass fraction is strongly affected by the changing feedback strength. In general, increasing (decreasing) the galactic wind speed or energy suppresses (increases) the stellar mass fraction. Turning off black hole feedback has the effect of increasing the stellar mass fraction, primarily at larger halo masses.

(iv) We show that varying the feedback prescription has an impact on the median halo shape parameters (Fig. 7): on average, stronger feedback prescriptions (increasing the galactic wind energy or speed) cause haloes to become less spherical and more prolate. Conversely, weaker feedback prescriptions (through the reduction of galactic wind speeds or suppression of black hole feedback) result in more spherical and oblate haloes. However, in the inner haloes, we find a maximum change in the median parameters $\langle q \rangle$, $\langle s \rangle$, and $\langle T \rangle$ of around

0.1 across the nine feedback-model variations: this is comparable to the halo-to-halo variation for the fiducial IllustrisTNG model, which has standard deviation $\sigma \approx 0.15$, and makes the results of any feedback model studied here closer to one another than to the DMO expectations.

(v) We demonstrate that a positive correlation is in place between the shape of the inner halo and the stellar mass fraction: haloes with larger stellar mass fractions are rounder and more oblate (Fig. 8), in agreement with previous numerical work (e.g. Bryan et al. 2013; Butsky et al. 2016; Cataldi et al. 2021). At a fixed stellar mass fraction, the predicted inner halo shapes are consistent across the simulations examined (TNG50, TNG100, and the feedback variations), within average variations of 0.1. In particular, we note that the effect of baryons is negligible when the stellar mass fraction is $m_*/M_{200} \lesssim 10^{-4}$. However, for stellar mass fractions larger than about 1–2 per cent, all feedback variation runs predict average halo shapes that are substantially different from those from DMO calculations and more compatible with one another than with the DMO results, even in the case of manifestly unrealistic feedback implementations.

(vi) Comparing with several observational estimates of the MW halo shape, we find good agreement between observations and the fiducial model predictions from TNG50 and TNG100 at the 1σ level (Fig. 9). With the small boxes, we demonstrate that the eight baryonic models (excluding the DMO model) are consistent with a value of $s_{\text{MW}} = 0.7\text{--}0.8$ for the MW halo shape, due to the large overlap in the shape distributions (Fig. 10). Although these values disfavour the models with stronger galactic feedback, the halo shape alone does not provide a very strong constraint on baryonic feedback models, but may do so in combination with complementary galaxy diagnostics and better statistics.

ACKNOWLEDGEMENTS

MV acknowledges support through NASA ATP grants 16-ATP16-0167, 19-ATP19-0019, 19-ATP19-0020, and 19-ATP19-0167, and NSF grants AST-1814053, AST-1814259, AST-1909831, and AST-2007355.

DATA AVAILABILITY

The snapshot data for TNG50 and TNG100 can be accessed at www.tng-project.org and for Illustris at www.illustris-project.org. The corresponding data release papers are Nelson et al. (2019a) and Nelson et al. (2015). Other data used in this paper will be shared upon reasonable request to the corresponding author.

REFERENCES

Abadi M. G., Navarro J. F., Fardal M., Babul A., Steinmetz M., 2010, *MNRAS*, 407, 435

Allgood B., Flores R. A., Primack J. R., Kravtsov A. V., Wechsler R. H., Faltenbacher A., Bullock J. S., 2006, *MNRAS*, 367, 1781

Bailin J., Steinmetz M., 2005, *ApJ*, 627, 647

Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013, *ApJ*, 762, 109

Bovy J., Bahmanyar A., Fritz T. K., Kallivayalil N., 2016, *ApJ*, 833, 31

Bowden A., Evans N. W., Williams A. A., 2016, *MNRAS*, 460, 329

Bryan S. E., Kay S. T., Duffy A. R., Schaye J., Dalla Vecchia C., Booth C. M., 2013, *MNRAS*, 429, 3316

Bullock J. S., 2002, in Natarajan P., ed., *The Shapes of Galaxies and their Dark Halos*. World Scientific Press, Singapore, p. 109

Butsky I. et al., 2016, *MNRAS*, 462, 663

Cataldi P., Pedrosa S. E., Tissera P. B., Artale M. C., 2021, *MNRAS*, 501, 5679

Chisari N. E. et al., 2017, *MNRAS*, 472, 1163

Chua K. T. E., Pillepich A., Vogelsberger M., Hernquist L., 2019, *MNRAS*, 484, 476

