

Reversing arrested development: a new method to address halo assembly bias

William J. Smith                                         <img alt

2012; Wetzel et al. 2014). Gravitational heating occurs when a low-mass halo cannot accrete particles in its vicinity because they are too energetic. Those particles are more energetic because a more massive neighbouring halo accelerates them. Thus, a method that addresses arrested development as a truncation of mass accretion due to a nearby more massive neighbour or neighbours addresses all three of these proposed causes of low-mass assembly bias simultaneously.

This work investigates arrested development as the primary cause of halo assembly bias. Hahn et al. (2009) first tested the hypothesis of arrested development as a cause of assembly bias using an N -body simulation and analysing the correlation between halo assembly history and environment, in this case defined not by a density field in the form of a two-point correlation function, but instead by a tidal/shear field. They argued that these tidal effects do indeed suppress halo growth. Though they did not go so far as to conclude that arrested development is the cause of assembly bias, they did conclude that it is an important factor contributing to its emergence in simulation data.

Another work lending evidence to arrested development as a source of halo assembly bias is Hearin, Behroozi & Van Den Bosch (2016). They analysed halo mass accretion rates of pairs of haloes as functions of the individual pair masses, pair separation, and pair environment. They found mass accretion conformity out to many times the haloes' virial radii. Halo pairs in overall denser regions have less mass accretion, and the lower the haloes' masses, the more pronounced this conformity. They state that this conformity finding caused by the halo tidal environment is an alternative quantification of other assembly bias findings.

A more recent and more direct investigation of sources of assembly bias is presented by Salcedo et al. (2018), who analysed a multitude of halo properties to disentangle which properties have the strongest effects on assembly bias. One of their main analyses studied halo properties as a function of distance to a halo's nearest more massive neighbour. Importantly, they found that, if one controls for 'neighbour bias', then bins haloes by mass and splits by age or concentration, the assembly bias signal is mostly removed. However, they used several different halo properties as both primary and secondary properties to bin and split haloes and found that controlling for neighbour bias does not successfully remove assembly bias in all cases, complicating their ability to discern a simple cause.

Though multiple explanations for assembly bias have been proposed, none alone has satisfactorily explained the phenomenon. Mansfield & Kravtsov (2020) attempted to assign a relative importance and a cohesive model for all three of the low-mass causes of assembly bias. They used a single set of simulations with multiple prescriptions and proxies for calculating splashback radii. They measured tidal strength from both the largest nearby halo and the general strength of the aggregate tidal field and measured tidal heating. From this, they determined the relative contributions of each factor by finding the percentage of haloes that need to be removed for a given proxy to eliminate the assembly bias signal. The results of Mansfield & Kravtsov (2020) support the connection between halo tidal forces halo assembly bias in agreement with the conclusions of Hahn et al. (2009), Hearin et al. (2016), and Salcedo et al. (2018), but they express a caveat that assembly bias is likely not an effect that can be approximated from the tidal force of a halo's single most massive neighbour alone.

We propose a new approach to investigate arrested development as a cause of assembly bias. The idea is to identify arrested development in haloes and predict the growth they would have had if they had not undergone a slow-down or truncation of mass accretion. Since arrested development mostly occurs at late times, when massive

neighbours have collapsed and produced strong tidal fields, we can attempt to extrapolate from the early part of a halo's history to predict its late evolution. To perform this extrapolation, we adopt prior work in using a functional form to fit mass accretion histories. Having removed the effects of arrested development from halo mass accretion histories, we can investigate to what extent the assembly bias signal is reduced.

This new approach offers more than an investigation of the cause of halo assembly bias; it also may offer a straightforward way of removing the halo assembly bias signal in simulations consistent with its underlying physical cause. Though prior work has been well motivated in addressing physical causes of assembly bias, the low-mass causes of assembly bias propose challenges for strategies to mitigate this bias. Because our proposed method attempts to correct the underlying cause of assembly bias – arrested development – it could be a useful tool that can be applied to any simulation, with practical applications, such as producing assembly bias-free mock catalogs.

Section 2 provides an overview of the simulation and how we extract mass accretion histories. Section 3 describes our procedure for fitting mass accretion histories and correcting for arrested development. Section 4 shows how we calculate the halo assembly bias signal via their two-point correlation function. Section 5 shows our main results from attempting to mitigate arrested development and investigating how this impacts the halo assembly bias signal. Section 6 discusses limitations and implications for future assembly bias investigations.

2 DATA

2.1 The Vishnu simulation

Vishnu is a large-scale cosmological high-resolution N -body simulation. The simulation contains 1000 snapshots of particles from $z = 99$ to 0 and was evolved using the GADGET-2 N -body TreeSPH algorithm (Springel 2005) in a Λ CDM cosmology with $h = 0.70$, $\Omega_m = 0.25$, $\Omega_\Lambda = 0.75$, and $\sigma_8 = 0.8$. The initial power spectrum of Vishnu was calculated using CAMB, a PYTHON code for cosmological calculations (Lewis, Challinor & Lasenby 2000). Initial positions and velocities were determined using the 2LPT code (Scoccimarro 1998) at $z = 99$. Vishnu has a boxsize of $130 h^{-1}$ Mpc and a particle mass of $3.2151 \times 10^7 h^{-1} M_\odot$ (Johnson et al. 2019). Haloes were identified using the ROCKSTAR algorithm (Behroozi, Wechsler & Wu 2012a) with the virial overdensity set to Δ_{vir} (Bryan & Norman 1998). Halo merger trees were constructed with the Consistent-Trees algorithm of Behroozi et al. (2012b).

The Vishnu simulation offers a high mass and temporal resolution compared to most other cosmological N -body simulations while still maintaining a large enough volume to measure the large-scale clustering of haloes. This makes it an ideal simulation for studying the connection between halo mass accretion histories and assembly bias. Since this analysis only uses the total mass of haloes and their clustering, and is not concerned with internal properties like density profiles, we are confident in using haloes with a fairly small number of particles. However, since the focus of this work is to analyse mass accretion histories, not just haloes in the present, we require haloes to be resolved at very early times when they were much less massive. This means that our final $z = 0$ mass threshold should be much higher than the minimum mass we can resolve. We adopt a final halo mass threshold for our analysis of $10^{10} h^{-1} M_\odot$, which corresponds to about 300 particles. This allows us to resolve halo histories over the period where they have grown by at least an order of magnitude. If, for example, a present-day halo of mass $10^{10} h^{-1} M_\odot$ has a half-

mass age at scale factor $a = 0.5$, then the halo at this earlier point would have roughly 150 particles, still enough for the analysis in this work and enough that we can probe its accretion history back to even earlier times.

The main idea in this paper is to fit the early portion of the mass accretion history (MAH) of haloes before arrested development has taken effect, and extrapolate the fit to predict how the haloes would have grown in the absence of arrested development. We thus need to adopt a truncation point in time from which to extrapolate. This point should be at a time that is early enough that arrested development has not yet taken place, but late enough that the halo is well resolved in a sufficient number of previous time-steps to give us confidence in the early accretion history fit. The earliest truncation point that we adopt in this analysis is $a = 0.5$, which satisfies these requirements. We expect that haloes that experience arrested development will have extrapolated final masses that are higher than their actual final masses in the simulation. Therefore, in order to have a complete sample of haloes with extrapolated final masses greater than $10^{10} h^{-1} M_\odot$, we must perform fits on haloes with actual final masses that are lower than that so that arrested development haloes are included in the sample. To satisfy this requirement, we create a sample of all haloes in the simulation with final $z = 0$ masses greater than $10^{9.5} h^{-1} M_\odot$, which provides enough of a buffer to ensure a complete sample above $10^{10} h^{-1} M_\odot$. Additionally, we adopt a high mass cut-off of $10^{11.6} h^{-1} M_\odot$ because, above this threshold, there are too few haloes per mass bin to calculate accurate clustering statistics. Finally, only haloes classified as central haloes at $z = 0$ are used in this analysis (i.e. not subhaloes). These conditions result in a sample of roughly 1.2×10^6 haloes from the Vishnu simulation that we fit for the analysis.

2.2 Extracting the mass accretion histories

The Consistent-Trees algorithm of Behroozi et al. (2012b) creates halo merger trees in Vishnu, which include every branch of every progenitor halo that results in a current halo. Since every halo at every time-step in the simulation is given a unique ID number, the Consistent-Trees algorithm records the descendant halo IDs for every halo. From this, the mass accretion histories of every halo at $a = 1$ are extracted by moving backwards in time and finding progenitor halo IDs and selecting the progenitor with the most mass. One branch is randomly chosen in the rare event that multiple progenitor haloes have the same mass. This process is repeated until a halo with no progenitor halo is reached. We refer to the scale factor at which the halo with no progenitor is reached as the ‘epoch of appearance’ of that halo. We note that this does not mean that the halo did not exist at earlier times, but that it had too few particles to be registered by the ROCKSTAR halo finder algorithm. Fig. 1 shows six randomly selected mass accretion histories from Vishnu. The haloes exhibit a large variety in their detailed histories. For example, the halo represented by the purple line exhibits sudden mass loss at around $a = 0.75$. However, all of these example haloes share roughly the same overall form. In their early histories, the haloes accrete at a rapid fractional rate, while later in their histories, the mass accretion histories level out.

