



Modelling galactic conformity with the colour–halo age relation in the Illustris simulation

Citation

Bray, Aaron D., Annalisa Pillepich, Laura V. Sales, Emily Zhu, Shy Genel, Vicente Rodriguez-Gomez, Paul Torrey, et al. 2015. “Modelling Galactic Conformity with the Colour–halo Age Relation in the Illustris Simulation.” *Monthly Notices of the Royal Astronomical Society* 455 (1) (November 2): 185–198. doi:10.1093/mnras/stv2316.

Published Version

doi:doi.org/10.1093/mnras/stv2316

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:34353287>

Terms of Use

This article was downloaded from Harvard University’s DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Modelling galactic conformity with the colour–halo age relation in the Illustris simulation

Aaron D. Bray,^{1★} Annalisa Pillepich,¹ Laura V. Sales,² Emily Zhu,^{3,4} Shy Genel,^{5†} Vicente Rodriguez-Gomez,¹ Paul Torrey,^{6,7} Dylan Nelson,¹ Mark Vogelsberger,⁶ Volker Springel,^{8,9} Daniel J. Eisenstein¹ and Lars Hernquist¹

¹Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

²University of California Riverside, 900 University Avenue, Riverside, CA 92521, USA

³Phillips Academy, 180 Main Street, Andover, MA 01810, USA

⁴Harvard College, Cambridge, MA 02138, USA

⁵Columbia University, 116th Street and Broadway, New York, NY 10027, USA

⁶Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

⁷TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

⁸Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

⁹Zentrum für Astronomie der Universität Heidelberg, ARI, Monchhofstr. 12-14, D-69120 Heidelberg, Germany

Accepted 2015 October 2. Received 2015 October 1; in original form 2015 August 24

ABSTRACT

Comparisons between observational surveys and galaxy formation models find that dark matter haloes’ mass can largely explain their galaxies’ stellar mass. However, it remains uncertain whether additional environmental variables, known as assembly bias, are necessary to explain other galaxy properties. We use the Illustris simulation to investigate the role of assembly bias in producing galactic conformity by considering 18 000 galaxies with $M_{\text{stellar}} > 2 \times 10^9 M_{\odot}$. We find a significant signal of galactic conformity: out to distances of about 10 Mpc, the mean red fraction of galaxies around redder galaxies is higher than around bluer galaxies *at fixed stellar mass*. Dark matter haloes exhibit an analogous conformity signal, in which the fraction of haloes formed at earlier times (old haloes) is higher around old haloes than around younger ones *at fixed halo mass*. A plausible interpretation of galactic conformity is the combination of the halo conformity signal with the galaxy colour–halo age relation: at fixed stellar mass, particularly towards the low-mass end, Illustris’ galaxy colours correlate with halo age, with the reddest galaxies (often satellites) preferentially found in the oldest haloes. We explain the galactic conformity effect with a simple semi-empirical model, assigning stellar mass via halo mass (abundance matching) and galaxy colour via halo age (age matching). Regarding comparison to observations, we conclude that the adopted selection/isolation criteria, projection effects, and stacking techniques can have a significant impact on the measured amplitude of the conformity signal.

Key words: galaxies: formation – galaxies: haloes – cosmology: theory – dark matter.

1 INTRODUCTION

Previous investigations of the demographics and distribution of dark matter haloes in a cold dark matter universe have found that the clustering properties of these haloes have a dependence on formation time, in addition to the more significant dependence on halo mass

(Gao, Springel & White 2005; Wechsler et al. 2006; Croton, Gao & White 2007; Li, Mo & Gao 2008). However, current observational frameworks for analysing the luminosity- and colour-dependent clustering of galaxies do not take into account this halo assembly bias (e.g. Zehavi et al. 2011). Rather, they use models that assume that galaxy clustering statistics can be modelled solely based on the mass of the halo (e.g. Berlind & Weinberg 2002; Yang, Mo & van den Bosch 2003; Conroy, Wechsler & Kravtsov 2006). This would be a correct assumption, as long as galaxy properties such as stellar mass and specific star formation rate (sSFR) are not also correlated

* E-mail: abray@cfa.harvard.edu

† Hubble Fellow.

with other dark matter properties at fixed halo mass. Otherwise, ignoring the effects of properties other than mass may lead to biased interpretations of the observational results (Zentner, Hearin & van den Bosch 2014).

Recent observations at low redshift have found a signal of *galactic conformity* in which the sSFR and gas fractions of neighbouring galaxies correlate with the respective properties of the central galaxy, both within and beyond the virial radius (e.g. Weinmann et al. 2006; Kauffmann et al. 2013; Lacerna, Padilla & Stasyszyn 2014; Hartley et al. 2015; Knobel et al. 2015). Such observations suggest that those galaxy properties may indeed be correlated with halo properties beyond mass, such that *halo* assembly bias may lead to *galaxy* assembly bias. Simultaneously, though, other observations have not found these same manifestations of galaxy assembly bias (e.g. Tinker, Wetzel & Conroy 2011; Lin et al. 2015). Moreover, there is considerable debate as to what the role of central and satellite galaxies play in the emergence of this signal (Knobel et al. 2015). Some models treat centrals and satellites identically (Hearin, Watson & van den Bosch 2015a), while others have satellites colours correlate directly with group-wide properties, such as halo concentration (Paranjape et al. 2015). Finally, the debate over the role of internal (e.g. Hartley et al. 2015) versus external (e.g. Hearin, Behroozi & van den Bosch 2015b) quenching mechanisms, and thus also the extent to which conformity is a product of assembly bias, relies heavily on the observed amplitude and radius out to which the conformity signal is observed (Knobel et al. 2015; Paranjape et al. 2015).

Semi-analytic models (Guo et al. 2011) can qualitatively produce the galactic conformity effect seen in observations, but it has been argued that such theoretical effects are not as large as in observations (Kauffmann et al. 2013; Hearin et al. 2015b). The same qualitative conformity signal can also be reproduced using semi-empirical halo occupation models (Hearin & Watson 2013; Hearin et al. 2015b; Watson et al. 2015) or with tunable extensions to the halo occupation distribution (HOD) framework (Paranjape et al. 2015).

In this paper, we investigate the presence of galactic conformity in Illustris (Vogelsberger et al. 2014b), a state-of-the-art cosmological simulation with full hydrodynamical and subgrid physics run with the AREPO code (Springel 2010). Galactic conformity has yet to be probed in a hydrodynamical simulation, given the limitations thus far in encompassed volumes, numerical resolution, and realism and statistical significance of the simulated galaxy populations. Illustris, on the other hand, combines a $75 h^{-1}$ Mpc per side cosmological volume at kpc resolution with a population of thousands of galaxies which compare well to observational constraints. By studying galaxy clustering in Illustris, we see whether a statistically significant galactic conformity signal arises in a realistic simulation of galaxy formation, and in particular, whether the conformity can be explained solely by differences in the halo masses of red and blue galaxies, or whether the additional information about the assembly history of the haloes is required. Furthermore, we explore the role that possible observational biases and selection criteria will have on the conformity signal.

This paper is organized as follows. In Section 2, we briefly review the properties of the Illustris simulation, describe our selection criterion, and explain how we calculate dark matter halo ages from the merger trees. We present the detection of both galactic and halo conformity in Section 3, and show the presence of a colour–halo age relation in Section 4. In Section 5, we apply abundance and age matching models from the literature to the Illustris galaxies to show how galactic conformity naturally arises in Illustris, and we discuss the differential importance of centrals and satellites, as well as the

effect of other observational choices on the strength and radial dependence of the conformity signal. We conclude and summarize in Section 6.

2 SIMULATIONS AND METHODS

In this paper, we use the Illustris suite, a set of simulations which form galaxies self-consistently, by combining an N -body treatment of gravity with the hydrodynamical, moving-mesh code AREPO (Springel 2010) to follow gas. AREPO solves the Euler equations on an unstructured Voronoi tessellation, in which the mesh-generating points advect with the baryonic flow. The code includes relevant physical processes such as gas cooling (Katz, Weinberg & Hernquist 1996), a photoionizing background (Faucher-Giguère et al. 2009), star formation (Springel & Hernquist 2003), black hole seeding and feedback (Di Matteo, Springel & Hernquist 2005; Springel, Di Matteo & Hernquist 2005; Sijacki et al. 2007), and chemical enrichment (Wiersma et al. 2009). Full details of the applied galaxy formation and feedback model are described in Vogelsberger et al. (2013) with multi-epoch galaxy population properties being tested and presented in Torrey et al. (2014). These simulations reproduce realistic populations of galaxies, as demonstrated in previous Illustris analyses (Genel et al. 2014; Vogelsberger et al. 2014a,b). While the suite includes realizations with different box sizes and at different resolutions, our primary results presented here are based on the highest resolution run (Illustris-1), where a $75 h^{-1}$ Mpc cosmological box is evolved from $z = 127$ to $z = 0$ with initial conditions consistent with WMAP-9 (Hinshaw et al. 2013). The mass resolution for the dark matter is $m_{\text{DM}} = 6.26 \times 10^6 M_{\odot}$, and for baryons it is roughly $m_b \sim 1.26 \times 10^6 M_{\odot}$. At $z = 0$, the softening lengths are roughly 1.42 kpc for dark matter particles and 0.71 kpc for stellar particles, being smaller at higher redshifts, and the hydrodynamics follows gas down to cell sizes as small as 48 pc.

All the data from the Illustris project and associated documentation is now publicly available (Nelson et al. 2015).¹

2.1 Galaxy sample and definitions

Haloes and subhaloes in Illustris are identified using the FOF and SUBFIND algorithms (Davis et al. 1985; Springel et al. 2001; Dolag et al. 2009) at 136 snapshots in time.