Conroy C., Wechsler R. H., 2009, *ApJ*, 696, 620

Crain R. A. et al., 2015, *MNRAS*, 450, 1937

Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371

Debattista V. P., Moore B., Quinn T., Kazantzidis S., Maas R., Mayer L., Read J., Stadel J., 2008, *ApJ*, 681, 1076

Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497

Dubinski J., 1994, *ApJ*, 431, 617

Dubinski J., Carlberg R. G., 1991, *ApJ*, 378, 496

Emami R. et al., 2021, *ApJ*, 913, 36

Genel S. et al., 2014, *MNRAS*, 445, 175

Hayashi E., Navarro J. F., Springel V., 2007, *MNRAS*, 377, 50

Ibata R., Lewis G. F., Irwin M., Totten E., Quinn T., 2001, *ApJ*, 551, 294

Jing Y. P., Suto Y., 2002, *ApJ*, 574, 538

Katz N., Gunn J. E., 1991, *ApJ*, 377, 365

Katz N., White S. D. M., 1993, *ApJ*, 412, 455

Kazantzidis S., Abadi M. G., Navarro J. F., 2010, *ApJ*, 720, L62

Law D. R., Majewski S. R., 2010, *ApJ*, 714, 229

Leauthaud A. et al., 2012, *ApJ*, 744, 159

Loebman S. R., Ivezic Ž., Quinn T. R., Governato F., Brooks A. M., Christensen C. R., Jurić M., 2012, *ApJ*, 758, L23

Lovell M. R. et al., 2018, *MNRAS*, 481, 1950

Macciò A. V., Dutton A. A., van den Bosch F. C., 2008, *MNRAS*, 391, 1940

Malhan K., Ibata R. A., 2019, *MNRAS*, 486, 2995

Marinacci F. et al., 2018, *MNRAS*, 480, 5113

Naiman J. P. et al., 2018, *MNRAS*, 477, 1206

Navarro J. F. et al., 2010, *MNRAS*, 402, 21

Nelson D. et al., 2015, *Astron. Comput.*, 13, 12

Nelson D. et al., 2018, *MNRAS*, 475, 624

Nelson D. et al., 2019a, *Comput. Astrophys. Cosmol.*, 6, 2

Nelson D. et al., 2019b, *MNRAS*, 490, 3234

Pillepich A. et al., 2018a, *MNRAS*, 473, 4077

Pillepich A. et al., 2018b, *MNRAS*, 475, 648

Pillepich A. et al., 2019, *MNRAS*, 490, 3196

Planck Collaboration XVI, 2014, *A&A*, 571, A16

Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, *MNRAS*, 338, 14

Prada J., Forero-Romero J. E., Grand R. J. J., Pakmor R., Springel V., 2019, *MNRAS*, 490, 4877

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

Schneider M. D., Frenk C. S., Cole S., 2012, *J. Cosmol. Astropart. Phys.*, 05, 030

Shen X. et al., 2020, *MNRAS*, 495, 4747

Shen X., Vogelsberger M., Nelson D., Tacchella S., Hernquist L., Springel V., Marinacci F., Torrey P., 2022, *MNRAS*, 510, 5560

Spergel D. N., Flauger R., Hložek R., 2015, *Phys. Rev. D*, 91, 023518

Springel V., 2010, *MNRAS*, 401, 791

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

Springel V., White S. D. M., Hernquist L., 2004, in Ryder S. D., Pisano D. J., Walker M. A., Freeman K. C., eds, *Proc. IAU Symp. 220, Dark Matter in Galaxies*. Astron. Soc. Pac., San Francisco, p. 421

Springel V. et al., 2018, *MNRAS*, 475, 676

Suresh J., Bird S., Vogelsberger M., Genel S., Torrey P., Sijacki D., Springel V., Hernquist L., 2015, *MNRAS*, 448, 895

Tenneti A., Mandelbaum R., Di Matteo T., Kiessling A., Khandai N., 2015, *MNRAS*, 453, 469

Tissera P. B., White S. D. M., Pedrosa S., Scannapieco C., 2010, *MNRAS*, 406, 922

Velliscig M. et al., 2015, *MNRAS*, 454, 3328

Vera-Ciro C., Helmi A., 2013, *ApJ*, 773, L4

Vera-Ciro C. A., Sales L. V., Helmi A., Frenk C. S., Navarro J. F., Springel V., Vogelsberger M., White S. D. M., 2011, *MNRAS*, 416, 1377

Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031

Vogelsberger M. et al., 2014a, *MNRAS*, 444, 1518

Vogelsberger M. et al., 2014b, *Nature*, 509, 177
 Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020a, *Nat. Rev. Phys.*, 2, 42
 Vogelsberger M. et al., 2020b, *MNRAS*, 492, 5167
 Wang L., Dutton A. A., Stinson G. S., Macciò A. V., Penzo C., Kang X., Keller B. W., Wadsley J., 2015, *MNRAS*, 454, 83
 Warren M. S., Quinn P. J., Salmon J. K., Zurek W. H., 1992, *ApJ*, 399, 405
 Weinberg D. H., Colombi S., Davé R., Katz N., 2008, *ApJ*, 678, 6
 Weinberger R. et al., 2017, *MNRAS*, 465, 3291
 Zaritsky D., Conroy C., Zhang H., Naidu R. P., Bonaca A., Caldwell N., Cargile P. A., Johnson B. D., 2020, *ApJ*, 888, 114
 Zemp M., Gnedin O. Y., Gnedin N. Y., Kravtsov A. V., 2011, *ApJS*, 197, 30

APPENDIX: RESOLUTION CONVERGENCE

In this work, we have primarily considered both the TNG100 and TNG50 runs with the highest resolution, e.g. TNG100-1 and TNG50-1. Here, we test the numerical convergence of the median halo shape profiles using the three resolution levels, and for haloes in different mass bins. The convergence of the shape profiles are shown in Fig. A1 for TNG100 and Fig. A2 for TNG50. Note that each lower resolution level is coarsened by a factor of 8 in mass resolution and a factor of 2 in the softening length. The dotted vertical lines show the median convergence radii r_{conv} as derived from equation (4) across the haloes in the mass bin. We find in the DMO cases (solid lines) good convergence of the median halo shape for $r > r_{\text{conv}}$. In Fig. A1, note that because of the low resolution of TNG100-DM-3, the halo shapes

of $10^{11} M_{\odot}$ haloes (left-hand panels) are not converged at any radii. A similar result can be observed for the smallest haloes in TNG50 (Fig. A2). We also plot the profiles from the smaller L25n512-DMO box (violet) in Fig. A1, which demonstrate good convergence with the TNG100-DM-1 profiles (blue), due to the similar resolutions of both simulations.

For comparison, we plot also results from the MHD runs (dashed curves). In these MHD runs, the halo shapes are not converged with resolution, even in the region $r > r_{\text{conv}}$. Instead, haloes from the higher resolution runs are more spherical, due to the resolution dependence of the baryonic physics. In the TNG galaxy formation model, lowering the simulation resolution results in lower galaxy formation efficiencies (Pillepich et al. 2018b), causing haloes to be less spherical as well.

Compared to circular velocities, the convergence of halo shapes is more demanding, with a larger minimum radius r_{conv} of the converged region. Although $\kappa > 7$ describes well the converged regions of halo shapes in DMO simulations, the situation is more complicated in hydrodynamic simulations since the baryonic physics implementations can also depend on resolution. For example, the TNG galaxy formation model when applied to a larger box of sidelength 300 Mpc results in the systematically lower stellar masses of galaxies compared to the smaller box TNG100 (Pillepich et al. 2018b), due to the lower mass and spatial resolution in TNG300. For this reason, we use equation (4) applied to the DMO runs to identify the regions where the effects of two-body relaxation can be neglected.

