Halo exclusion criteria impacts halo statistics

Rafael García[★] and Eduardo Rozo

Department of Physics, University of Arizona, Tucson, AZ 85721, USA

Accepted 2019 August 30. Received 2019 August 30; in original form 2019 April 1

ABSTRACT

Every halo-finding algorithm must make a critical yet relatively arbitrary choice: it must decide which structures are parent haloes, and which structures are subhaloes of larger haloes. We refer to this choice as *percolation*. We demonstrate that the choice of percolation impacts the statistical properties of the resulting halo catalogue. Specifically, we modify the halo-finding algorithm ROCKSTAR to construct three different halo catalogues from the same simulation data, each with identical mass definitions, but different choice of percolation. The resulting haloes exhibit significant differences in both halo abundance and clustering properties. Differences in the halo mass function reach 6 per cent for haloes of mass $10^{13} h^{-1} M_{\odot}$, larger than the few per cent precision necessary for current cluster abundance experiments such as the Dark Energy Survey. Comparable differences are observed in the large-scale clustering bias, while differences in the halo–matter correlation function reach 30 per cent on translinear scales. These effects can bias weak-lensing estimates of cluster masses at a level comparable to the statistical precision of current state-of-the-art experiments.

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

1 INTRODUCTION

In the halo model, the abundance and distribution of galaxies and clusters are linked to the abundance and distribution of dark matter haloes (Cooray & Sheth 2002). Predicting the properties of haloes requires large computer simulations that map the matter distribution of the Universe. The output of simulations is then analysed using a halo finder to find gravitationally bound dark matter structures.

Every halo-finding algorithm makes two critical yet relatively arbitrary choices. The first has received plenty of attention, and is the definition of halo mass. Halo mass is typically defined as the mass enclosed within some specific spherical aperture, chosen such that the mean density of the halo within that sphere is equal to some factor of either the critical density or the mass density of the Universe. However, other definitions are also commonly used (e.g. friends-of-friends) (see e.g. Knebe et al. 2013). For this reason, one can find calibrations of the halo mass function for multiple different halo mass definitions (Tinker et al. 2008; Bhattacharya et al. 2011; McClintock et al. 2019). The second arbitrary choice has received little attention to date, namely, how a halo-finding algorithm decides which structures are parent haloes, and which are subhaloes that ‘belong’ to a larger halo. We refer to the criteria for categorizing structures as parent haloes versus subhaloes as *percolation* or *exclusion criteria*. There is currently no standard percolation scheme, with different halo finders applying different halo exclusion criteria when constructing halo catalogues.

In this paper, we show that the choice of percolation impacts the statistical properties of the resulting halo population at a non-negligible level. To do so, we modify ROCKSTAR (Behroozi, Wechsler & Wu 2013), a state-of-the-art halo-finding algorithm, to generate halo catalogues with identical mass definitions, but different halo exclusion criteria. For each such halo catalogue we measure the halo mass function, correlation function, and halo–mass correlation function. By comparing these properties of the resulting halo catalogues to the properties of the fiducial ROCKSTAR catalogue we quantify the level of systematic uncertainty in current theoretical predictions associated with the choice of percolation algorithm implemented in the construction of the catalogue.

2 SIMULATION DATA

We use a cosmological N -body simulation run using the publicly available code GADGET2 (Springel 2005). This simulation is similar to the simulations used for the Aemulus project (DeRose et al. 2018). Specifically, the simulation has a box size of $1050 h^{-1}$ Mpc with 1400^3 particles and was run using periodic boundary conditions with a force softening scale of $20 h^{-1}$ kpc. The cosmology of the simulation we use is $h = 0.6704$, $\Omega_m = 0.318$, $\Omega_{\Lambda} = 0.682$, $\Omega_b = 0.049$, $\sigma_8 = 0.835$, $n_s = 0.962$. The particle mass is $3.7275 \times 10^{10} h^{-1} M_{\odot}$. Since the goal of this study is to highlight the underappreciated impact of halo exclusion on halo statistics, a single simulation suffices for our purposes.