In what follows, we work exclusively at redshift $z = 0$ and select a sample of galaxies by imposing $M_{\text{stellar}} > 2 \times 10^9 M_{\odot}$ (corresponding to a minimum of roughly 2000 stellar particles or 2800 stellar, dark-matter or gas elements). The sample includes both central and satellite galaxies, with satellites being SUBFIND subhaloes which are members of their parent FOF group regardless of their distance from the FOF centre. Thus, in our parlance, central galaxies include field galaxies with no satellites of their own, and we call any SUBFIND-identified object a *halo* unless the distinction between haloes and subhaloes is relevant. Moreover, all galaxy properties (stellar masses, SFRs, colours) are derived from SUBFIND-identified stellar particles or cells within twice the stellar half-mass radius of the galaxy under consideration. Halo or total masses are defined as the peak mass of each halo’s mass accretion history (see Section 2.2 for details), including all gravitationally bound resolution elements.

In Fig. 1, the colour–magnitude diagram and the sSFR distribution as a function of stellar masses are given for Illustris galaxies at the current epoch (both centrals and satellites; see also Vogelsberger

¹ <http://www.illustris-project.org>




Figure 1. The galaxy population in Illustris. Left: the colour–magnitude diagram for the $M_{\text{stellar}} > 2 \times 10^9 M_{\odot}$ galaxy sample used in this work. The solid black line divides the red and blue populations (shown in corresponding colours). Contours for 50 and 95 per cent inclusion in absolute M_r versus $g - r$ colour space are shown for our 18 243 galaxies. To allow for a statistically meaningful comparison of redder and bluer galaxies, we divide the sample such that ‘green valley’ galaxies are grouped with red sequence galaxies, despite having some ongoing star formation. Right: sSFRs as a function of stellar mass. We show how the red–blue colour–luminosity cut in Illustris translates into the sSFR–stellar mass plane. While the 50 and 95 per cent inclusion contours correspond to the distribution of points in the sSFR–stellar mass plane, the colours are inherited from the left-hand plot. For clarity, we randomly plot galaxies with no SFR randomly distributed about $2 \times 10^{-13} \text{ yr}^{-1}$.

et al. 2014a; Sparre et al. 2015). These are in qualitative agreement with observations except for the lack of a clear bimodality between red and blue galaxies, and an overpopulation of the green valley and the blue cloud with respect to the red sequence. Encouragingly, the colour distribution of *satellite* galaxies alone is in good agreement with observations (Sales et al. 2015).

For the purposes of our conformity analysis, we divide the selected galaxies into binary *red* and *blue* subsamples, rather than using a continuous distribution of sSFR as a proxy for conformity (see Kauffmann et al. 2013). To the extent that the full distribution of sSFR in Illustris differs from observations, we believe such a binary division better allows us to investigate the emergence of a conformity signal in Illustris, and the effects of observational choices on the observed signal. Unless otherwise stated, our cut will be based on stellar colours, as follows:

$$\text{Red : } (g - r)_{\text{galaxy}} > 0.04(r + 20) + 0.43 \quad (1)$$

$$\text{Blue : } (g - r)_{\text{galaxy}} < 0.04(r + 20) + 0.43. \quad (2)$$

As shown in the right-hand panel in Fig. 1, this corresponds to a slightly increasing cut in sSFR with increasing stellar mass, with some star formation still ongoing in the red population, especially at higher masses. Moreover, in the following, we will use the terms *red* (*blue*) and *quenched* (*star-forming*) galaxies interchangeably. Finally, the fraction of red galaxies is a strong function of stellar mass, and satellite galaxies are more often red than their analogue central galaxies at similar masses, in agreement with observational findings (see left-hand panel of Fig. 2).

Following observational practice (e.g. Kauffmann et al. 2013; Hearin et al. 2015a), in the following sections, we will also adopt an isolation criterion for our sample of galaxies and divide them into *primaries* and *secondaries*. Unlike observations, we have full spatial knowledge of our galaxies, and so we define the isolation criterion based on the 3D real-space locations within the simula-

tion volume, rather than relying on a 2D projection and redshift-space cut for the line-of-sight dimension. An Illustris galaxy is *isolated* if, given its stellar mass M_{stellar} , no other galaxy with stellar mass greater than $M_{\text{stellar}}/2$ is present within a 3D distance of 500 kpc. The *primary* sample is constituted by those galaxies with $\log M_{\text{stellar}} > 9.61 \approx 4.07 \times 10^9 M_{\odot}$ that are also isolated. The adopted mass cut is the minimum for which we can consistently apply the isolation criterion, given our minimum mass threshold of $M_{\text{stellar}} > 2 \times 10^9 M_{\odot}$. We use the term *secondary* or *neighbouring* to refer to all galaxies in the vicinity (in our case, out to 10 Mpc) of a primary galaxy.

The purpose of the isolation criterion is to reduce the number of interloping satellite galaxies in the primary sample when the distinction between centrals and satellites is not available (as often is the case in observations). This can be seen in Fig. 2, left-hand panel, where the isolation criterion serves to lower the satellite fraction from ~ 30 down to about 10 per cent, with Illustris total satellite fraction falling in between the estimates from semi-analytical models and observations (Kauffmann et al. 2013; Guo et al. 2011; Wang et al. 2013, at higher masses, the satellite fraction drops quickly). While a secondary galaxy need not be isolated, it may be. Thus, secondary galaxies around a particular primary galaxy may themselves be members of the primary sample, since the conformity signal is measured out to radii well beyond the 500 kpc radius used for the isolation criterion.

For the conformity itself, we measure the quenched fraction of secondary galaxies in every 3D real-space radial bin R around primaries. We then report the mean value of this red fraction for primary galaxies that have at least one galaxy in radial bin R . Thus, if a particular primary has no neighbouring secondary galaxy in a particular radial bin, this does not count towards the mean. Another way of saying this is that the mean red fraction of primary galaxies is not equivalent to the red fraction of the stack of all primary galaxies. This distinction means that conformity, as we measure it, will equally weight galaxies with a few satellites and galaxies with

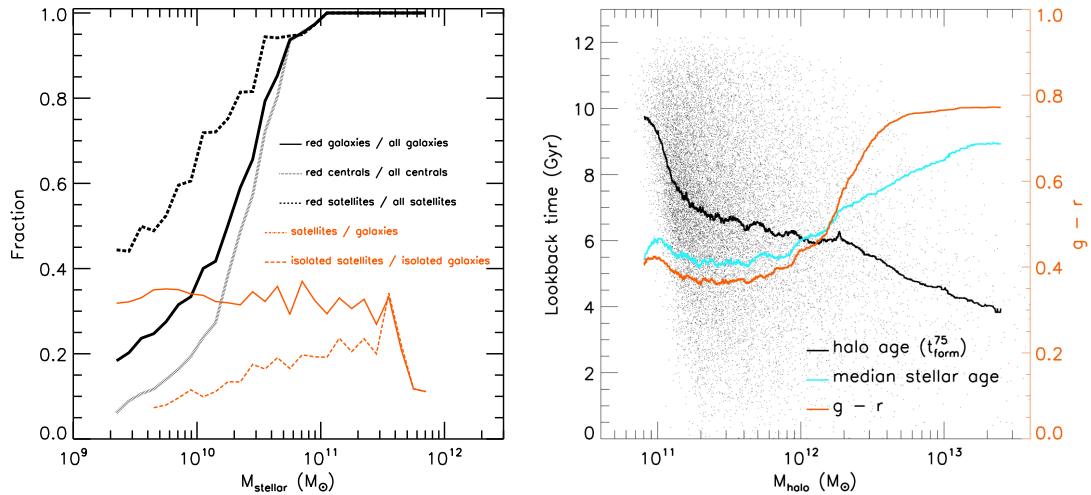


Figure 2. Left: the red fraction and satellite fraction of galaxies as a function of stellar mass. We note that the isolation criterion has the effect to reduce the satellite fraction from ~ 30 to ~ 10 per cent, although the effect is stronger at lower masses. Both centrals and satellites show a strong relationship between red fraction and halo mass; however, especially at lower masses, satellites are significantly more quenched than centrals. Right: halo age (i.e. look-back halo formation time) as a function of halo mass, alongside both the colour and median stellar age as a function of halo mass. Despite the fact that lower mass haloes are older than their more massive counterparts, they host galaxies with younger and bluer stellar populations.

many satellites, rather than letting the signal be dominated by a few primary galaxies with the highest halo-to-stellar mass ratios. We report our galactic conformity results in three bins in primary stellar mass and three bins in primary halo mass, so that we can discern effects on the conformity signal caused by differences in stellar-to-halo mass relation of red versus blue galaxies.

2.2 Halo merger tree and assembly histories

To follow the evolution of individual haloes and galaxies, we use the `SUBLINK` merger tree catalogues from Rodriguez-Gomez et al. (2015). These merger trees provide the evolution of any `SUBFIND` property along the main branch of all haloes and galaxies at $z = 0$. Specifically, the *main branch* is defined as the sequence of progenitors with the most massive history behind them (rather than the sequence of progenitors which maximize the mass at every time step). While different definitions of main branch are on average consistent, the addition utilized here provides a safeguard against spurious defects in the halo finding algorithms, such as subhalo swapping.