3 FITTING THE MASS ACCRETION HISTORIES

A search for a generalized mass accretion history function that can be derived from an analytic foundation and verified by simulations is ongoing. It has been recognized that, in general, most haloes

fractionally accrete much of their mass early, then experience a period of gradually slower fractional accretion (Ludlow et al. 2013) (see Fig. 1), with this effect being more or less pronounced depending on other halo properties (notably overall halo mass) (Wechsler et al. 2002; Tasitsiomi et al. 2004; McBride, Fakhouri & Ma 2009). Early work attempted to derive a generalized halo mass accretion history profile from mass assembly histories (e.g. Ryden & Gunn 1987; Avila-Reese, Firmani & Hernandez 1998). In a seminal paper, Wechsler et al. (2002) extracted and analysed structural merger trees of haloes from simulation data. Prior work by Bullock et al. (2001) had used a simulation to study the density profiles of dark matter haloes, and Wechsler et al. (2002) attempted to find a correlation between these density profiles and the assembly history of the haloes. In doing so, they proposed an exponential fitting function for mass accretion histories of the form:

$$M(a) = M_0 e^{-\alpha z}, \quad (1)$$

where $M(a)$ is the mass of the halo at some scale factor a , M_0 is the final mass of the halo, z is the redshift, which is related to the scale factor as $a = (1 + z)^{-1}$, and α is a free parameter. Fig. 2 shows this functional form, fit to one of our Vishnu haloes. Wechsler et al. (2002) notably restricted fitting this function to haloes above $10^{12} h^{-1} M_\odot$ at $z = 0$ to have reliable fits for most haloes. This choice created a better set of halo fits for their analysis, but it raises doubt on this function’s appropriateness for fitting haloes of lower masses.

Though this functional fit for mass accretion histories is simple and powerful, subsequent researchers recognized its limitations and suggested improved functional forms. For example, van den Bosch (2002) suggested a two-parameter function and Tasitsiomi et al. (2004) proposed a two-parameter, more general function of:

$$\tilde{M}(\tilde{a}) = \tilde{a}^p \exp \left(\alpha \left(1 - \frac{1}{\tilde{a}} \right) \right), \quad (2)$$

in which $\tilde{M}(\tilde{a}) = \frac{M}{M_0}$, $\tilde{a} = \frac{a}{a_0}$, and M_0 and a_0 are the virial mass and the scale factor of the halo at the time (or epoch) of observation, and p and α are free parameters. This parametrization allows haloes to be fit with a functional form that is a combination of an exponential of $\frac{1}{\tilde{a}}$ and a power law of base \tilde{a} . For cases in which $p = 0$, this function for fitting mass accretion histories simplifies to the one-parameter exponential function of Wechsler et al. (2002). Tasitsiomi et al. (2004) proposed and applied this form to 14 haloes for analysis. McBride et al. (2009) then tested this form over a much larger sample of approximately 500 000 haloes from the Millennium Simulation and found it to be ‘a reasonable fit’. They rewrote the function identically as:

$$M(z) = M_0 (1 + z)^\alpha e^{\beta z} \quad (3)$$

with α and β as free parameters, which is the form used for this analysis, though we convert redshift z to scale factor a for this work. Some sources using this two-parameter functional form have a negative exponential term, though this is just a sign convention choice. Furthermore, McBride et al. (2009) suggest a four-tiered classification system based on the fitted values of the two free parameters to classify halo types. This classification scheme refers to haloes with $\alpha < 0.35$ as ‘good exponential’ haloes, haloes with $\alpha - \beta < -0.45$ as ‘steep late growth’ haloes, haloes with $-0.45 < \alpha - \beta < 0$ as ‘shallow late growth’ haloes, and haloes with $\alpha - \beta > 0$ as ‘late plateau/decline’ haloes. Though this categorization is not used explicitly in this analysis, it is a consideration for future work.

This specific two-parameter functional fit was further investigated by Wong & Taylor (2012), who noted a concern that for most fits α and β are correlated and exhibit some degeneracy (meaning multiple

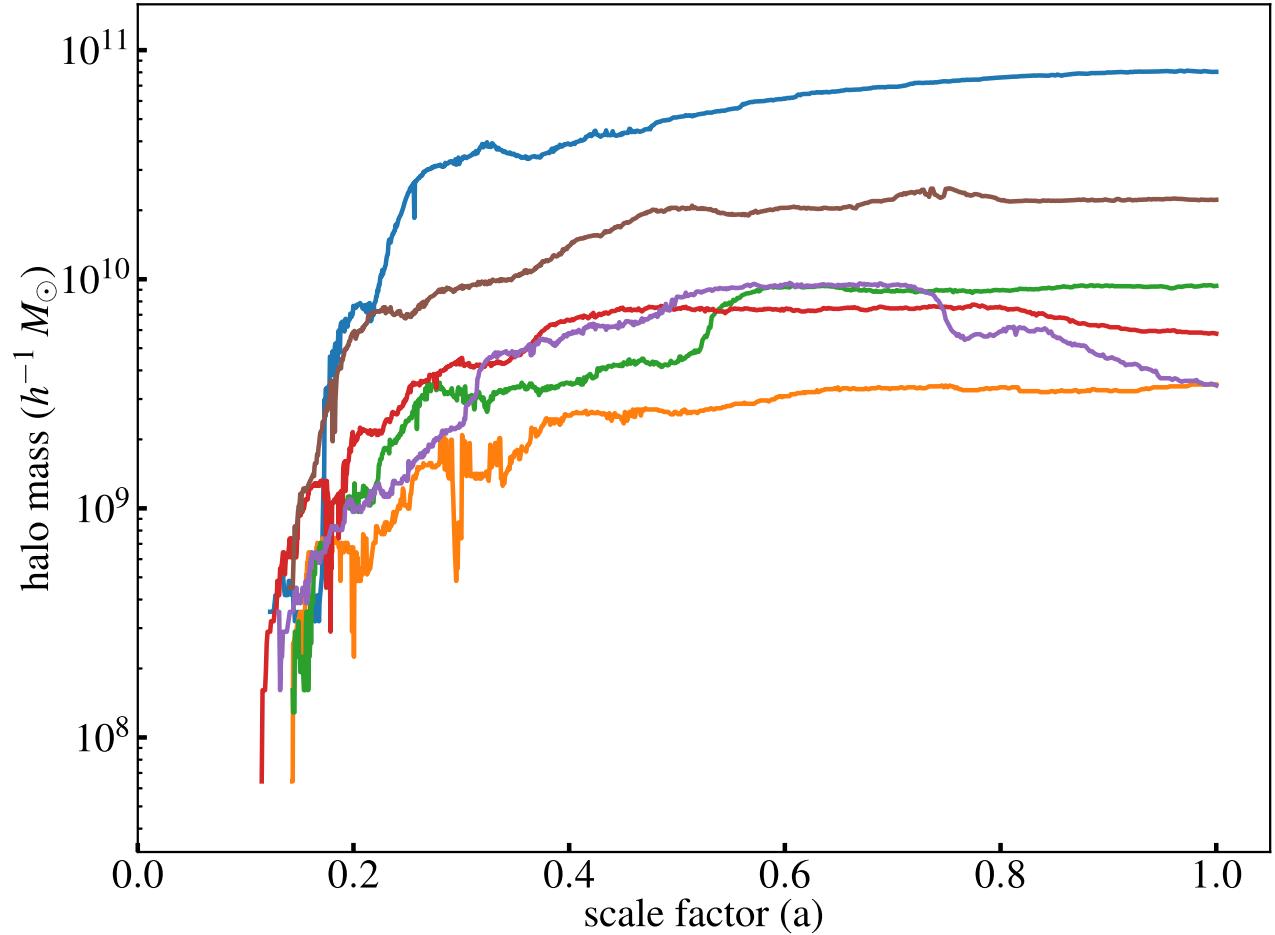


Figure 1. Six mass accretion histories from the Vishnu simulation. Each line shows mass as a function of scale factor, a , for a particular randomly chosen halo. The haloes exhibit a large variety in their histories. For example, the halo represented in purple shows a mass loss in its later history ($a \sim 0.75$). Most show a levelling off of mass accretion in their late histories.

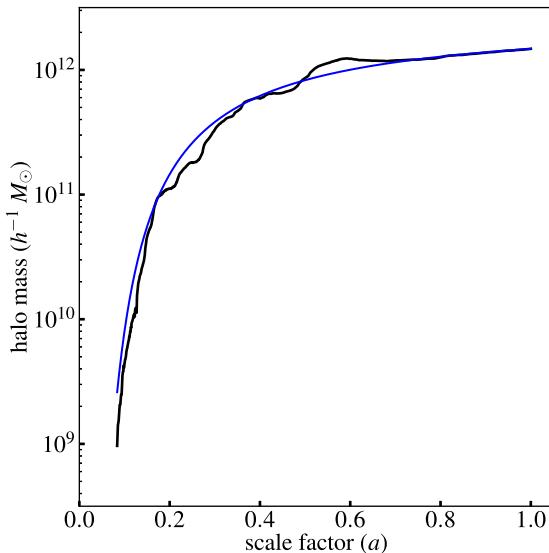


Figure 2. Example of a mass accretion history, always following the most massive progenitor back in time, for a halo from the Vishnu simulation (black), fit with an exponential function of the form $M(a) = M_0 e^{(-\alpha a)}$ (solid blue). This exponential fit is a simple parametrization that fits many haloes well.

value combinations of α and β can produce the same functional fit). They further proposed a model where the primary free parameter is $(\beta - \alpha)$. Correa et al. (2015) explained this correlation and provided a more comprehensive justification for this functional form by fitting to an analytic mass history model based on the EPS formalism and connecting α and β to the initial power spectrum. Though continued work investigates the best functional fit for mass accretion histories, this is currently the most well-researched and tested function for this purpose.