We compare the statistics of halo catalogues generated from the same simulation box, using the same halo mass definition, namely M_{200m} , but different halo exclusion criteria. The specific statistics we consider are

[★] E-mail: rgarciamar@email.arizona.edu

- (i) the halo mass function,
- (ii) the large-scale halo clustering bias,
- (iii) the halo–mass correlation function,
- (iv) and the halo–halo correlation function.

The different halo catalogues are created by modifying the publicly available code ROCKSTAR (Behroozi et al. 2013).

ROCKSTAR uses a friends-of-friends algorithm in six-dimensional phase space to find seed dark matter structures. It then iteratively assigns particles in the simulation to each seed group, merging seeds into a single halo when the separation between haloes is sufficiently small (equation 2 in Behroozi et al. 2013). ROCKSTAR then removes all unbound particles from each halo, and computes the spherical mass and radius as defined using a virial overdensity criteria. Specifically, the mass M_Δ and spherical radius R_Δ associated with each halo are selected such that they satisfy the constraint equation

$$\frac{4}{3}\pi R_\Delta^3 \Delta \bar{\rho}_m = M_\Delta. \quad (1)$$

Here, M_Δ is the mass contained within the radius R_Δ , and Δ is the overdensity calculated using the spherical collapse model of Bryan & Norman (1998). While the virial overdensity criterion is the default for ROCKSTAR, given the final halo catalogue one can readily recompute strict spherical overdensity masses, i.e. masses using all particles, without any unbinding procedure. In our work, we always define halo mass using strict spherical overdensity masses with a fixed overdensity threshold of $\Delta = 360$ relative to the mean matter density. Finally, ROCKSTAR percolates the seed structure catalogue to generate a final halo catalogue by determining which seed structures are subhaloes of the parent halo centred on a larger substructure. The classification of a seed structure as a halo or a subhalo is dependent upon the phase-space distance of the seed structure to all larger seeds, and incorporates information from the previous time-step when available. It is the impact of this percolation step on the statistical properties of the resulting halo catalogue that we investigate in this work.

3 PERCOLATION SCHEMES

We create alternate halo catalogues starting from the seed structures identified by ROCKSTAR by changing the default percolation in the code. First, we trim the list of seed structures to those above a mass threshold of $M_{360m} \geq 10^{12.5} h^{-1} M_\odot$ (~ 100 particles). Seed structures are then classified as haloes or subhaloes using a simple spherical exclusion criterion. Specifically, we rank order all seed structures according to their maximum circular velocity, defined as the maximum of the circular velocity profile

$$V(r) = [GM(< r)/r]^{1/2}. \quad (2)$$

Initially, all seed structures are considered candidate haloes. Starting from the top-ranked (largest) candidate halo, we apply a spherical exclusion criteria to identify substructures of the halo centred on the top-ranked seed structure. Specifically, given two structures of mass M_1 and M_2 , the two structures are considered to fall within the same parent halo if the separation between the two structures x_{12} satisfies

$$x_{12} \equiv |\mathbf{r}_1 - \mathbf{r}_2| \leq d(M_1, M_2), \quad (3)$$

where d is the halo exclusion function. All seed structures identified as substructures of a larger parent halo are removed from the candidate halo list, and the procedure is iterated with the next highest ranked candidate halo until no more candidate haloes remain.

We consider three different choices for halo exclusion:

(i) **Soft-sphere halo exclusion:** Two seed structures are considered to be in the same parent halo if their separation is less than the radius of the larger structure, i.e. $d(M_1, M_2) = R_{360}(\max(M_1, M_2))$. This is the halo exclusion criteria used in ROCKSTAR. Tinker et al. (2008) and McClintock et al. (2019) use this percolation scheme with a different mass definition: $d(M_1, M_2) = R_{200}(\max(M_1, M_2))$.

(ii) **Point-mass exclusion:** Two seed structures are considered to be in the same parent halo if their separation is less than the radius of a structure of mass $M_1 + M_2$, i.e. $d(M_1, M_2) = R_{360}(M_1 + M_2)$. This halo exclusion criterion self-consistently enforces strict spherical overdensity mass definitions and exclusion when the haloes can be approximated as point particles.