We use the total `SUBFIND` mass accretion histories to calculate the halo formation time or halo age, t_{form} , of every halo, as well as the halo mass. In practice, for every object within our sample at $z = 0$, we first run the sequence of masses at subsequent snapshots through a median box filter of full width of five snapshots (or three, if fewer snapshots exist), and then we spline this mass accretion history to obtain a fine-grained mass evolution as a function of redshift. Moreover, in order to avoid spurious identifications, we require that each object has existed as either a central or a satellite for at least three consecutive snapshots. The halo formation time t_{form}^x is the earliest moment in cosmic time at which the splined total mass accretion history reaches x per cent of the peak mass of a halo (we usually express it here in terms of look-back time from the present day, in Gyr). The halo mass is the maximum mass value reached along the main branch: for central haloes the peak mass is usually very close to their mass at $z = 0$ and provides a reasonable approximation of the virial mass, generally overestimating it by roughly 10 per cent; for satellite subhaloes, the current-epoch mass is usually much lower

than the peak mass, because of mass-loss due to stripping after accretion on to the parent haloes. This procedure ensures that we have a standard definition of mass and halo formation time that is identical for both central and satellite haloes; however, by construction, subhaloes' ages will always be biased high compared to central haloes' ages.

In what follows, we will adopt various choices for the halo formation time, with, e.g. t_{form}^{25} , t_{form}^{50} , and t_{form}^{75} being the age at which a halo has assembled 25, 50, and 75 per cent of its peak halo mass. More massive haloes formed more recently than lower mass haloes (see Fig. 2, right-hand panel, black curve; and also e.g. Wechsler et al. 2002). However, the stellar populations of galaxies residing in more massive haloes at $z = 0$ are *older* than the stellar population of galaxies residing in less massive haloes (Fig. 2, right-hand panel, cyan curve), or equivalently, redder (orange curve – see e.g. Heavens et al. 2004; Jimenez et al. 2005; Nelan et al. 2005; Thomas et al. 2005 for the first observational claims of archaeological downsizing).

3 GALACTIC AND HALO CONFORMITY

3.1 Galactic conformity

We now present the measurements of galactic conformity in the Illustris simulation. In Fig. 3, we plot the mean red fraction of secondary galaxies around their isolated primaries (see definitions and methods in Section 2.1). Each bin is a spherical shell of width 500 kpc, centred on real-space distances from $r = 0.75$ Mpc to $r = 9.75$ Mpc. Red and blue squares are used to represent red and blue primaries, respectively, as defined in Section 2, and the error bars define the standard error on the mean, as determined from 1000 bootstrap resamplings. The top panel shows the results in three bins in stellar mass for the primary sample, while the lower panel shows the results in three bins in total halo mass. The red (blue) primaries, by stellar mass bin, low to high, have 655 (2921), 1114 (1639), and 1171 (191) galaxies, respectively. The red (blue) primaries, by halo mass bin, low to high, have 559 (3195), 602 (1690), and 721




Figure 3. Galactic conformity in Illustris, i.e. mean red fraction of neighbouring galaxies as a function of real-space distance in 500 kpc bins around primary galaxies divided in three different stellar (halo) mass bins in the top (bottom) rows. The red solid line is the fraction of red galaxies around the red population of primary galaxies, while the blue solid line is the fraction of red galaxies around the blue population of primaries. The galactic conformity signal is present at all distances, but it is particularly strong at $r < 3$ Mpc. Standard errors on the means are calculated from bootstrapping. Here, the separation between red/quenched and blue/star-forming populations is based on a colour cut (see Section 2.1). Moreover, the primaries are isolated galaxies which are mostly centrals, with roughly 20 per cent fraction of satellite interlopers (see Fig. 2); secondaries can be either centrals or satellites, isolated or not.

(654) galaxies. Note that secondary galaxies can have any mass in all panels.

We clearly see that red primaries have a higher fraction of red neighbours than their bluer counterparts. Furthermore, we see two significant trends. First, there is a near-field, higher amplitude conformity signal out to roughly 3 Mpc, and then a plateau of a far-field effect that extends out to at least 5 Mpc. Secondly, lower mass primaries have both a higher amplitude conformity effect, in both the near and far fields, and the far-field effect continues out to larger radii, remaining present out to 10 Mpc in the lowest mass bin. We have confirmed that the signal disappears entirely by 15 Mpc in all cases. We note that for such low-mass galaxies, the virial radii of the primaries are significantly lower than even the radii at which we see the near-field effect (the typical virial radii spanning from 150 to about 370 kpc across the three adopted mass bins).

One possibility for the presence of the conformity signal is that red and blue galaxies, selected in fixed stellar mass bins, are none the less hosted by halo masses of substantially different size. In this case, the higher red fraction around red primaries would be due to a halo-mass quenching effect. However, as shown by the lower panel of Fig. 3, selecting our primary galaxies in halo mass bins does not reduce the amplitude of our signal. Thus, clustering observations that use stellar mass as a proxy for halo mass are unlikely to be biased by this selection technique. We return to the role of mass in the conformity signal in Section 5.

We have also tested these results for robustness to different definitions of quenched/red galaxies. We have used both sSFR and stellar ages, in which quenched galaxies are defined either as being below the median sSFR or above the median stellar age; as having sSFR lower than finite values across stellar mass (e.g. $3, 6, 8 \times 10^{-11} \text{ yr}^{-1}$); or in which quenched galaxies are defined using sSFR or stellar age, but at a cut-off to reproduce the same red to blue ratio

we have in our fiducial definition. In all cases, a galactic conformity remains, including the near- and far-field distinction: namely, the specific cut to separate red versus blue galaxies does not affect whether there exists a statistically significant signal in a particular mass bin. However, the exact location of, for example, a flat cut in sSFR may affect the presence and magnitude of a dependence of the conformity signal on primary mass. For example, choosing a constant cut in sSFR that is unreasonably high would cause some less massive blue galaxies that are part of ‘blue cloud’ in Illustris to be classified as red, which would inflate the overall red fraction around more massive galaxies. In relation to this consideration, we argue that a reasonable colour cut should follow the general number density contours in colour–luminosity space, as opposed to, for example, splitting the blue cloud. Yet we caution that whether or not the conformity signal exhibits a trend with primary mass may depend on the adopted split between star-forming and quenched galaxies.

3.2 Dark matter halo conformity

In Fig. 4, we now present the dark-matter halo conformity effect in the Illustris simulation. In order to be able to compare the relative amplitude of the halo and the galaxy conformity effects, we split Illustris haloes into two samples of haloes, old and young, which are chosen to match the sample sizes of the red and blue populations. Here, we adopt the t_{form}^{75} definition for halo formation time. In the case of Illustris-1, the global split between old and young haloes occurs at $z = 0.965$.

We keep the same isolation criterion as before defined by stellar mass, but now the binning of primary galaxies is by halo mass. In this way, we check only for the influence on the conformity signal of using stellar versus halo mass, rather than introducing another

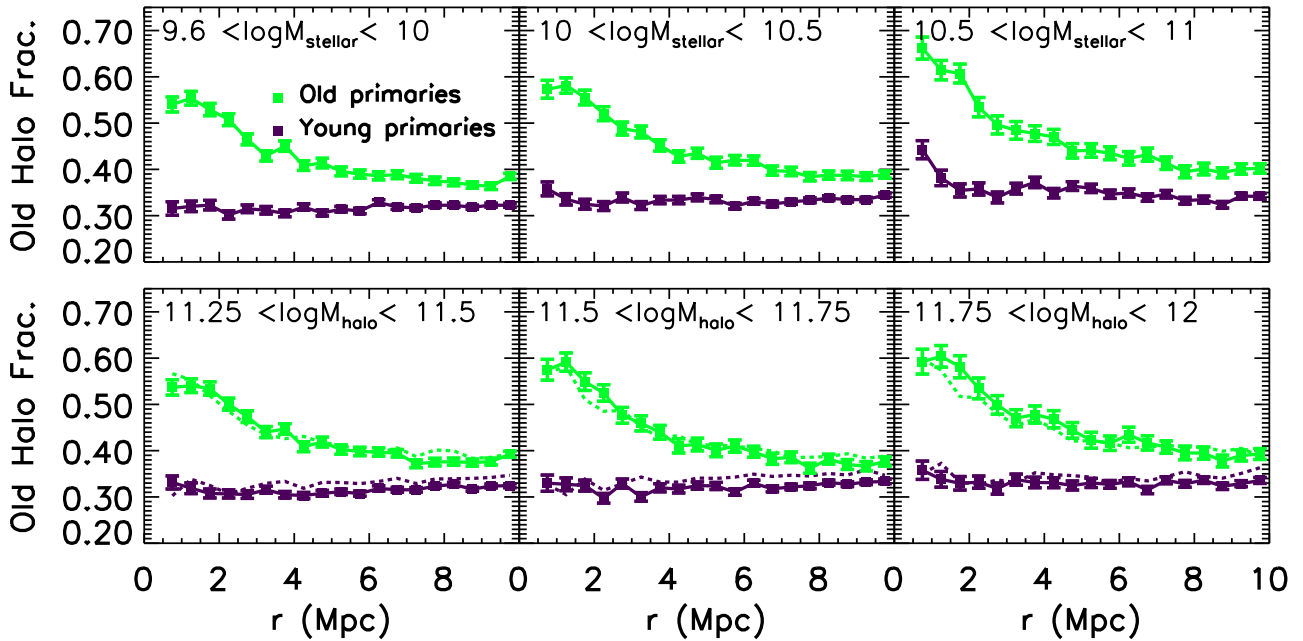


Figure 4. Halo conformity in Illustris. We show how the effect in Fig. 3 is paralleled by the conformity between old dark matter haloes. We keep the same isolation criterion, but now, instead of looking at the red fraction around red (blue) galaxies, we plot the fraction of old haloes around old (young) haloes (where halo age is based on t_{form}^{75}). We see the same qualitative signal as before. The split between old and young haloes is defined such that the number of old haloes is the same as the number of red galaxies; this ensures that the relative amplitudes of the conformity signals are the same. In the top row, the stellar mass cuts are for the galaxies that are in each subhalo. In the bottom row, we also overplot as dotted lines the results from Illustris-Dark-1, the dark matter only run with identical initial conditions (without error bars, for clarity). The minimum subhalo mass of secondary subhaloes used to measure the conformity signal in the Illustris-Dark-1 run is $8.8 \times 10^{10} M_{\odot}$, and we set the division between old and young subhaloes to match the ratio in the Illustris-1 hydrodynamical run.

difference due to a new isolation criterion. For the top panel, the stellar masses refer to the galaxies that inhabit those haloes, as we make the measurements in the fiducial Illustris-1, so haloes and galaxies have direct counterparts.