We use the parametrization of McBride et al. (2009) (equation (3)) to functionally model all the mass accretion histories in our simulation: once over the entire history and again over a portion of the early history out to specified truncation points. The fit over the entire history checks how well the function models the mass accretion over the entire history of the haloes and serves as a comparison to the functional fit over the early histories of the haloes. The equation being fit is equation (3), where M_0 , α , and β are free parameters. Since the independent time variable is in units of scale factor instead of redshift, the function is converted to be a function of scale factor a instead of redshift z . Though this function is identical to the functional form of mass accretion histories used by Correa et al. (2015) and prior works, one difference is that we add a third free parameter, the final halo mass. In past works, the final mass M_0 was assumed to be known and fixed to the final mass of the halo in the simulation data. However,

in this analysis, we assume that the final halo mass, had the halo not undergone arrested development, is not known from the simulation itself but can be inferred from extrapolating the early halo accretion data. Thus, this is the first investigation into halo accretion that allows M_0 to be a free parameter when fitting haloes.

The fit is performed using the *optimize* subpackage of the SCIPY package. For each halo fit, boundary conditions are imposed on the free parameters. Specifically, the parameter M_0 , being a mass, is constrained to be positive, and both α and β are constrained to only be negative. In the case of β , this constraint arises from the functional form. A positive β would cause the function to have an asymptote at positive infinity as it approaches scale factor zero regardless of the values of the other two parameters, representing an unphysical halo with infinite mass at the beginning of the Universe. The α parameter has no physical motivation for imposing a boundary condition; however, a negative value of α imposes a restriction against mass loss for the halo. Importantly, for some halo history forms, a range of positive values for α can also result in mass accretion histories that do not exhibit mass loss, but the specific value of α for which the halo mass function has a ‘turnover point’ and exhibits mass loss is itself dependent on the value of β . For this work, a boundary condition is imposed that α be negative, preventing all halo fits from exhibiting mass loss. We restrict the possibility of mass loss when fitting haloes because it more accurately accounts for the arrested development hypothesis, which states that assembly bias is caused by haloes not accreting mass as they would without tidal forces. Whether a halo undergoing arrested development gains mass more slowly than it otherwise would or whether it actually loses mass, in the absence of the tidal forces that result in arrested development, haloes can only gain mass. Thus, we expect that, whether the halo itself loses mass or not, the ‘corrected’ mass history from our functional fit should only increase in mass.

Once the optimal values of the three parameters of M_0 , α , and β are calculated using Scipy’s curve fit module, a model halo’s fitted mass history can be calculated from the function and parameters. From the model mass accretion history, the model half mass age is calculated as the scale factor of the model halo when it has acquired half its final mass. Fig. 3 shows the results of this fit for two example haloes from the Vishnu simulation, one ‘normal’ halo (top panel) and one arrested development halo (bottom panel). The solid black lines show the actual mass accretion histories of the haloes as extracted from the simulation. The dot-dashed orange curves show the best-fit model accretion histories, using the fitting function in equation (3), when the fit is performed over the entire mass accretion histories. The half-mass ages, estimated from the actual and fitted histories, are denoted by the red circle and orange X symbols, respectively. In both cases, the fitted history tracks the overall shape of the actual history fairly well and has a similar final mass and half-mass age.

Next, mass accretion histories are fit using only the early portion of their histories. This results in a best-fit analytic history with new best-fit values of M_0 , α , and β . The later part of the halo’s history can be extrapolated from the fit to the earlier portion of the history. In this work, we refer to the latest epoch that is used in the fit as the ‘projected scale factor’. We begin with a projected scale factor of $a = 0.5$. This value was chosen because it is early enough in most halo histories to allow the extrapolation to correct for arrested development but late enough that haloes are resolved in a sufficient number of earlier epochs of the simulation in order for a trustworthy fit to be made. The blue curves in Fig. 3 show halo fits using a projected scale factor of 0.5 for the examples of ‘normal’ and arrested development haloes. The corresponding half-mass ages from these fits are denoted by the blue diamond symbols. The figure shows that in the case of the

‘normal’ halo, this new fit matches the previous fit well, meaning that the late-time extrapolation of the mass accretion history represents the true history well even though it is only based on the early history of the halo. In the case of the arrested development halo, however, the extrapolated accretion history overpredicts the late-term mass of the halo by a large factor, with a $z = 0$ mass that is more than two times higher than the actual mass of the halo. Likewise, the half-mass scale factor of the ‘corrected’ history is ~ 50 per cent larger (i.e. younger) than the original age. Our interpretation of these differences is that the early history of the halo is not affected by arrested development, and so the extrapolation of the fit to this early history is predicting how the halo would have grown in the absence of arrested development.

We can use Fig. 3 to demonstrate the importance of the chosen projected scale factor. As the projected scale factor (represented in Fig. 3 by the black arrows) is set to earlier times, the number of simulation time-steps over which to fit the halo between its appearance in the simulation and its truncation point decreases. Conversely, as the projected scale factor moves to later time-steps, the truncated halo fit (the blue curve) naturally moves closer to the model halo fit (the orange curve). In this case, the projected halo converges to the model halo fit, undermining the goal of accounting for arrested development in the halo. In addition to a projected scale factor of $a = 0.5$, we also analyse projected scale factors of $a = 0.6, 0.7, 0.8$, and 0.9 . We fit each halo from its scale factor of first appearance in the simulation to the projected scale factor. Haloes that formed after the projected scale factor are excluded from the analysis.

Fig. 4 demonstrates the appropriateness of our chosen projected scale factor of $a = 0.5$. In choosing a projected scale factor, we strive to ensure two criteria. First, most haloes in the simulation should have an epoch of appearance before the projected scale factor. Clearly, haloes whose oldest resolved progenitor haloes appear after our chosen projected scale factor cannot be fit. Such haloes would be excluded by our analysis, potentially biasing the result. The left-hand panel of Fig. 4 shows the distribution of epoch of appearance for all haloes with final masses above $10^{9.5} h^{-1} M_\odot$ in the Vishnu simulation. The overwhelming majority of haloes appeared before $a = 0.5$, with most appearing just before $a = 0.2$. Thus, we are confident that our choice in the projected scale factor of $a = 0.5$ does not entail a significant source of bias due to excluded haloes. Second, we look at the data in each mass accretion history, focusing on how many of the simulation time-steps fall before the projected scale factor. If there are too few time-steps between the halo’s epoch of appearance and the projected scale factor, we risk extrapolating the halo’s late history from too brief a period in the halo’s history, yielding unreliable results. Conversely, if most of the time-steps are between the halo’s epoch of appearance and the projected scale factor, then there is little to extrapolate. In this case, the truncated halo will yield the same result as the model fit. The right-hand panel of Fig. 4 shows the distribution of the fraction of the time-steps between the epoch of appearance of the halo and the projected scale factor of $a = 0.5$. To help understand this, consider the sharp peak at roughly 0.6. This signifies that for a large number of haloes, about 60 per cent of the time-steps in the simulation fall between the halo’s epoch of appearance and the projected scale factor of $a = 0.5$, while 40 per cent of the time-steps fall between $a = 0.5$ and 1. For most haloes, the truncated fits hit the optimal middle ground in which we are not extrapolating over too little or not extrapolating at all.

After haloes are fitted, we return the final halo mass from the simulation, the model final mass (the best-fit parameter M_0 fitted over the entire halo), and a projected final mass (the best-fit parameter M_0 fitted over only the halo’s early history). Only haloes above $10^{10} h^{-1} M_\odot$ are analysed for each of the three mass categories.