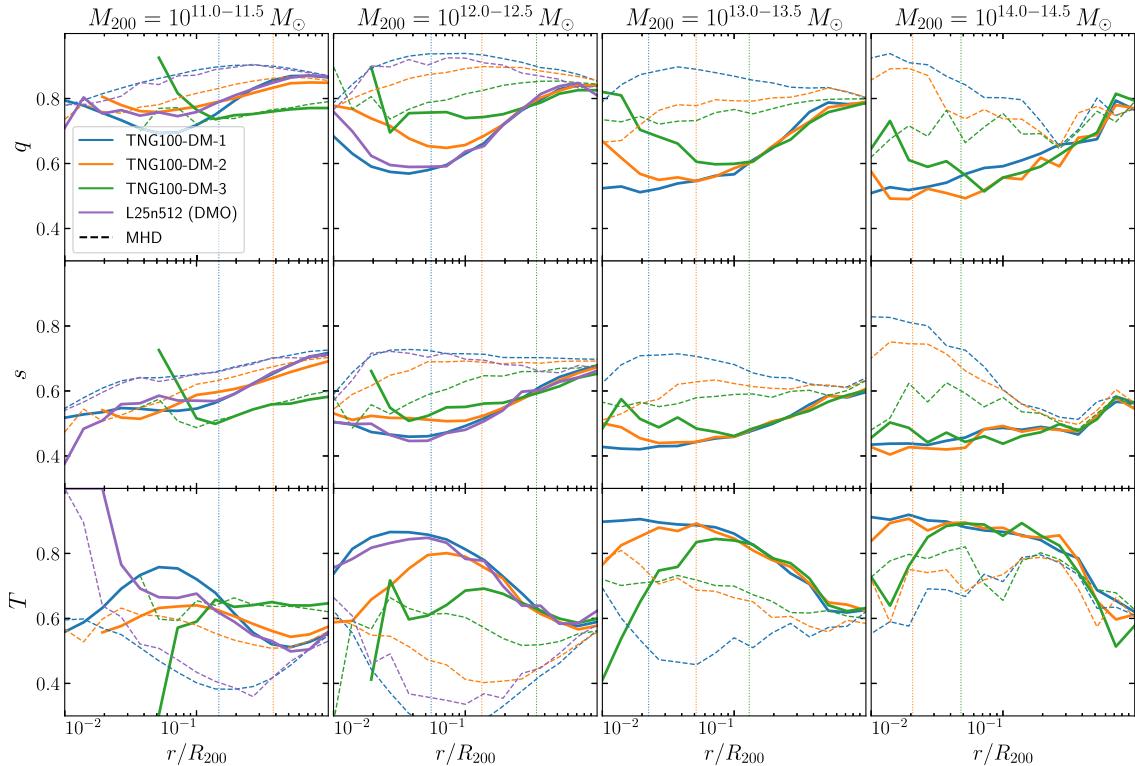


Figure A1. Effect of resolution on the median DM halo shapes in TNG100-DM (solid lines) and TNG100 (MHD, dashed lines) in various mass bins. Dotted lines show the median convergence radii r_{conv} as derived from $\kappa(r) = 7$ (equation 4). For $r > r_{\text{conv}}$, the median DMO halo shapes are converged, i.e. the lower resolution runs match the highest resolution TNG100-DM-1. The halo shapes of $10^{11} M_{\odot}$ haloes (left-hand panels) are not converged at all in the lowest resolution TNG100-DM-3. In the MHD runs (dashed lines), halo shapes do not appear to converge with resolution, due to the additional resolution dependence of the galaxy formation physics. Results from the smaller L25n512 box are also shown for comparison.

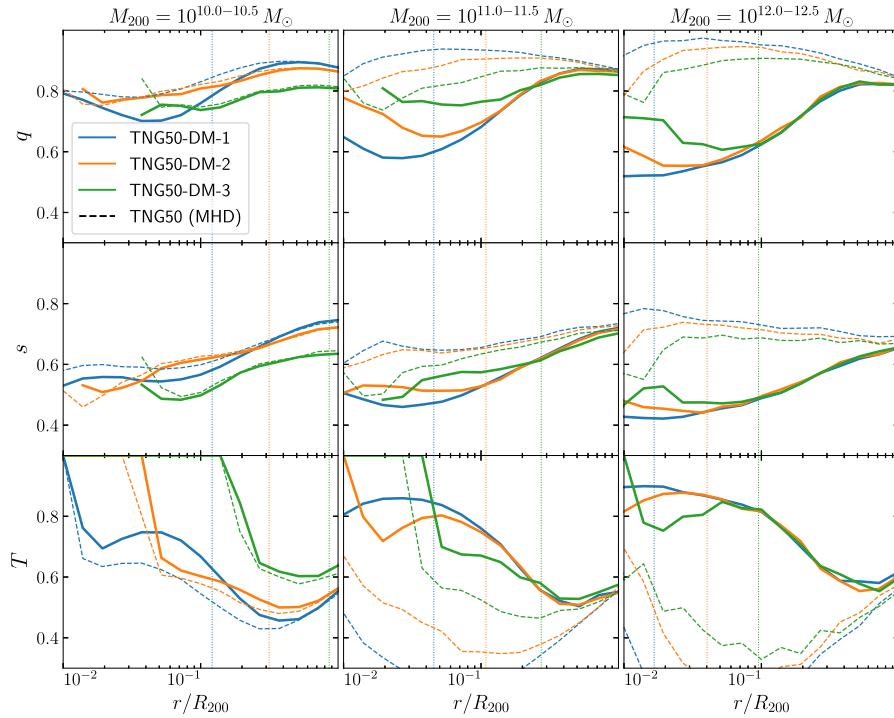


Figure A2. Effect of resolution on the median DM halo shapes in TNG50-DM (solid lines) and TNG50 (MHD, dashed lines), for haloes between 10^{10} and $10^{12.5} M_{\odot}$. In particular, note that $r = 0.15R_{200}$ is resolved in $10^{10} M_{\odot}$ haloes for the highest resolution run.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.