A dark-matter halo conformity signal is detected. Old haloes are preferentially surrounded by other old haloes. Similarly, there is the same near- and far-field split as in the galactic effect, namely the halo conformity signal is strongest at separations < 3 Mpc. Halo assembly bias, and its effect on clustering, has been studied for many years in other dark matter-only simulations and with semi-analytic models (e.g. Gao et al. 2005; Wechsler et al. 2006; Croton et al. 2007; Dalal et al. 2008; Lacerna & Padilla 2011; Wang et al. 2013). Thus, it is not particularly novel that we find such a signal in Illustris. What is notable is that, as the Illustris simulation has full baryonic physics, we can compare the qualitative shape of the galactic conformity, as seen in Fig. 3, with the halo conformity as seen here in Fig. 4.

Overall, the galaxy and halo conformity signals are qualitatively similar. Both show a higher amplitude effect out to ~ 3 Mpc, and the signal plateaus at larger radii. One difference is that the halo conformity signal shows no obvious decline with increasing halo mass. At first glance, this is in contradiction with the mass dependence seen first in Gao et al. (2005), in which the relative bias of the oldest to the youngest haloes was a function of mass. That is, the relative bias was found to be higher at lower masses. Three differences likely account for this apparent discrepancy. First, the range of halo mass that we probe directly in this analysis, from $\log M_{\text{halo}} = 11.25$ to $\log M_{\text{halo}} = 12$, is much smaller than that probed by Gao et al. (2005). Our mass range corresponds roughly to the four lowest mass bins in Fig. 3 of that work, over which the relative bias is nearly constant. Any remaining difference might be attributable to the facts

that (a) our split between young and old haloes is global and not halo mass dependent; and (b) the bias there is measured between 6 and $25 h^{-1}$ Mpc, whereas most of our measurements lie at smaller radii, $\lesssim 10$ Mpc.

Another difference between the galactic and halo conformity signals is that the decline in old fraction with radius is more gradual than the decline in red fraction, with the sharp feature seen especially prominently in the lowest mass bin not present in Fig. 4. This suggests that there could be a particular scale at which quenching becomes efficient. This could be related to internal processes, such as stellar or AGN feedback, akin the one halo effect offered by Hartley et al. (2015), or it could be due to external processes, such as the influence of the tidal radii of neighbouring haloes (Dalal et al. 2008; Hearin et al. 2015b). Finally, for the halo conformity signal, the old fraction around young and old haloes does not return to the global average until almost 15 Mpc, whereas in the galactic conformity case, the fractions are equal by 6–10 Mpc, depending on mass bin.

These results are robust to the effect of baryonic physics: indeed, we confirm the presence of an equivalent halo conformity signal in the Illustris-1-Dark simulation. This is overplotted for old and young primaries in the lower panels of Fig. 4 as dashed lines (green and purple, respectively). The Illustris-1-Dark simulation was run with the same initial conditions and box size as Illustris-1, but with only dark matter. Because this run does not have stellar masses, we cannot keep the identical isolation criteria for this test. We instead define a halo mass cut of $8.8 \times 10^{10} M_{\odot}$, which is more than two times below the minimum primary halo mass in our sample. We can thus apply an equivalent isolation criterion to our dark matter haloes – we require that primary haloes have no neighbour within 500 kpc that has more than half the primary’s halo mass – and

obtain a similar sample size with which to compare results between Illustris-1 and Illustris-1-Dark. We set the division between old and young haloes to match the global ratio of old to young haloes in the Illustris-1 run.

The results show that the halo conformity signal is also present in the dark matter-only run. Specifically, the shape of the halo conformity signal is essentially identical between the Illustris-1 and Illustris-1-Dark runs. We thus find that this measure of halo assembly bias is not affected by the inclusion of baryons via hydrodynamics and feedback.

4 THE COLOUR-AGE RELATION

Having seen that Illustris exhibits both a galactic and a halo conformity, and that they appear to be similar in qualitative shape and amplitude, we turn our attention to the connection between the two. In this section, we emphasize the relationship between galaxy colour and halo age, which may explain how the baryonic signal arises from the dark matter one. As Sparre et al. (2015) showed for central galaxies in Illustris, at least some measures of formation time (in that case, when half the stellar mass is formed) correlate well with galaxy colour. In this work, we examine formation times based on the build-up of the full halo, including mass from dark matter, stars, and gas. The total mass is chosen because the underlying ansatz of basic HOD models is that mass is the only contributing factor to galaxy occupation. By adding in the halo formation time, we simply add one additional parameter, which attempts to encompass the time evolution of that halo mass. We do not argue that any particular formation time is a priori superior, but we will show subsequently that not only are the correlations different, but some do a better job at reproducing the measured two-point statistics.

In Fig. 5, we show that this correlation between galaxy colour ($g - r$) and halo age, t_{form}^{75} , exists in the Illustris simulation. We see that older haloes tend to host redder galaxies. However, the trend is split into two distinct regions. Namely, the reddest galaxies exist almost exclusively in the oldest haloes, but the reverse is not true; there are old haloes that contain bluer galaxies. This creates a sharp break in the correlation. As can be seen by the running median (solid black line), in the lower stellar mass bins, however, there is still a weak correlation between halo age and colour even among the bluer galaxies. The trend is no longer visible at increasingly high stellar masses. In the two highest mass bins, there is substantial scatter, with a population of red galaxies hosted by young haloes. The halo masses of galaxies in these stellar masses range from roughly $3.3 \times 10^{11} M_{\odot}$ to $1.25 \times 10^{13} M_{\odot}$. Especially at the high halo mass end, internal quenching may be responsible for destroying the tighter correlation seen at lower stellar masses. In particular, radio-mode feedback from the AGN will be a contributing factor to the colour of the galaxies.

To guide interpretation, we also show the maximal correlation between galaxy colour and halo age (orange dashed line). This is similar to the ‘age matching’ of Hearin et al. (2015a). For each 0.25 dex stellar mass bin, we rank order the galaxies by our halo age. We then assign the reddest galaxy in that stellar mass bin to the oldest halo in that stellar mass bin. We continue this process, such that the bluest galaxy is eventually assigned to the youngest halo in this mass bin. We then repeat this process for each mass bin. If the relationship between colour and halo age were in fact monotonic, with redder galaxies always present in older haloes, we would expect the galaxies in Illustris to lie along the overplotted, orange dashed line. Instead we see that, due to scatter in the correlation, older haloes

are bluer and younger haloes are redder than the most simplistic model would predict.

Satellite galaxies serve to enhance the overall trend in the colour–halo age plane. By construction, and because of their infall times, they tend to populate the older halo end rather than the young one. Because they also have a higher red fraction at fixed halo (or stellar) mass, this drives the sharp upward shape of the correlation. Certainly, though, the functional form and scatter at fixed halo age of the colour–age relation depends on the exact choice of the definition for halo formation time: with t_{form}^{75} , for example, values necessarily extend to more recent halo formation times than the analog t_{form}^{25} values. An equivalent relation exists in Illustris between halo age and sSFR, and halo age and stellar age.

5 DISCUSSION

5.1 Modelling and interpretation

The parsimonious explanation for why the halo and galaxy conformity signals are qualitatively the same is to invoke a correlation between halo age and galaxy colour. If such a relationship holds, then in the presence of dark-matter halo conformity based on halo formation time (i.e. halo age), we would expect a colour-dependent galaxy conformity signal to necessarily arise. In the previous sections, we have shown that this is indeed the case in the Illustris hydrodynamical simulation, with the galactic conformity signal resulting from the combination of the aforementioned two effects.

First, initial conditions, plus time evolution under gravity, produces halo conformity. Studies in dark matter only simulations (Keselman & Nusser 2007; Dalal et al. 2008; Hearin et al. 2015b) provide possible physical mechanisms for this effect that we do not explore, since we are concerned primarily with the role played by baryons. Secondly, due to the correlation between stellar mass and dark matter mass build-up, Illustris naturally exhibits a colour–halo age relation at $z = 0$. Together, they produce a galactic conformity effect.

To test whether this qualitative picture works in practice, we present in Fig. 6 the measured galactic conformity signal with the same method as used in Fig. 3, except that now we have reassigned the galaxies to different haloes according to a toy semi-empirical model. The fiducial model (‘No Scatter, t_{75}^{75} ’) works as follows. First, we assign galaxies to dark matter haloes via subhalo abundance matching (Conroy et al. 2006; Vale & Ostriker 2006), with a scatter of $\sigma_M = 0.18$, which falls in between various literature values (e.g. Han et al. 2015; Zu & Mandelbaum 2015) and which is roughly consistent with the actual scatter at fixed halo mass in the stellar-to-halo mass Illustris relation (this varies within 0.13–0.23 over the range $\log M_{\text{halo}} = [11.3, 13.5]$). Next – following the same basic procedure as Hearin & Watson (2013), but using our 75 per cent formation time, t_{form}^{75} – in bins of 0.25 dex in $\log M_{\text{stellar}}$, we rank galaxies according to colour, and we assign the reddest galaxy to the oldest halo. This would be equivalent to the orange line in Fig. 5, except that, rather than using stellar masses directly from Illustris, we first perform subhalo abundance matching.