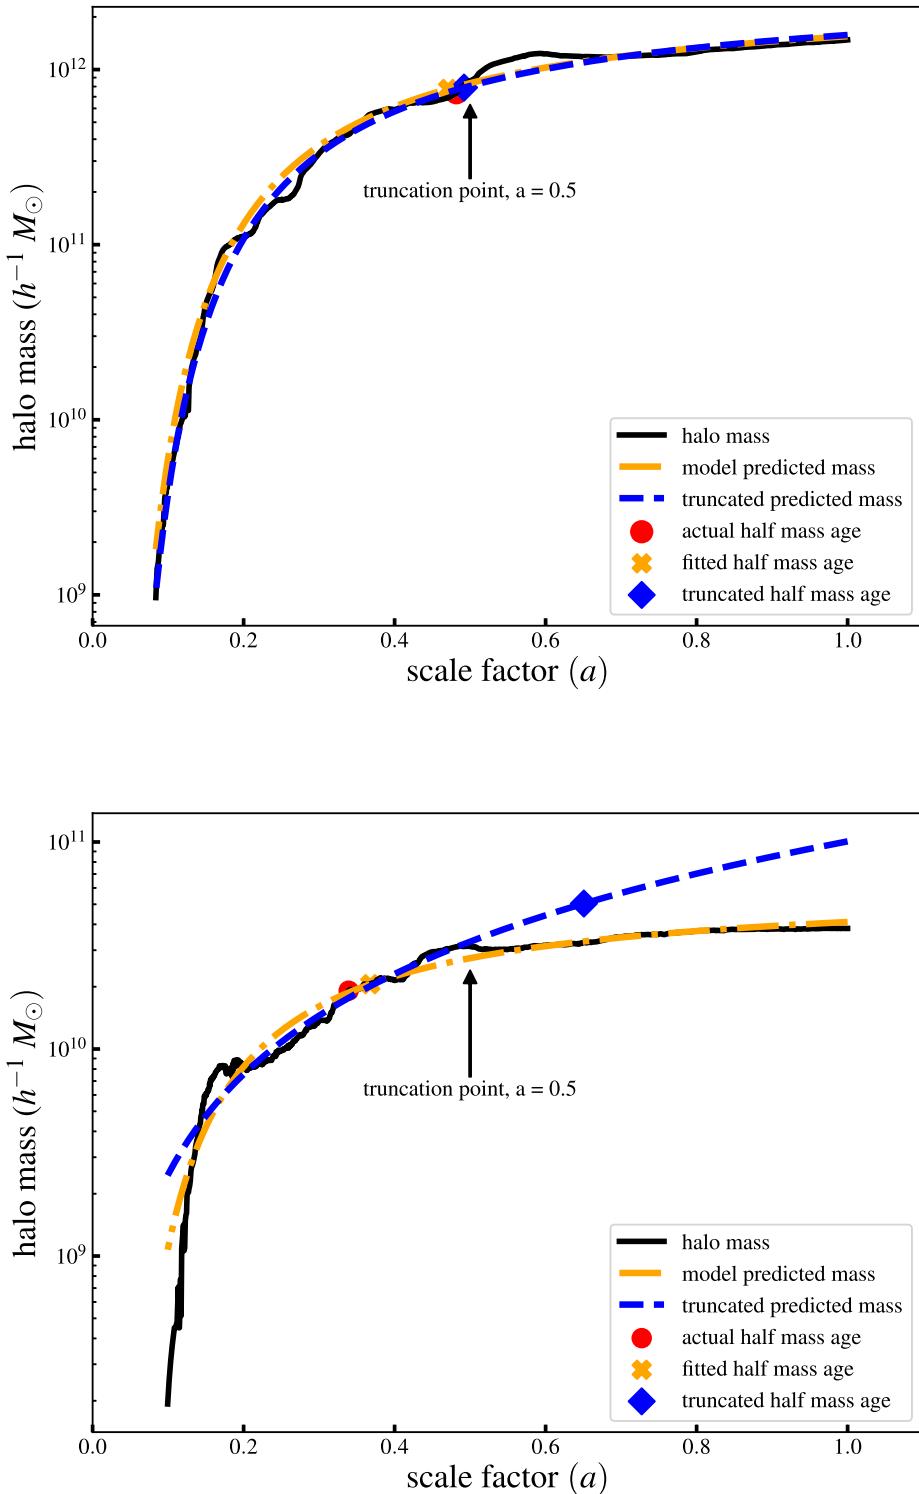


Figure 3. Two fits for a randomly chosen ‘regular’ halo (top panel) and arrested development halo (bottom panel) in the Vishnu simulation. The solid black curves are the actual mass accretion histories of the haloes as extracted from the simulation. The orange (dot-dashed) curves show the best-fit model accretion histories, using the fitting function in equation (3), when the fit is performed over the entire mass accretion histories. The blue (dashed) curves show the fit results when the fit is only performed over the earlier portion of the histories, up to scale factor $a = 0.5$. This scale factor is denoted by the black arrows. The blue curves at later times thus represent an extrapolation from the earlier history fits. The large points show the scale factors corresponding to the half-mass ages of the three histories. Fitting these haloes through both methods results in nearly identical fits with similar final predicted masses and ages for the ‘regular’ halo (top panel), but very different fits for the arrested development halo (bottom panel). In this latter case, the fitting function accurately predicts the final halo mass and age when fit over the entire history, but it predicts a higher final halo mass and a younger age when fit over only the early history.

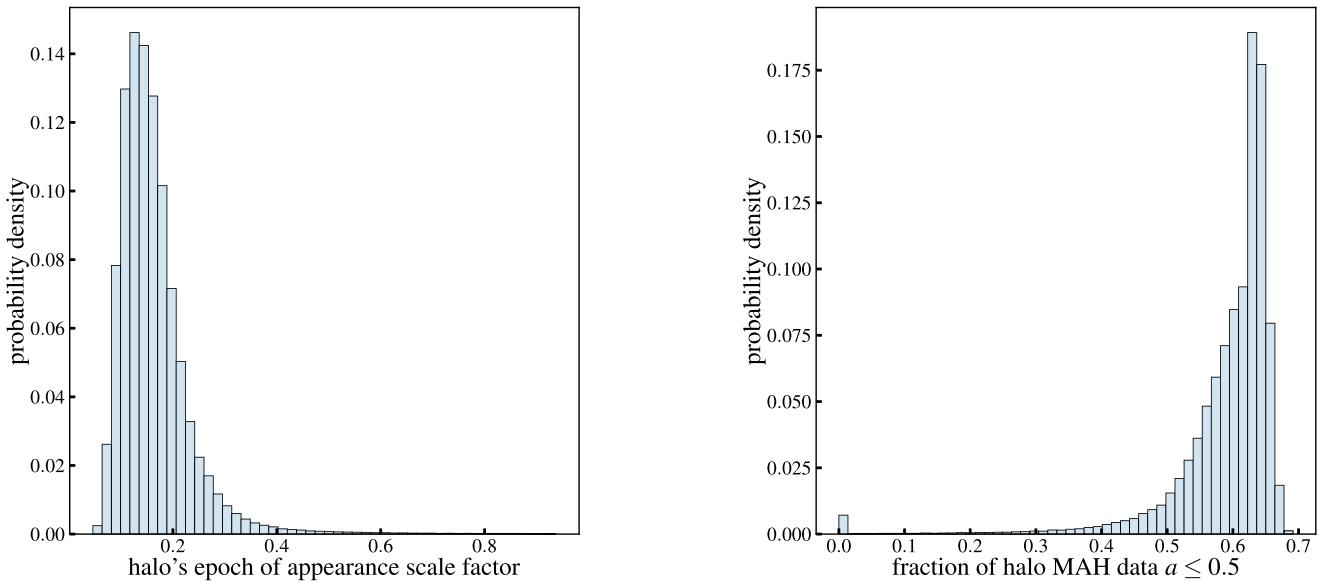


Figure 4. Tests to check whether haloes have enough of their history in place before the projected scale factor of $a = 0.5$, in order to robustly fit this early history and extrapolate halo growth to late times. *Left-hand panel:* the distribution of halo ‘epochs of appearance’, which is the earliest scale factor when the progenitor of a present-day halo was first resolved in the Vishnu simulation (following the main, most massive, trunk of the merger tree back in time). The distribution shows that the vast majority of haloes first appear well before $a = 0.5$, meaning that we are not excluding a significant number of haloes from our analysis due to them not existing before our chosen projected scale factor. *Right-hand panel:* the distribution of the fraction of time steps in Vishnu halo histories that occur before $a = 0.5$. The distribution shows that, for almost all haloes, a large fraction of their resolved history occurs before our chosen projected scale factor (with a peak near 60 per cent). This means that our fits to the early histories of haloes are based on a large enough number of epochs to be robust. The results in both panels apply to all haloes above $10^{9.5} h^{-1} M_\odot$.

However, we fit all haloes down to final masses of $10^{9.5} h^{-1} M_\odot$ to allow for the haloes that have actual masses below $10^{10} h^{-1} M_\odot$ but may have model and/or projected masses above $10^{10} h^{-1} M_\odot$ to be included in the analysis, as not including these haloes would bias the model and projected halo mass data. The $10^{9.5} h^{-1} M_\odot$ mass cut threshold was chosen to provide a wide enough mass range to reasonably include most haloes that fall into this category but without compromising the reliability of the halo fits from low-mass limits on halo resolution. We support this choice by noting that only roughly 5 per cent of the model haloes with a final mass greater than $10^{10} h^{-1} M_\odot$ were fitted from haloes with an original final mass between $10^{9.5}$ and $10^{10} h^{-1} M_\odot$, and only 0.3 per cent of model haloes with a mass greater than $10^{10} h^{-1} M_\odot$ had original halo masses less than $10^{9.8} h^{-1} M_\odot$.