The results from the model are shown as solid lines in Fig. 6, overplotted to the points from Fig. 3. It bears mentioning first that when we perform only abundance matching – and no age matching of any kind – and thus erase any connection between colour and halo age, we obtain a null result, in which the red fraction is indistinguishable for red and blue primaries in all mass bins. On the other hand, when we add in age matching, we see that the results are qualitatively the same as those found directly from the hydrodynamical




Figure 5. Colour–halo age relation in Illustris. We show the galaxy colour versus halo formation time in bins of stellar mass. Here, the halo formation time is defined as the accretion time of 75 per cent of the maximum mass for both centrals and satellites (see Section 2.2 for details.) Black points are each galaxy in a given stellar mass bin. The solid black line is a running median of the black points. The dashed orange line defines the monotonic relationship between colour and halo formation time, if we rank order the simulated galaxies by assigning the reddest galaxies to the oldest haloes. We use this model in Section 5 to reproduce the conformity results from the hydrodynamical simulation.




Figure 6. Galactic conformity versus semi-empirical modelling. Data points with errors are Illustris measurements from Fig. 3, while dashed curves represent the results from our semi-empirical abundance + age matching model, using two different definitions of halo age: t_{form}^{75} and t_{form}^{50} . We also show the effect of adding scatter to the colour–halo age relation for use in the age matching. More recent definitions of halo age produce larger conformity amplitudes at all separations, while adding scatter to the colour–halo age relation lowers the amplitude of the conformity signal.

simulation, with some small differences especially at the lowest mass bin. Indeed, we know from Fig. 5 that the colour–age relation in Illustris is not, in fact, strictly monotonic and exhibits a large scatter in colour at fixed halo age: therefore, our fiducial model is implementing an extreme version of a colour–halo age relation.

As we noted in Section 4, the functional shape of the colour–age relation depends on the exact choice for the definition of halo

age. We have therefore checked the output of the abovedescribed semi-empirical model for different definitions of halo formation time. Moreover, we have tested the dependence of the modelled conformity signal also on the adopted amount of scatter in the colour–age relation.

Three of these additional models for different choices of formation time and scatter in the colour–age relation are shown for

comparison and insight in Fig. 6. In all three, the subhalo abundance matching remains fixed. The ‘0.1, t75’ model adds 0.1 Gaussian scatter in $g - r$ at fixed halo age. The ‘No Scatter, t50’ model has no scatter, but ranks galaxies according to halo age as defined by the 50 per cent formation time. Finally, the ‘Variable’ model has weaker scatter at low stellar mass, and increasingly strong scatter at high stellar mass. Specifically, using the same 0.25 dex bins that we use to assign colours, we increase the scatter by 0.02 per bin, beginning with 0.02 scatter in the $\log M_{\text{stellar}} = [9.25, 9.5]$ bin. Thus, this model has 0.1 scatter (equivalent to the ‘0.1, t75’ model) in the $\log M_{\text{stellar}} = [10.25, 10.5]$ bin, and higher scatter at higher masses than that. These adopted scatter choices encompass the actual values from Fig. 5, even though with somewhat different functional dependences on halo age and mass.

The models all have quite similar consequences for the conformity signal in the two larger primary stellar mass bins. However, for the lowest stellar mass bin, the ‘No Scatter, t50’ and ‘0.1, t75’ models both underpredict the conformity signal below 3 Mpc. Meanwhile, both the ‘No Scatter, t75’ and the ‘Variable Scatter, t75’ models overpredict the conformity signal beyond 3 Mpc.

More importantly, two general trends can be taken away from such a comparison. First, more recent definitions of halo age produce larger conformity amplitudes at all separations. Compared to t_{form}^{75} , using t_{form}^{25} (not shown) lowers the amplitude of the signal significantly, while using t_{form}^{50} lowers it by an intermediate amount. No particular formation time measure is a priori better than any other, but in the Illustris simulation, t_{form}^{75} leads to the greatest spread in halo ages for central galaxies. We argue that this allows for a less noisy relationship between halo ages and galaxy colours, and therefore results in a larger inferred galactic conformity signal (at least for the specific colour-based cut adopted in this work). Secondly, as should be expected, adding scatter to the colour–halo age relation lowers the amplitude of the conformity signal. This trend continues with higher values of the scatter until the entire signal is washed away. The implication of this is substantial: galactic conformity cannot exist without a reasonably tight relationship between the galaxy property (in this case, colour) and the underlying dark matter property (in this case, halo age), which itself is clustered. The fundamental clustering is between dark matter haloes; galactic conformity comes along for the ride.

5.2 The roles of centrals and satellites

In presenting our results, we have measured the galactic conformity signal of all neighbouring galaxies above a stellar mass threshold around isolated primary galaxies. However, a subset of satellite galaxies remains in the primary sample despite the isolation criterion. Furthermore, the secondary sample includes *all* galaxies, regardless of whether they are central or satellite. In this section, we present the effect of satellites in Illustris on the galactic conformity signal, with its possible implications for observational searches for assembly bias.

In Fig. 7, we show two analogue versions of the measurements presented in Fig. 3 but for different selections of the primary/secondary samples. In the top panels, we require the primary galaxies not only to be isolated, but also to be centrals, as defined in the Illustris catalogues. That is, the galaxy must belong to the central `SUBFIND` halo in a given `FOF` halo, *and* it must still be isolated, according to our previous definition. We thus ensure that we are not introducing any added effect from nearly equal mass, merging subhaloes, and remove the effect of the 10–20 per cent of satellites that would otherwise remain in the primary sample (see Fig. 2).

Compared to Fig. 3, we see that the quenched fraction around red galaxies is reduced in the near field (out to 3 Mpc) by about a third in the lower and middle stellar mass bins, and almost eliminated in the highest mass bins. Moreover, the signal amplitude beyond 3 Mpc is also reduced, although as can be seen for the lowest mass bins, it remains statistically significant out to 10 Mpc nonetheless.

The influence of satellites becomes enlarged when the isolation (or central) criterion is extended to the secondary neighbours. If one is interested primarily in the galaxy assembly bias within central haloes, then this is the best way to control for any possible satellite effect. In Fig. 7 (bottom panel), we have thus limited our measurement to only isolated, central galaxies for the *primaries* and central galaxies for the *neighbouring secondaries*. While a conformity signal now remains out to ~ 3 Mpc – one which is nearly identical to the top panel – the far-field conformity effect between 3 and 10 Mpc has completely been suppressed. In addition, the overall quenched fraction has also been reduced, as satellites are preferentially redder than central galaxies.

In Illustris, we thus find that while satellites have a significant role in the emergence of a conformity signal, central galaxies *on their own* can produce a signal beyond the virial radius and out to roughly 3 Mpc in real-space distance. Therefore, it is possible that repeated analyses of observations using different methods for classifying isolated galaxies will obtain somewhat different quantitative results.

Interestingly, these results from Illustris imply that the terms 1-halo and 2-halo used to describe the two conformity regimes may be misleading. Specifically, in Illustris, we find a conformity signal out to 3 Mpc that, while enhanced by the inclusion of satellites, is still present when using only centrals, at least for lower primary masses. Thus, the near-field effect is in fact a 2-halo correlation. Conversely, eliminating satellites entirely from both the primary and secondary samples, as opposed to just the primary sample, eliminates the far-field signal beyond 3 Mpc. Thus, this portion of the signal in Illustris appears due to the correlation between centrals and satellites around *other* centrals, a 2-halo term, but one that will be diminished if there is no 1-halo correlation between centrals and their own satellites.

We postpone to future analysis the task to identify the physical mechanisms responsible for the emergence of the near- and far-field effects. In fact, whether or not it makes sense to distinguish between centrals and satellites in understanding assembly bias depends of the question being asked, but as noted by Hearin et al. (2015b), many satellites may have only recently been centrals themselves, and likewise, centrals may have been formerly satellites during close interactions.

5.3 Tests of mass dependence and the central-satellite split

In Fig. 8, top-left, we show that for the lowest stellar mass bin, there is a difference in the mean (thick lines), as well as the 25th and 75th percentiles (thin lines) in the masses of the neighbours of red and blue galaxies. It is thus natural to ask whether this difference – which is maximal for this stellar mass bin, and disappears at the highest stellar mass bin – is sufficient to reproduce the conformity signal. Given the role that satellite galaxies play in producing the full galactic conformity signal, however, it is also sensible to ask whether the conformity signal can be modelled knowing only that a particular halo is a central versus a satellite, in addition to knowing its mass, rather than the particular epoch of formation. In this scenario, the exact definition of halo formation is unimportant; merely we question whether knowing something of the halo’s environment is a sufficient proxy for its assembly history. We tackle these issues in this section, by further noticing that the red fraction for centrals

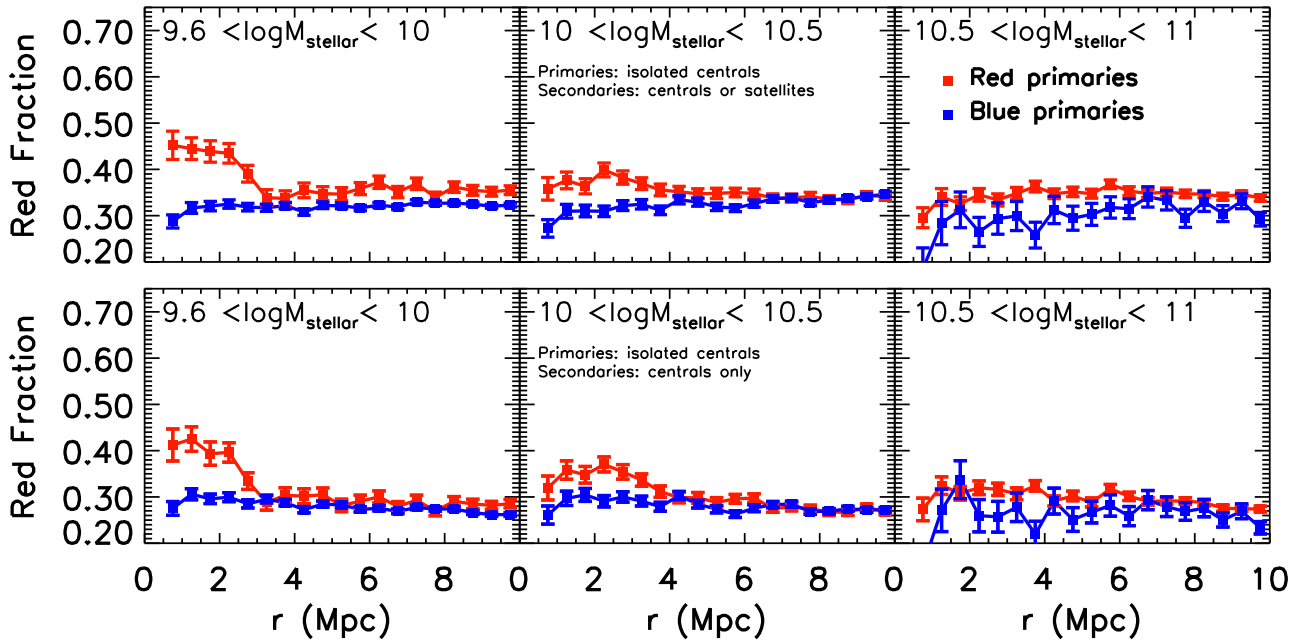


Figure 7. Galactic conformity in Illustris for different galaxy selections. Top panel: we show the conformity signal as calculated selecting only central galaxies, in addition to the isolation criterion, in stellar mass bins. A conformity signal remains at $r < 3$ Mpc, but the signal out to 10 Mpc is substantially reduced. Thus, while the isolation criterion selects mostly central galaxies for the primaries, much of the observed conformity signal at larger scales appears due to the correlation between low-mass primaries and satellites around other central galaxies. Bottom panel: we show the conformity signal as calculated using only central galaxies, for both the primary sample and the secondary neighbours. With no satellites included in either sample, a conformity signal remains at $r < 3$ Mpc, but the signal out to 10 Mpc is completely eliminated.

and satellites as a function of halo mass is substantially different, with satellites being on average much redder than centrals (Fig. 8, top-right): therefore, even if mass alone is insufficient to explain the conformity signal, the central-satellite division could produce an effect in the correct direction.

In Fig. 8, bottom-left, we confirm that the clustering dependence we measure in Illustris is not merely a product of low-mass quenched primaries being in the large-scale environment of more massive secondaries. To test this, we measure a weighted conformity signal, in which the weights are assigned solely as a function of the peak halo mass. First, we measure the mean quenched fraction as a function of halo mass in appropriately small bins (Fig. 8, solid line, top-right.) Then, by fitting a spline over the halo masses our galaxy sample, we have assigned an average ‘quenched weight’ to each halo based on its total halo mass. We then measure the equivalent mean weighted quenched fraction as a function of radius.

If the conformity signal was caused by higher red fractions in higher mass bins, combined with the fact that smaller mass quenched galaxies were more often found in the vicinity of these haloes, then we would expect to reproduce our conformity signal from Fig. 3. In this case, however, we see in Fig. 8 (bottom-left) that erasing the spatial information and measuring the conformity signal due to the average red fraction completely eliminates all trace of conformity. This is because it is not merely the halo masses making the galaxies around quenched galaxies preferentially quenched, but some other property associated with halo assembly. However, it also suggests that satellite-specific baryonic processes may be partially degenerate with long-term environmental effects for creating the colour–halo age relation. Still, we note that based on the previous subsection, satellite quenching mechanisms cannot be solely responsible, because galactic conformity remains, even when only central galaxies are considered, albeit out to smaller scales.

Now, we add the assembly information: we repeat the above procedure for assigning quenched weights independently for centrals and satellites. Each central is assigned a weight based on its red fraction as a function of central peak halo mass, and each satellite is assigned a weight based on its red fraction as a function of satellite peak halo mass. Once again, we measure the average mean weighted quenched fraction around our primary galaxies. The results (shown again for the $9.6 < \log M_{\text{stellar}} < 10$ stellar mass bin; Fig. 8, bottom-right panel) replicate the qualitative results seen in the top-left plot of Fig. 3: namely, we can model galactic conformity also by using the central versus satellite split. This is not surprising: satellites are, by construction, significantly older at fixed stellar mass than central galaxies. Thus, using the central versus satellite information rather than the formation times gets at similar information regarding the assembly history. Both formation time and whether a galaxy is a satellite are proxies for the much more complicated information encoded in the full history of a halo’s location within its environment. However, the amplitude of this signal is roughly similar to the ‘No Scatter, t50’ model in Fig. 6. Thus, like the 50 per cent formation, the central-satellite split contains less information about the colour of a galaxy than its 75 per cent formation time.

5.4 Toward comparisons with observations

In this paper, we have focused on the measurement of the conformity signal in the Illustris simulation almost exclusively from a theoretical perspective. However, in the previous sections we have touched upon typical observational choices that can affect the measured amplitude and shape of the galactic conformity effect. Here, we attempt a summary of such approaches and present future analysis directions.