We define half-mass age separately for each of the three mass accretion history types. The half-mass age of the halo is the point at which, following the most massive progenitor along the halo’s merger tree, the halo reaches half its present mass. The model half-mass age is the point at which the function fitted over the entire halo reaches half the final fitted mass, and the projected half-mass age is the point at which the projected function, fitted over the early portion of the mass accretion history, reaches half its final projected mass. All half-mass ages are given as scale factors, which means that larger values correspond to younger ages.

4 MEASURING THE ASSEMBLY BIAS SIGNAL

Following most previous works, we measure the halo assembly bias signal by comparing the clustering of old versus young haloes at fixed mass. To effectively fix mass, clustering statistics must be calculated in bins of halo mass that are narrow enough so that there is no residual

dependence of clustering on mass within the bin. At the same time, mass bins should be wide enough to contain a large enough number of haloes for calculating robust clustering statistics. This is especially pertinent given that we will make additional cuts on halo age within the mass bins, thus further reducing the halo sample size. In this work, we adopt eight logarithmic bins of mass from 10^{10} to $10^{11.6} h^{-1} M_\odot$ with each bin having a width of 0.2 dex. Within each bin, we rank haloes by their half-mass age and construct ‘old’ and ‘young’ samples that contain the 25 per cent oldest and 25 per cent youngest haloes. We find that this choice of binning strikes a good balance between the competing needs to control for mass while also having large enough samples of haloes for robust measurements of clustering. The age samples in our largest mass bin contain 152 966 haloes.

We use all three definitions of halo mass in this analysis, described in the previous section: the actual simulation mass, the model mass that results from the fit to the entire mass accretion history, and the projected mass that is extrapolated from the fit to the early portion of the halo’s history. Each definition of mass is accompanied by its own definition of half-mass age, as described in the previous section. For each definition of mass, we construct one sample containing all haloes in a mass bin, as well as the 25 per cent oldest and 25 per cent youngest samples. Thus, for each mass bin, we have nine sets of haloes for which to measure clustering.

To measure the clustering of a given halo sample, we calculate the two-point correlation function, $\xi(r)$, which is the excess probability of finding a pair of haloes separated by a physical distance r , over what would be expected for a random distribution. To calculate $\xi(r)$, we use the equation

$$\xi(r) = \frac{DD(r)}{RR(r)} - 1, \quad (4)$$

where $DD(r)$ is the number of halo pairs at separation distance r in the data, and $RR(r)$ is the number of pairs at separation distance r in a random set of data of equivalent size and volume. We use the CorrFunc package (Sinha & Garrison 2020), a suite of highly optimized and robustly tested codes for calculating clustering statistics. Since the Vishnu volume is a periodic cube, RR is calculated analytically.¹

To facilitate the interpretation of our results, we compute relative bias factors for old and young haloes in each mass bin. First, we divide $\xi(r)$ for old haloes by that for all haloes and take the square root of that ratio to obtain a relative bias as a function of scale, $b(r)$. Since assembly bias is a large-scale effect and bias functions become scale independent at scales larger than a few Mpc (Scherrer & Weinberg 1998; Narayanan, Berlind & Weinberg 2000), we average the values of bias in all bins corresponding to physical separations in the range $r: 5-30 h^{-1}$ Mpc. This yields a single relative bias factor b for the 25 per cent oldest haloes in a mass bin. We then do the same thing for the young haloes and for all our mass bins. The end result is a bias for old and young haloes as a function of halo mass, and for each age definition. An assembly bias signal corresponds to a value of b that is different from unity. Specifically, $b > 1$ for old haloes and $b < 1$ for young haloes (e.g. Salcedo et al. 2018).

To obtain uncertainties for our measurements, we assume that the main source of noise in the correlation function is Poisson noise from the DD term (since RR is calculated analytically). The Poisson uncertainty in the pair counts is \sqrt{DD} . We then propagate this error all the way to our final bias measurements.

5 RESULTS

Fig. 5 shows the resulting masses and ages from our fits compared to the actual masses and ages from the simulation, for all haloes in our analysis. Specifically, the figure shows the bivariate distribution of the fractional difference in age (fitted half-mass scale factor minus actual half-mass scale factor) on the x -axis versus the fractional difference in mass (fitted mass minus actual mass) on the y -axis. Each panel shows results for fits done using a different fit truncation point from $a = 0.5$ in the top-left-hand panel to $a = 1$ in the bottom-right-hand panel. Haloes for which fits produce masses and ages that match their actual masses and ages (for example, the ‘normal’ halo shown in the top panel of Fig. 3) appear in the middle of each panel, close to $(0, 0)$. On the other hand, arrested development haloes for which fits should produce masses that are larger and ages that are younger than their actual masses and ages (like the halo shown in the bottom panel of Fig. 3), appear in the top-right quadrant of each panel.

Fig. 5 shows several interesting features. First, we can see that the distributions of mass and age differences have a significant amount of scatter, meaning that haloes exhibit a large range of behaviours when we fit their mass accretion histories. However, some clear trends can be seen. Looking at the top left-hand panel where fits are performed using a truncation scale of $a = 0.5$, we see that most haloes lie in on a sequence running from the lower-left to the upper-right. Haloes on the upper part of this sequence have extrapolated fitted masses that are larger than their actual simulation masses, and ages that are younger (i.e. larger half-mass scale factors) than their actual simulation ages. These haloes have likely experienced arrested development, and our fits to their early histories are possibly estimating what their growth would have been without this effect. Correcting these haloes is the main goal of this study. However, there is also a significant population

of haloes that live in the lower-left part of the sequence. These haloes have extrapolated fitted masses that are smaller and ages than are older than their actual simulation masses and ages. One likely cause for this effect may be an inability of the fitting function to account for a major merger late in the halo’s history. A sudden mass spike at late times would cause the extrapolated fit to underestimate the final mass and overestimate the age. The sequence that we see in the figure thus represents different types of late-time mass growth, with haloes that experience late-time mergers on the bottom-left, and arrested development haloes on the top-right.

We now look at how fit results change as we shift the truncation scale to later times. From the different panels in Fig. 5, we see that the scatter is reduced for both mass and age, and the strong sequence we see for a truncation scale of $a = 0.5$ gradually disappears. For truncation scales close to $a = 1$, most haloes have fitted ‘model’ masses and ages that are close to their actual values from the simulation. This behaviour is expected since the fits in these cases take the haloes’ late-time evolution into account. Of particular note is that even for the case of the full model fit (the bottom right-hand panel of figure 5), in which haloes are fit over their entire history, most haloes have fractional mass and age differences that skew slightly positive. This means that our fitting function overpredicts the final mass of haloes by ~ 5 per cent, and overpredicts the half-mass scale factor of haloes by ~ 10 per cent. This could be a failure of our simple fitting function to properly describe the global shape of halo mass accretion histories, or it could be a systematic effect in the fits caused by small-scale variability in the accretion histories. Either way, we need to make sure that this small systematic effect does not impact our assembly bias analysis.

The primary result of this work is shown in Fig. 6, where we investigate how the assembly bias signal changes after ‘correcting’ halo masses and ages. The dashed and solid lines in Fig. 6 show the relative halo bias as a function of halo mass for the 25 per cent oldest and 25 per cent youngest haloes, respectively. The difference in clustering between old and young haloes at fixed mass is the assembly bias signal. Blue lines show results using the actual halo masses and ages from the simulation. This shows what previous authors have found (e.g. Salcedo et al. 2018), though our results extend to significantly lower halo masses. We find that the oldest 25 per cent haloes have a large-scale clustering bias that is ~ 1.65 times greater than that of all haloes, while the 25 per cent youngest haloes have a large-scale clustering bias that is ~ 0.6 times that of all haloes, at the smallest masses we study. The effect decreases slightly in higher mass bins.

The other sets of lines in Fig. 6 show results when we use halo masses and ages from our accretion history fits. The yellow lines show results for fits that are made over the entire accretion history (i.e. the ‘model’ fits). We note that using these model masses and ages does not alter the assembly bias signal, as is evidenced by the fact that the yellow and blue lines are in agreement. This is important because it demonstrates that the small systematic shifting of masses and ages that we observed in Fig. 5 does not introduce a significant change to the assembly bias signal. Consequently, any reduction of this signal from fitting to the early part of haloes’ histories and extrapolating to $a = 1$ can be explained by our correcting for arrested development, rather than systematic errors in the fitting process itself. The remaining sets of lines in Fig. 6 show results when we use masses and ages from fits to the early parts of halo histories. The lines show results for progressively earlier fit truncation scales, with brown, red, and cyan lines corresponding to truncation scales of $a = 0.9, 0.7$, and 0.5 , respectively. The figure shows clearly that the assembly bias signal is substantially reduced when we predict halo growth from

¹For more information, see: https://corrfunc.readthedocs.io/en/master/modules/rr_autocorrelations.html#rr-autocorrelations