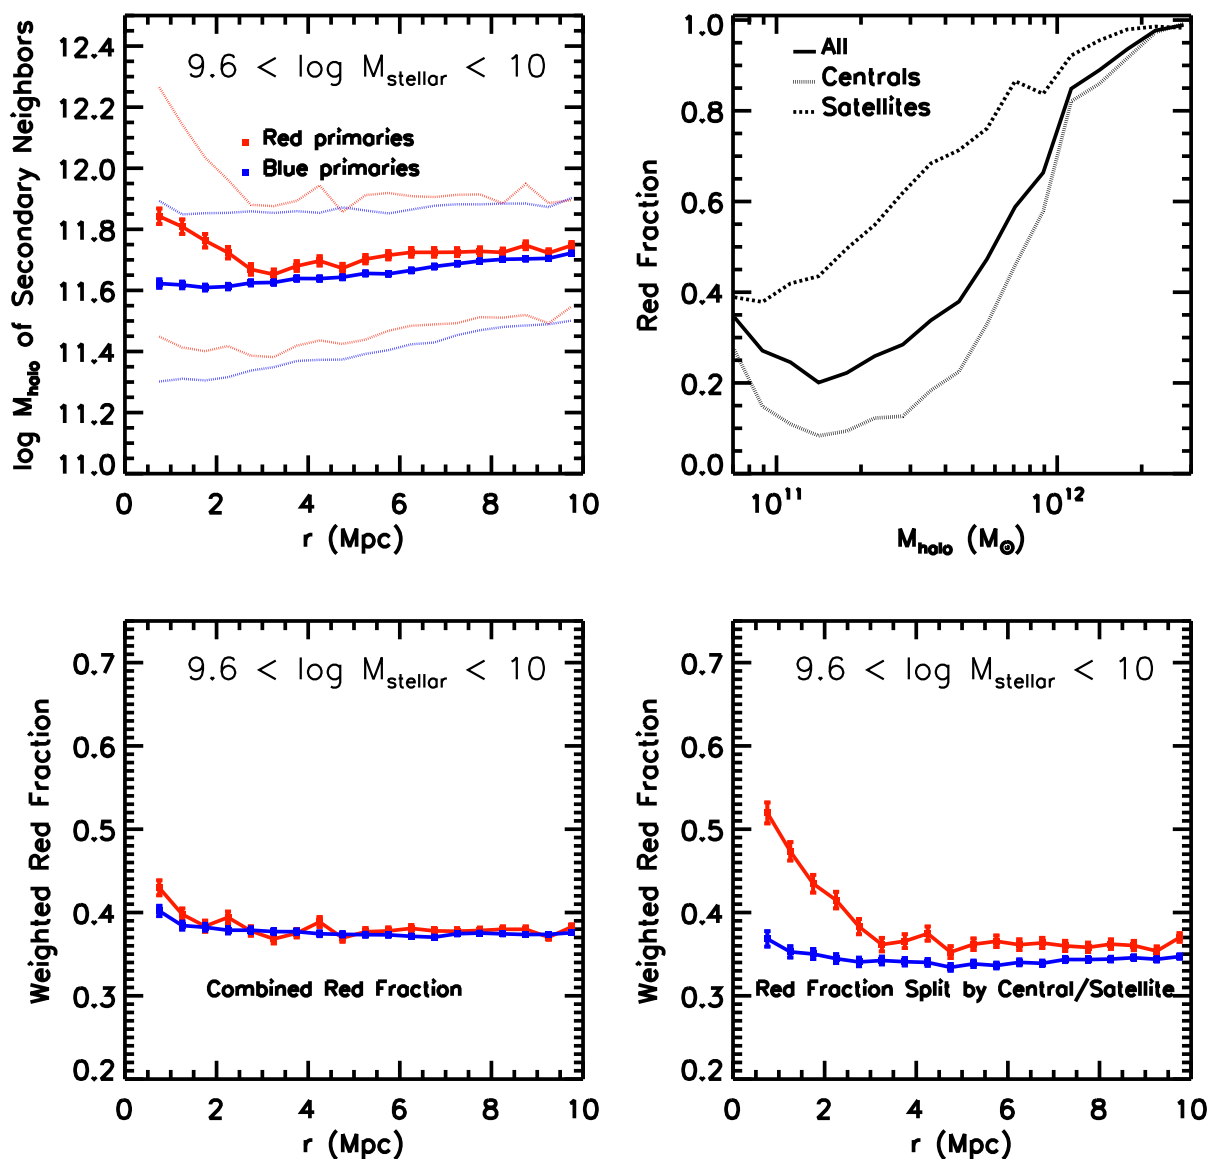


Figure 8. In the top-left panel, we show the mean halo mass, along with the 25th and 75th percentiles, of the neighbours around primary galaxies in the $9.6 < \log M_{\text{stellar}} < 10$ stellar mass bin. We choose this bin for demonstration purposes, because the difference between red and blue galaxies is maximal compared to the other stellar mass bins. (In the largest mass bin, the reverse is actually true: blue primaries are surrounded by slightly more massive galaxies than red primaries.) Because red fraction generally increases with stellar mass (see Fig. 2), we test whether this mass difference is sufficient to account for the conformity signal. In the top-right panel, we show the red fraction as a function of halo mass for both centrals and satellites, as well as both the joint population. We note that halo mass is defined as the peak mass of the accretion history. We use this relationship to test the effect the halo mass has, on its own, on the conformity signal compared to the effect of using additional properties that are sensitive to assembly history. In the bottom-left plot, we assign to each galaxy the value of the mean quenched fraction at its halo mass, as defined by a spline over narrow bins in halo mass. We plot the mean of these weights for neighbouring galaxies as a function of real-space distance in 500 kpc bins for the $9.6 < \log M_{\text{stellar}} < 10$ stellar mass bin of the primary galaxies. If the conformity signal is due entirely to quenched galaxies being surrounded by more massive galaxies than unquenched galaxies, then the signal should be of the same magnitude as in Fig. 3, whereas if the signal is not due to the masses of the secondary galaxies, but rather dependent on assembly history, then there should be no signal, because all that information has been discarded by considering only mass. As can be clearly seen, there is no remaining signal. In the bottom-right panel, we can add back in the dependence on halo assembly, but without directly using t_{form} as a proxy. Rather we weight the centrals and satellites separately, in each case by mean red fraction as a function of peak halo mass, and we qualitatively recover the conformity signal. This reinforces the idea that there is nothing special about formation time, per se. Rather, any additional measure of halo assembly, including splitting by central and satellite galaxies, introduces a bias not included by a single definition of mass.

A few choices distinguish this work from observations-based ones (e.g. Kauffmann et al. 2013) and make direct comparisons difficult to interpret. These include: (1) different isolation and selection criteria, which affect satellite galaxy inclusion; (2) 3D- versus 2D-projected separations with additional redshift-space cuts; (3) different averaging approaches for the measurement of the signal,

i.e. non-stacking versus stacking techniques of the primary galaxies; (4) different splittings of the galaxy population into quenched and star-forming samples, as fixed binary groupings (red and blue galaxies) versus a more continuous split by sSFR percentiles typically done in observations; and finally (5) different apertures for the measurement of the relevant galaxy properties themselves, either

colours or sSFR, i.e. within the whole galaxy or within a multiple of the stellar half-mass radius versus within the aperture of the SDSS ~ 1 kpc fibres (or consistently corrected measurements).

In Section 5.2, we have demonstrated that the amplitude of the galactic conformity signal is enhanced with the inclusion of satellite galaxies in the primary and secondary samples, and that the far-field term is due exclusively to satellite galaxies. We have therefore inferred that different methods for classifying isolated galaxies and separating centrals from satellites might imply different quantitative results. This is in agreement with the findings of Campbell et al. (2015), who have more systematically quantified to what extent errors in group finders can affect colour-dependent occupation statistics measured directly with galaxy group catalogues. For example, Kauffmann et al. (2013) find a conformity effect for quenched galaxies predominately out to a projected 3 Mpc: given our findings from Fig. 7, this could be associated with a central–central correlation alone, but some of the amplitude is likely due to the inclusion of isolated satellites as well.

We have also tested (but not shown) that measuring conformity in Illustris as a function of 2D projected separations instead of 3D real-space distances reduces the amplitude of the signal, and so, for example, we would expect that a projected conformity signal at any distance would require a 3D signal out to an even larger radius.

Relatively to the averaging technique, Kauffmann et al. (2013), for example, stack their galaxies in each primary bin together, thus weighting more heavily galaxies with more satellites. On the other hand, as noted above, in this work we average the red fractions for each galaxy together, thus equally weighting primary galaxies that, when binned in stellar mass, can differ in halo mass by up to a factor of 10. When we stack instead, the primary difference is that the plateau at large radii in the conformity signal truncates earlier. This is interesting, as it suggests that the farthest-field effects are being driven by primary galaxies in low-density environments (as defined by number of neighbours), whereas the nearest-field are being driven by galaxies in high-density environments. Paranjape et al. (2015) suggest that only at the large (>8 Mpc) scales is the difference between 1-halo and 2-halo conformity discernible. If the choice of whether to stack or not affects the amplitude at these scales, then great care must be taken to understand what signal is expected from assembly bias, in addition to the interpretation complications associated with scale-mixing, given that large-radius effects in 3D can affect, at reduced amplitude, 2D conformity on smaller scales.

Furthermore, in this work, we have split our sample into *red* and *blue* populations, rather than using the full sSFR distribution. As commented in Section 3, this choice can per se affect the detection of a primary-mass dependence of the conformity signal. Splitting sSFR by percentiles has the potential to reveal the overall trend in conformity in a single-mass bin – in particular, whether the trend is continuous or exhibits sharp features –, while using a global colour cut highlights any change in red fraction as a function of mass. Both choices are useful, but comparisons need to be done with care. In this work, measuring conformity based on sSFR percentile binning would accentuate possible differences in the overall sSFR distribution between the simulation and observations, potentially obscuring the nonetheless robust connection between halo age conformity and galactic conformity through the colour–halo age relation.

If the simulated and actual one-point sSFR distributions do not match each other, two-point statistics will not match each other either. In Illustris, both the colour–magnitude diagram and the sSFR distribution as a function of stellar mass are in reasonable agreement with the observations. However, it has been shown that a significant

fraction of red and quenched galaxies is in fact missing in Illustris at the intermediate- and large-mass range (Genel et al. 2014; Vogelsberger et al. 2014a), making the Illustris colour distribution less bimodal than observed. In fact, when we split the Illustris sample by sSFR percentiles, as Kauffmann et al. (2013) do with the Guo et al. (2011) semi-analytical model, we also find a weaker conformity signal than the one inferred from SDSS – effectively by construction given the narrower range of sSFR. Such discrepancy can point either to differences in the way colour or sSFR are being measured, or to choices of the subgrid and galaxy formation models that are adopted in Illustris but that do not fully capture galaxy processes. With regard to the actual measurement technique, it would be indeed ideal to be able to directly measure sSFR in the inner cores of Illustris galaxies using mock fibres, rather than to compare some choice for a simulated global value (e.g. the half-mass or twice-half-mass radius SFR values) with the corrected global SDSS SFRs (Brinchmann et al. 2004). We have tried this approach, but at $z = 0$, the mock fibres are barely larger than Illustris’ resolution of 700 pc, and so the comparison relative to point (5) above is simply not feasible at this time. As for the latter point, we are planning to explore in future analyses the emergence of galactic conformity and of the colour–age relation, and specifically to address how the actual shape of the *galaxy colour distribution* at fixed mass affects the amplitude and shape of the galactic conformity signal: this will be possible thanks to upcoming realizations of the Illustris volume with different prescriptions for the underlying galaxy formation and subgrid models.

Finally, the discussion proposed thus far serves to emphasize the degree to which methodology impacts the perceived amplitude and shape of galactic conformity. In a recent work, Kauffmann (2015) has presented an independent comparison of conformity-like signals in Illustris and SDSS, by making use of the publicly available Illustris data (Nelson et al. 2015) and by concluding that the simulated conformity effects do not show comparable strength to those seen in observations. It bears emphasizing clearly that our primary results and the conformity signal from Illustris presented by Kauffmann (2015) in Fig. 8 are not in conflict, but two different analyses of the same underlying simulation. Whereas Kauffmann (2015) presents Illustris data to directly compare to observations, our focus is on understanding the conformity signal in the Illustris simulation. In fact, the apparent qualitative discrepancies are due to observationally motivated choices as the one discussed above. The difference in sSFR distributions, together with the employment of projected separations and stacking techniques, is bound to suppress the measured conformity signal, therefore explaining why the Illustris-based results in Kauffmann (2015) appear weaker than the one presented here. In turn, the fact that the conformity signal in fig. 8 of Kauffmann (2015) is weaker than the one inferred from SDSS data might indeed imply that in our simulation galactic conformity is underestimated: it is reasonable to speculate that a stronger bimodality in the simulated Illustris galaxy population or simply a different galaxy colour distribution that better matches observations could enhance the theoretically measured conformity signal herein presented.

6 SUMMARY AND CONCLUSIONS

In this paper, we have presented a series of phenomenological measurements of the galactic and halo clustering signals and of the relation between galaxy colours and dark-matter halo ages from the redshift $z = 0$ snapshot of the Illustris simulation. This is a full-volume cosmological hydrodynamical simulation from which we have selected a sample of about 18 000 galaxies with stellar

mass $>2 \times 10^9 M_{\odot}$ in a $75 h^{-1}$ Mpc box. Rather than compare directly with observations, we have chosen to analyse the conformity signal found in Illustris and sought to explain it in the context of dark-matter halo clustering. Further, we have attempted to understand what confounding effects may affect the inferred amplitude and shape of the galactic conformity signal obtained from actual observations.

Our primary results are as follows.

(i) We find that the colours of galaxies at fixed stellar or halo mass are spatially correlated for pair separations that go well beyond the virial radius. Namely, we find in Illustris that red isolated galaxies are surrounded by a higher fraction of red neighbours than their bluer counterparts, in qualitative agreement with recent observational claims.