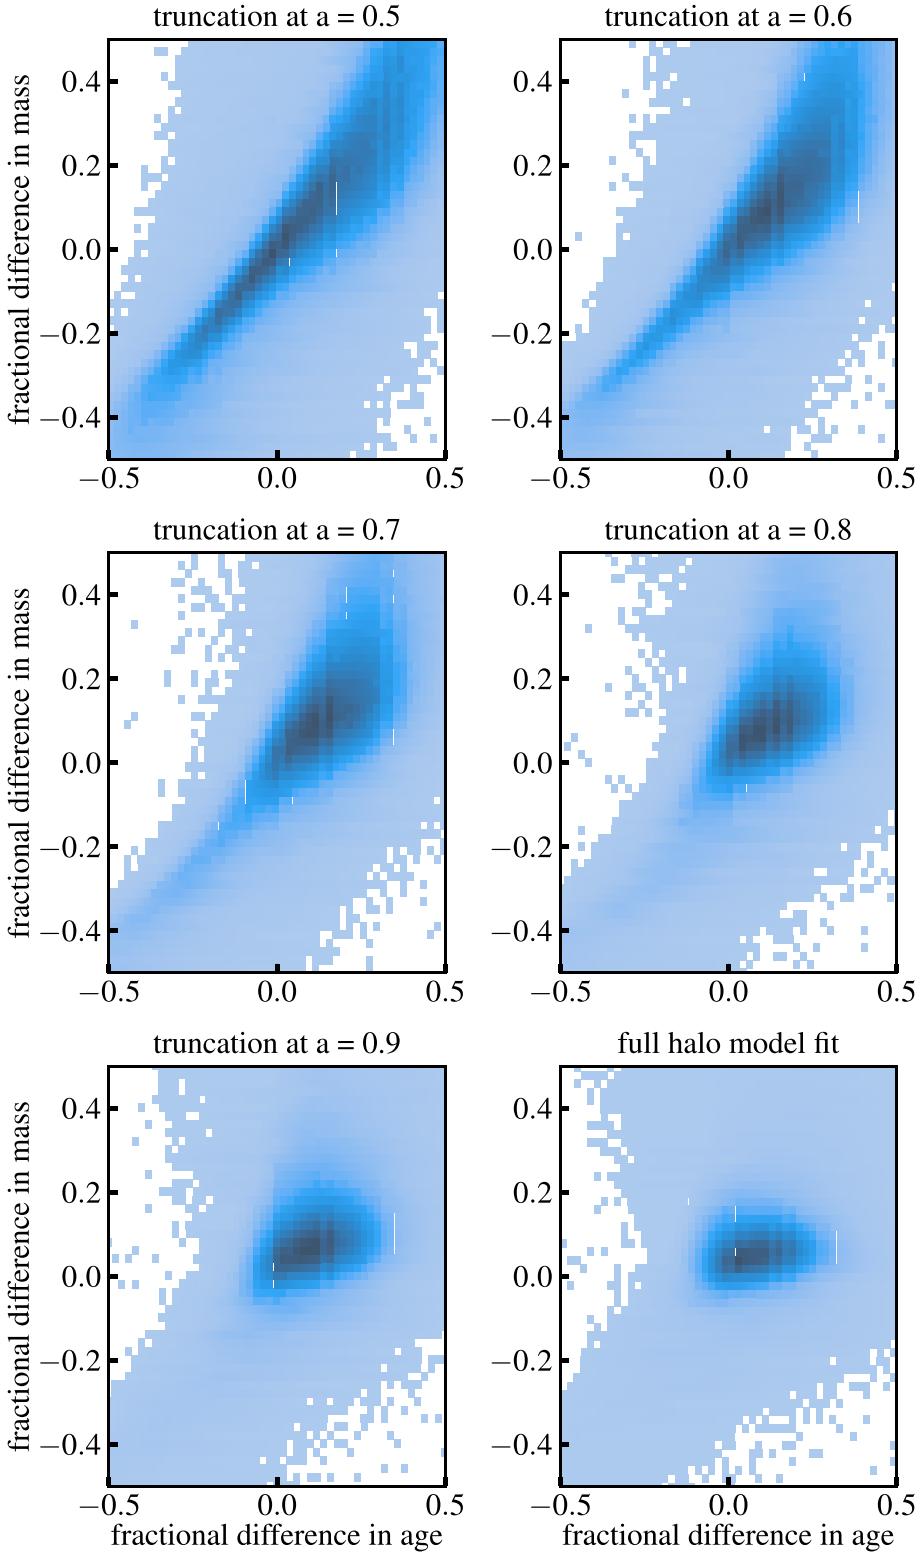


Figure 5. Bivariate distributions of fractional age and fractional mass differences between fitted and actual haloes. Each panel shows fits for different projected scale factor from $a = 0.5$ (top left) to $a = 1$ ('model fit', bottom right). Each panel shows the bivariate distribution of the fractional difference in the age (x-axis) versus fractional difference in mass (y-axis). Fractional differences are defined by taking the mass (or half-mass scale factor) from the fit minus the actual mass (or half-mass scale factor) from the simulation. Arrested development haloes are thus expected to appear in the top-right quadrant of each panel, since their fitted masses and half-mass scale factors should be larger than their actual values (see the example in Fig. 3). As expected, the fractional differences in mass and age decrease for most haloes when fits are done adopting later truncation times.

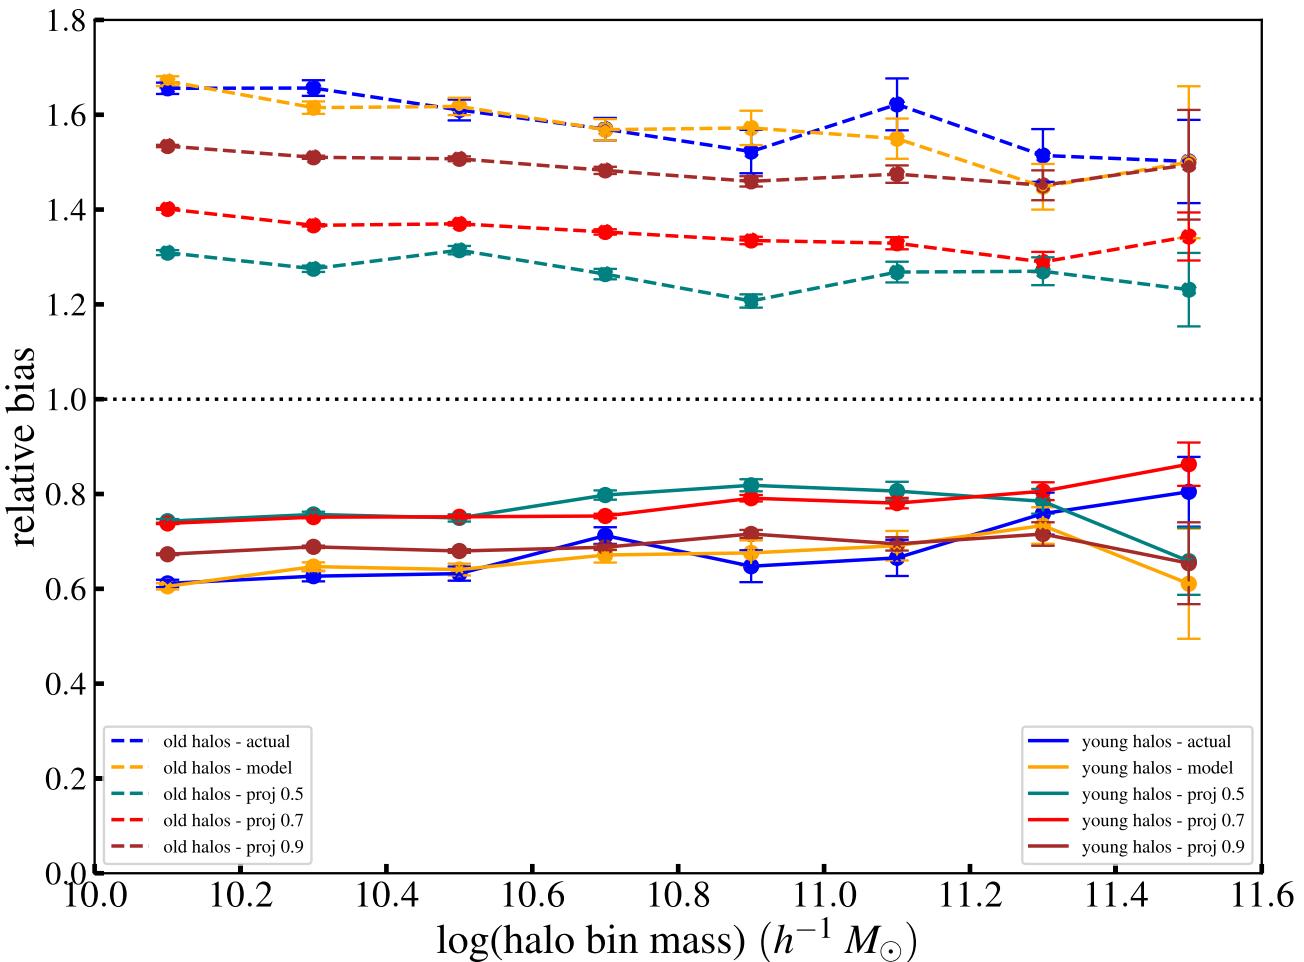


Figure 6. Relative bias of halo populations (y-axis) as a function of final halo mass (x-axis). Dashed lines show the large-scale clustering bias of the oldest 25 per cent haloes in a set of eight mass bins divided by the bias of all haloes in the same bins. Solid lines show the same, but for the 25 per cent youngest haloes in each mass bin. The difference between the relative bias of old and young haloes at fixed mass represents the halo assembly bias signal. Blue lines show results for when actual simulation halo masses and ages are used, while the other lines show results for when halo masses and ages come from fits to the haloes' mass accretion histories. Yellow lines show results for fits to haloes' entire histories, while brown, red, and cyan lines show results for fits out to progressively earlier truncation scales, but extrapolated to $a = 1$. This represents our attempt to correct for arrested development. The cyan lines show results for the earliest truncation scale of $a = 0.5$. There is a significant reduction in the assembly bias signal when haloes' masses and ages are predicted from their early histories.

its early history. In the case of the earliest fit truncation scale that we consider, $a = 0.5$, the signal for old haloes is reduced by about a half, while the signal for young haloes is reduced by a bit more than a third. The reduction in the assembly bias signal is less for fit truncation scales larger than $a = 0.5$, as expected. These results are consistent with our hypothesis that arrested development is at least partially responsible for the assembly bias of low-mass haloes and that we can mitigate this effect by fitting the early growth of haloes and extrapolating to predict what the late-time growth would have been in the absence of arrested development.