(ii) This galactic conformity signal extends out to ~ 10 Mpc for low-mass primary galaxies (far-field term), has a larger amplitude at separations smaller than about 3 Mpc (near-field term), and decreases in overall amplitude towards higher primary mass (Fig. 3).

(iii) In concordance with previous gravity-only numerical studies, we find a qualitatively similar dark-matter halo conformity signal, in which older haloes are preferentially neighbored by older haloes, at fixed mass (Fig. 4).

(iv) We show that a galaxy colour–halo age relation naturally arises in the Illustris simulation, with the reddest galaxies being preferentially found in the oldest dark-matter haloes (Fig. 5).

(v) We demonstrate that a plausible interpretation of the galactic conformity effect can be given as the combination of the halo conformity signal with the relation between galaxy colours and halo ages.

(vi) In practice, we show that by using abundance and age matching, we can reproduce the signal measured directly in the Illustris simulation, demonstrating that once a reasonably tight relationship exists between galaxy property (in this case, colour) and the underlying halo property itself clustered (in this case, halo age), then galactic conformity is, to first approximation, a consequence of halo conformity. Conversely, increasing the scatter in the colour–halo age relationship lowers the amplitude of the so-modelled conformity, suggesting that strong feedback mechanisms that remove this galaxy–halo correlation may eliminate entirely any galactic conformity signal. In this picture, the underlying dark matter clustering is the primary contributing factor to the galactic conformity signal, rather than a spatial correlation induced by baryonic processes. A colour–halo age relation must necessarily be present for a galactic conformity signal to be visible.

(vii) In fact, different proxies for the assembly history of haloes (different halo formation time choices, as well as knowledge of whether a galaxy is a central or satellite) can give substantially different quantitative results, even as the qualitative picture remains the same. Thus, understanding the primary way in which assembly history affects galaxy properties remains an outstanding issue.

(viii) Finally, we show that the conformity signal out to ~ 3 Mpc arises from both correlations between central galaxies and from central-satellite correlations. Therefore, there is a 2-halo effect in Illustris accounting for the near-field conformity signal, as well as a 1-halo term. At larger radii, the signal is driven by the correlation between central galaxies and the satellites of other centrals.

We note that the actual quantitative amplitude of the theoretically predicted conformity signal ultimately depends on the simulated colour–magnitude diagram (or equivalently, the sSFR versus stellar mass distribution). Given the lack of a strong bimodality of Illustris

galaxies colours at intermediate masses and given the lack of a significant fraction of quenched galaxies towards the high-mass end, it is likely that the quantitative galactic conformity results presented in Fig. 3 underpredict the observational conformity. None the less, this paper shows that *even with* a significantly smaller bimodality, halo assembly bias is an effective mechanism for producing a galactic conformity signal in Illustris. Spatial correlations or interactions between baryons in different haloes are not needed; all that is required is a strong enough correlation between halo age and galaxy colour. However, we do not rule out the possibility that the colour–halo age relation could be induced by baryonic physics that affect volumes beyond the virial radius. As new measurements are made in different spatial fields and at higher redshifts, the selection criteria, and the choice of stacking techniques may have a dramatic effect on the amplitude and shape of the conformity signal inferred from observations, and therefore, on its interpretation.

ACKNOWLEDGEMENTS

ADB, AP, and SG thank Andrew Hearin for helpful discussions and comments on the manuscript and thank Frank van den Bosch, Jeremy Tinker, Guinevere Kauffmann, Alexie Leauthaud, and Cameron McBride for useful discussions. ADB thanks the OSU CCAPP for hosting its 2014 HOD workshop. AP acknowledges support from the *HST* grant HST-AR-13897. SG acknowledges support provided by NASA through Hubble Fellowship grant HST-HF2-51341.001-A awarded by the STScI, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. PT acknowledges support from NASA ATP Grant NNX14AH35G. VS acknowledges support through the European Research Council under ERC-StG grant EXAGAL-308037. LH acknowledges support from NASA grant NNX12AC67G and NSF grant AST-1312095.

The Illustris-1 simulation was run on the CURIE supercomputer at CEA/France as part of PRACE project RA0844, and the SuperMUC computer at the Leibniz Computing Centre, Germany, as part of GCS-project pr85je. The further simulations were run on the Harvard Odyssey and CfA/ITC clusters, the Ranger and Stampede supercomputers at the Texas Advanced Computing Center through XSEDE, and the Kraken supercomputer at Oak Ridge National Laboratory through XSEDE. The analysis presented in this paper was conducted on the Harvard Odyssey and CfA/ITC clusters.

REFERENCES

- Berlind A. A., Weinberg D. H., 2002, *ApJ*, 575, 587
 Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, 351, 1151
 Campbell D., van den Bosch F. C., Hearin A., Padmanabhan N., Berlind A., Mo H. J., Tinker J., Yang X., 2015, *MNRAS*, 452, 444
 Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201
 Croton D. J., Gao L., White S. D. M., 2007, *MNRAS*, 374, 1303
 Dalal N., White M., Bond J. R., Shirokov A., 2008, *ApJ*, 687, 12
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 Di Matteo T., Spingel V., Hernquist L., 2005, *Nature*, 433, 604
 Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
 Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, 703, 1416
 Gao L., Springel V., White S. D. M., 2005, *MNRAS*, 363, L66
 Genel S. et al., 2014, *MNRAS*, 445, 175
 Guo Q., Cole S., Eke V., Frenk C., 2011, *MNRAS*, 417, 370
 Han J. et al., 2015, *MNRAS*, 446, 1356

- Hartley W. G., Conselice C. J., Mortlock A., Foucaud S., Simpson C., 2015, *MNRAS*, 451, 1613
- Hearin A. P., Watson D. F., 2013, *MNRAS*, 435, 1313
- Hearin A. P., Watson D. F., van den Bosch F. C., 2015a, *MNRAS*, 452, 1958
- Hearin A. P., Behroozi P. S., van den Bosch F. C., 2015b, preprint ([arXiv:1504.05578](https://arxiv.org/abs/1504.05578))
- Heavens A., Panter B., Jimenez R., Dunlop J., 2004, *Nature*, 428, 625
- Hinshaw G. et al., 2013, *ApJS*, 208, 19
- Jimenez R., Panter B., Heavens A. F., Verde L., 2005, *MNRAS*, 356, 495
- Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
- Kauffmann G., 2015, *MNRAS*, 454, 1840
- Kauffmann G., Li C., Zhang W., Weinmann S., 2013, *MNRAS*, 430, 1447
- Keselman J. A., Nusser A., 2007, *MNRAS*, 382, 1853
- Knobel C., Lilly S. J., Woo J., Kovač K., 2015, *ApJ*, 800, 24
- Lacerna I., Padilla N., 2011, *MNRAS*, 412, 1283
- Lacerna I., Padilla N., Stasyszyn F., 2014, *MNRAS*, 443, 3107
- Li Y., Mo H. J., Gao L., 2008, *MNRAS*, 389, 1419
- Lin Y.-T., Mandelbaum R., Huang Y.-H., Huang H.-J., Dalal N., Diemer B., Jian H.-Y., Kravtsov A., 2015, preprint ([arXiv:1504.07632](https://arxiv.org/abs/1504.07632))
- Nelan J. E., Smith R. J., Hudson M. J., Wegner G. A., Lucey J. R., Moore S. A. W., Quinney S. J., Suntzef N. B., 2005, *ApJ*, 632, 137
- Nelson D., Genel S., Vogelsberger M., Springel V., Sijacki D., Torrey P., Hernquist L., 2015, *MNRAS*, 448, 59
- Paranjape A., Kovac K., Hartley W. G., Pahwa I., 2015, preprint ([arXiv:1503.08212](https://arxiv.org/abs/1503.08212))
- Rodriguez-Gomez V. et al., 2015, *MNRAS*, 449, 49
- Sales L. et al., 2015, *MNRAS*, 447, L6
- Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
- Sparre M. et al., 2015, *MNRAS*, 447, 3548
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Springel V., Di Matteo T., Hernquist L., 2005, *MNRAS*, 361, 776
- Thomas D., Maraston C., Bender R., Mendes de Oliveira C., 2005, *ApJ*, 621, 673
- Tinker J. L., Wetzel A. R., Conroy C., 2011, preprint ([arXiv:1107.5046](https://arxiv.org/abs/1107.5046))
- Torrey P., Vogelsberger M., Genel S., Sijacki D., Springel V., Hernquist L., 2014, *MNRAS*, 438, 1985
- Vale A., Ostriker J. P., 2006, *MNRAS*, 371, 1173
- 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
- Wang L., Weinmann S. M., De Lucia G., Yang X., 2013, *MNRAS*, 433, 515
- Watson D. F. et al., 2015, *MNRAS*, 446, 651
- Wechsler R. H., Bullock J. S., Primack J. R., Kravtsov A. V., Dekel A., 2002, *ApJ*, 568, 52
- Wechsler R. H., Zentner A. R., Bullock J. S., Kravtsov A. V., Allgood B., 2006, *ApJ*, 652, 71
- Weinmann S. M., van den Bosch F. C., Yang X., Mo H. J., 2006, *MNRAS*, 366, 2
- Tornatore, LWiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., 2009, *MNRAS*, 399, 574
- Yang X., Mo H. J., van den Bosch F. C., 2003, *MNRAS*, 339, 1057
- Zehavi I. et al., 2011, *ApJ*, 736, 59
- Zentner A. R., Hearin A. P., van den Bosch F. C., 2014, *MNRAS*, 443, 3044
- Zu Y., Mandelbaum R., 2015, *MNRAS*, 454, 1161

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