It is worth thinking about how our arrested development corrections are causing a reduction in the assembly bias signal. It is not immediately obvious how changing the masses and ages of haloes will affect their clustering within age quartiles because it depends on several factors, such as (i) how the change in haloes' masses and ages depends on their larger environment, (ii) how the relative change in mass and age changes a halo's location in the mass–age relation, and (ii) how the clustering of the overall halo population depends on mass.

Let us start by considering how our arrested development corrections depend on the environment. As we discussed in Section 1, several previous authors have demonstrated that assembly bias for low-mass haloes is caused by their tidal field environment (e.g. Hahn et al. 2009; Mansfield & Kravtsov 2020). It is thus interesting to investigate the tidal fields of haloes as a function of the degree of arrested development that they exhibit. To study the tidal field, we use the 'tidal force' quantity provided by the ROCKSTAR halo finding algorithm. For each halo, the algorithm calculates the tidal force from all its neighbours and finds the one that exerts the maximum tidal force on the halo, which is typically its closest more massive neighbour. ROCKSTAR then outputs the value of $R_{\text{halo}}/r_{\text{Hill}}$, where R_{halo} is the radius of the halo and r_{Hill} is the Hill radius due to the massive neighbour, which is defined as the distance from the halo at which the gravitational influence of the neighbour exceeds that of the halo. Small values of $R_{\text{halo}}/r_{\text{Hill}}$ thus correspond to weak tidal fields while large values correspond to strong tidal fields. Haloes that have a ratio larger than unity live in such strong tidal fields that matter in their own outskirts is more bound to the massive

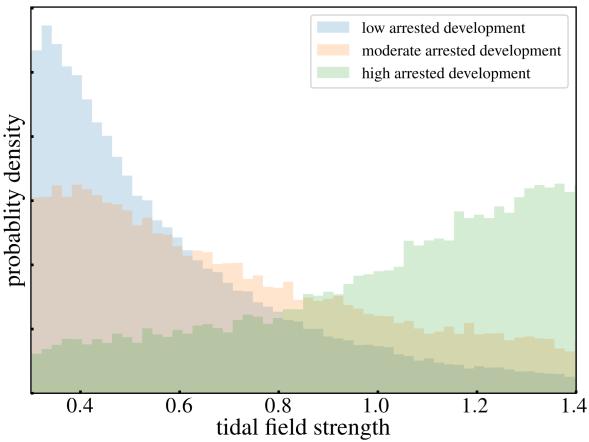


Figure 7. Distributions of tidal field strength for haloes with different degrees of arrested development. The shaded histograms show the normalized distributions of tidal field strength, where the strength of the tidal field affecting each halo is defined as the ratio of the halo's radius over the Hill radius due to the nearest massive neighbour. The three histograms show distributions for haloes that have experienced low, moderate, and high degrees of arrested development in their mass accretion histories. Haloes that have experienced a high degree of arrested development tend to live in regions containing strong tidal fields.

neighbour than to themselves. These haloes are likely in the process of being tidally stripped. To study how this tidal parameter depends on arrested development, we consider the mass accretion history fits using a projected scale factor of 0.5 and select subsamples of haloes based on the fractional difference in mass and age that they have between these extrapolated fits and their original histories (the top-left-hand panel of Fig. 5). We first make a sample where the fractional difference in mass and age is between -0.1 and 0.1 . These haloes exhibit little to no arrested development. We then make a ‘moderate’ sample by selecting all haloes whose fractional difference in mass and age is between 0.4 and 0.6 . Finally, we make a ‘high’ sample where the fractional difference in mass and age is greater than 1 . Fig. 7 shows the normalized distributions of tidal field strength ($R_{\text{halo}}/r_{\text{Hill}}$) for these three samples. The figure shows a dramatic difference between the samples. Haloes that exhibit a high degree of arrested development tend to live in very strong tidal field environments, while most haloes with little to no arrested development live in much weaker tidal fields. Since strong tidal fields are caused by massive neighbours, it is safe to assume that haloes in strong tidal fields live in highly clustered regions because they inherit the strong clustering of their massive neighbours.

Let us now examine how our mass and age corrections change the location of a halo in the mass–age relation. Fig. 8 shows the masses and ages (half-mass scale factors) of all our haloes (grey points), as well as the median, 25th and 75th percentiles of age as a function of mass (solid and dotted black lines, respectively). The plot shows the well-known mass–age relation for haloes whereby more massive haloes tend to be younger (i.e. have assembled their mass more recently) than less massive haloes. The red X symbol on the plot shows the original location of the example halo whose history is shown in the bottom panel of Fig. 3. This halo’s age places it in the oldest age quartile for its mass. The fact that this halo likely lives in a strong tidal field and is thus in a highly clustered region perfectly demonstrates the assembly bias effect: old haloes are more clustered than young haloes at fixed mass. Now we consider the mass and age corrections that we apply with the intention of ‘erasing’ arrested

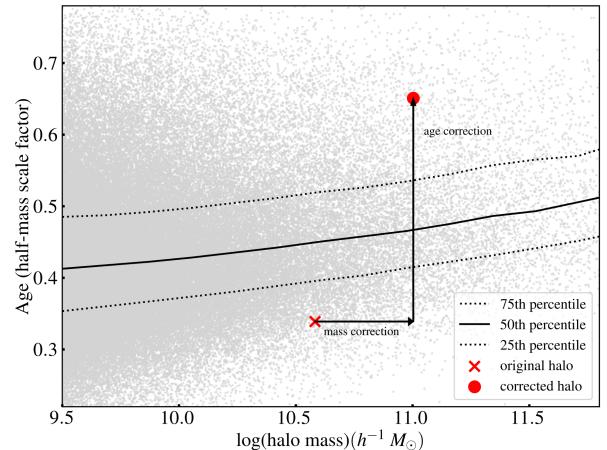


Figure 8. Illustration of the effect of our arrested development correction on the mass and age of a halo relative to the full halo population. Grey dots show the masses and ages (half-mass scale factors) of a random subset of all haloes in our sample. The median age as a function of mass is shown by the solid black line, while the 25th and 75th percentiles are shown by the dotted lines. The red X symbol shows the original mass and age of the example halo whose mass accretion history is shown in the bottom panel of Fig. 3. The red circle shows the adjusted mass and age of this halo after applying our arrested development corrections, and the arrows illustrate the mass and age corrections separately. Correcting for arrested development results in the halo having a higher mass and younger age, such that it ends up in a different age quartile for its mass.

development. First, the mass of this halo will increase, which is shown by the horizontal arrow. The effect of this change on the assembly bias signal is complicated. Due to the sloped mass–age relation, this halo that lives in a highly clustered region is now even older in the age ranking of haloes at fixed mass. This would make the assembly bias signal even stronger than before. On the other hand, the clustering of haloes increases with mass, which means that, at this new mass, the halo’s clustering will not be as enhanced relative to the total population as it was before. This would make the assembly bias signal weaker than before. Since these two effects work in opposite directions, the net effect on assembly bias depends on how strongly halo clustering increases with mass. This has been studied extensively by many past works, which find that the large-scale clustering of haloes is fairly constant at low masses and starts to increase exponentially at masses close to the non-linear mass M_* (e.g. Sheth & Tormen 1999; Seljak & Warren 2004; Tinker et al. 2010). The low-halo masses that we are considering here are well below M_* , so the clustering of haloes is mostly independent of mass. As a result, a mass correction that moves highly clustered haloes to larger masses should have a net effect of strengthening the assembly bias signal. However, we also correct halo ages, which become correspondingly younger as we increase their masses. The vertical arrow in Fig. 8 shows the age correction that we apply to our example halo, which places it at a corrected location indicated by the red circle. The halo is now in the youngest age quartile at fixed mass, which will certainly contribute to a reduction of the assembly bias signal. This occurs because the ratio of the age shift to the mass shift for this halo exceeds the slope in the halo mass–age relation. The reduction in the global assembly bias signal that we see in Fig. 6 takes place because the arrested development corrections that we obtain from our extrapolated mass accretion history fits behave in a qualitatively similar way as in the case of the example halo we considered here.

6 SUMMARY AND DISCUSSION

In this paper, we have tested a new method for probing halo arrested development as a cause of halo assembly bias in the low mass regime. Furthermore, our method provides a potential way to mitigate the effects of assembly bias in N -body simulations. We define halo arrested development as the late-time slow-down of mass accretion in some low-mass haloes due to the tidal field of the surrounding mass distribution. Arrested development thus results in haloes having lower final masses and older measured ages than they would have had in the absence of arrested development. These haloes live in dense large-scale environments and are thus more clustered and older than other haloes of the same final mass, leading to an assembly bias signal. The new idea that we introduce and test is that we can use the early part of a halo's mass accretion history, before arrested development has kicked in, to predict what the late-time growth of the halo would have been in the absence of arrested development. We do this by fitting an analytic function to the mass accretion history at early times, for example before $a = 0.5$, and extrapolating the fit to $a = 1$. We then calculate a 'corrected' mass and age for the halo based on this fit. Using these corrected masses and ages of haloes, we check to what extent the assembly bias signal is affected. Our results show that this method leads to a substantial decrease in the assembly bias signal. This supports the hypothesis that arrested development is a major cause of assembly bias, in agreement with (Mansfield & Kravtsov 2020), and validates our new approach to mitigating the effect.

The mass accretion fitting method cuts the assembly bias signal by a half for the 25 per cent oldest haloes and by a little more than a third for the 25 per cent youngest haloes, in our lowest mass bins of $\sim 10^{10} h^{-1} M_{\odot}$. Though this result is promising, our method does not yet produce the magnitude of reduction that would lend it useful as a method of fully mitigating assembly bias in cosmological N -body simulations. It is thus worth thinking about why the method does not fully remove assembly bias. One possibility is that arrested development is not solely responsible for low-mass assembly bias. This is certainly possible, but before reaching this conclusion, we must first gain confidence that our methodology is able to fully remove the effects of arrested development and does not introduce any additional spurious assembly bias signal.

There are several potential reasons why our fitting method may not be fully accounting for and removing arrested development. One possibility is that the effects of arrested development begin to impact a halo at very early times, before the scale $a = 0.5$ that represents the earliest fit truncation scale we adopt. Fig. 6 shows a decrease in the assembly bias signal for old haloes when we shift the truncation scale from $a = 0.7$ to 0.5 so perhaps the signal would further reduce by restricting the fit to even earlier times. However, fits might start to become unreliable if we restrict them to too small a portion of the history. Also, the assembly bias signal for young haloes does not change from $a = 0.7$ to 0.5 , making it unlikely that pushing to earlier times would affect young haloes. Another possibility is that the early mass growth of a halo has limited power to predict its late-time growth. The extrapolation of our fits to $a = 0$ may thus contain a fair amount of stochasticity or even systematic effects that fail to properly remove the effects of arrested development in many haloes. Alternatively, it may be the case that the early growth of haloes does indeed have predictive power, but that our choice of fitting function is the limiting factor. A more sophisticated model that takes into account a halo's entire merger tree rather than just the one-dimensional mass history following the most massive progenitor, might be more successful in predicting late-time growth.

A different possibility is that our fitting methodology is introducing a spurious assembly bias signal. Fig. 5 revealed that many haloes have 'corrected' masses that are smaller than the original masses from the simulation. These are likely haloes that experience fast growth at late times, e.g. through a late major merger, which is the opposite of arrested development. Applying our mass and age correction to these systems might be inducing a fake assembly bias signal. An investigation of ways to identify which haloes require a correction deserves further study.

These are difficult issues to disentangle because arrested development is not a simple binary process, where a halo either experiences it or not. In that case, one could identify a set of 'normal' haloes that do not experience arrested development and use it to validate that a given correction method does not alter those haloes. However, the reality is likely more complex, with every halo being exposed to some degree of tidal field and having its mass accretion impacted by its larger-scale environment to some extent. One could perhaps approximately identify a sample of 'normal' haloes by selecting haloes that do not have a nearby massive neighbour. This is worth future study. For haloes that do experience arrested development, we have no way of knowing how they would have evolved in the absence of this effect, making it impossible to test how well our correction method works on a halo-by-halo basis. We can only test this in the aggregate by examining the assembly bias signal of all haloes – if the assembly bias signal vanishes then our correction method must be working. Of course, this only holds true under the assumption that arrested development is the sole cause of assembly bias. The main goal of this work was to test a new idea for probing and mitigating the effect of assembly bias. Our preliminary results are encouraging and thus warrant further study along the various lines discussed in this section.

ACKNOWLEDGEMENTS

Parts of this research were conducted at Vanderbilt University through the Fisk-Vanderbilt Bridge to Ph.D. Program and the Establishing Multimessenger astronomy Inclusive Training (EMIT) program. Parts of this research were conducted by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013. Computational facilities used in this project were provided by the Vanderbilt Advanced Computing Center for Research and Education (ACCRE). This research made use of NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), MATPLOTLIB, a PYTHON library for publication quality graphics (Hunter 2007), and corrfunc (Sinha & Garrison 2020). We thank the anonymous referee for asking probing questions that led to an improvement of the paper.

DATA AVAILABILITY

The data underlying this article is available on request and will be shared on reasonable request to the corresponding author.

REFERENCES

- Avila-Reese V., Firmani C., Hernandez X., 1998, *ApJ*, 505, 37
- Behroozi P. S., Wechsler R. H., Wu H.-Y., 2012a, *ApJ*, 762, 109
- Behroozi P. S., Wechsler R. H., Wu H.-Y., Busha M. T., Klypin A. A., Primack J. R., 2012b, *ApJ*, 763, 18
- van den Bosch F. C., 2002, *MNRAS*, 331, 98
- Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80

Bullock J. S., Kolatt T. S., Sigad Y., Somerville R. S., Kravtsov A. V., Klypin A. A., Primack J. R., Dekel A., 2001, *MNRAS*, 321, 559

Correa C. A., Wyithe J. S. B., Schaye J., Duffy A. R., 2015, *MNRAS*, 450, 1514

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

Diemer B., 2021, *ApJ*, 909, 112

Faltenbacher A., White S. D. M., 2010, *ApJ*, 708, 469

Gao L., White S. D. M., 2007, *MNRAS*, 377, L5

Gao L., Springel V., White S. D. M., 2005, *MNRAS*, 363, L66

Hahn O., Porciani C., Dekel A., Carollo C. M., 2009, *MNRAS*, 398, 1742

Harris C. R. et al., 2020, *Nature*, 585, 357

Hearin A. P., Behroozi P. S., Van Den Bosch F. C., 2016, *MNRAS*, 461, 2135

Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90

Johnson J. W., Maller A. H., Berlind A. A., Sinha M., Holley-Bockelmann K., 2019, *MNRAS*, 486, 1156

Lacerna I., Padilla N., 2012, *MNRAS*, 426, L26

Lazeyras T., Musso M., Schmidt F., 2017, *J. Cosmol. Astropart. Phys.*, 2017, 059

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

Li Y., Mo H. J., Gao L., 2008, *MNRAS*, 389, 1419

Ludlow A. D. et al., 2013, *MNRAS*, 432, 1103

Mansfield P., Kravtsov A. V., 2020, *MNRAS*, 493, 4763

Mcbride J., Fakhouri O., Ma C.-P., 2009, *MNRAS*, 398, 1858

Narayanan V. K., Berlind A. A., Weinberg D. H., 2000, *ApJ*, 528, 1

Ryden B. S., Gunn J. E., 1987, *ApJ*, 318 15

Salcedo A. N., Maller A. H., Berlind A. A., Sinha M., Mcbride C. K., Behroozi P. S., Wechsler R. H., Weinberg D. H., 2018, *MNRAS*, 475, 4411

Scherrer R. J., Weinberg D. H., 1998, *ApJ*, 504, 607

Scoccimarro R., 1998, *MNRAS*, 299, 1097

Seljak U., Warren M. S., 2004, *MNRAS*, 355, 129

Sheth R. K., Tormen G., 1999, *MNRAS*, 308, 119

Sheth R. K., Tormen G., 2004, *MNRAS*, 350, 1385

Sinha M., Garrison L. H., 2020, *MNRAS*, 491, 3022

Sinha M., Holley-Bockelmann K., 2012, *ApJ*, 751, 17

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

Sunayama T., Hearn A. P., Padmanabhan N., Leauthaud A., 2016, *MNRAS*, 458, 1510

Tasitsiomi A., Kravtsov A. V., Gottlöber S., Klypin A. A., 2004, *ApJ*, 607, 125

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

Villarreal A. S. et al., 2017, *MNRAS*, 472, 1088

Virtanen P. et al., 2020, *Nature Methods*, 17, 261

Wang H. Y., Mo H. J., Jing Y. P., 2007, *MNRAS*, 375, 633

Wechsler R. H., Tinker J. L., 2018, *ARA&A*, 56, 435

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

Wetzel A. R., Tinker J. L., Conroy C., Bosch F. C. v. d., 2014, *MNRAS*, 439, 2687

Wong A. W. C., Taylor J. E., 2012, *ApJ*, 757 102

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