(iii) **Hard-sphere exclusion:** Two seed structures are considered to be in the same parent halo if the spherical volumes associated with each structure overlap at all, i.e. $d(M_1, M_2) = R_{360}(M_1) + R_{360}(M_2)$.

The three percolation schemes are illustrated in Fig. 1. There we apply each of our proposed percolations to the same set of halo seeds in an illustrative example in which a large mass halo is surrounded by several smaller halo seeds. The boundaries we draw are circles of radius $R_{360m}(M)$. The large halo is the top-ranked parent halo. The classification of the rest of the halo seeds as haloes or subhaloes depends on the percolation scheme: soft-sphere (left), point-mass exclusion (centre), or hard-sphere exclusion (right). The crossed seeds are removed from the final halo catalogue and are identified as substructures of the bigger halo in each of the percolation schemes. In terms of the amount of haloes removed from the candidate halo list, the soft-sphere scheme (left) is the most conservative scheme, and the hard-sphere scheme (right) is the most aggressive one.

4 IMPACT ON HALO STATISTICS

We characterize the impact of the percolation scheme used to generate the halo catalogue on four different halo statistics: the halo mass function, the halo–matter correlation function, the halo–halo correlation function, and the large-scale clustering bias. In all cases, we use strict spherical overdensity masses to define the mass of a halo. Because the mass definition itself is constant, any differences in the statistics of the four halo catalogues we generated must be the direct result of the different percolation schemes.

4.1 Halo mass function

Fig. 2 compares the halo mass functions of the halo catalogues generated using each of the three different percolation algorithms (fiducial, soft-sphere, point-mass, and hard-sphere). The lower panel show the fractional difference in the halo mass function relative to the fiducial percolation. The impact of percolation is clearly negligible at the high-mass end, but can become significant at low halo masses. This makes sense. Since high-mass haloes dominate their environment, the impact of percolation on these haloes is negligible: these haloes are never assigned as subhaloes of more massive systems. By contrast, the more aggressive percolation schemes remove small haloes from the immediate vicinity of large haloes, thereby suppressing the resulting halo abundance at low masses. The largest difference is that between the fiducial ROCKSTAR percolation algorithm and the hard-sphere exclusion, reaching ~ 6 per cent (4 per cent) differences for mass $M \sim 10^{13} h^{-1} M_\odot$ ($M \sim 10^{14} h^{-1} M_\odot$) haloes.

Figure 1. Qualitative illustration of different percolation schemes. The red line indicates the removal of a halo seed from the halo catalogue. *Left:* Soft sphere scheme, which removes all of the seeds whose centre resides within the radius of another halo seed. This is the ROCKSTAR fiducial percolation scheme. *Right:* Hard-sphere scheme, which removes all of the seeds that overlap with a bigger halo seed. *Centre:* Point-Mass Approximation (PMA) scheme, which takes into account the mass of both halo seeds in question. This scheme allows overlaps between haloes but may also remove seeds whose centre resides outside of other haloes.

Figure 2. The halo mass function for distinct percolation schemes. The choice of percolation has a significant impact on halo abundance. *Top:* Halo mass functions. *Bottom:* Fractional difference of the halo mass functions with respect to the fiducial. The shaded regions represent 68 per cent confidence intervals as determined by jackknifing.

The impact of halo definition on halo abundance has been pointed out in previous work (e.g. Bocquet et al. 2016, and references therein). They show that the differences between their halo mass function and Tinker et al. (2008) are due to the use of different halo definitions. Specifically they compare halo mass functions at the same overdensity threshold Δ_{500c} , but different percolation schemes. They observe large differences at the low-mass regime (25 per cent at $M_{500c} \simeq 10^{12} h^{-1} M_{\odot}$) and a negligible impact at high masses, consistent with our results.

The differences illustrated in Fig. 2 are larger than the ≈ 1 per cent precision necessary for stage III dark energy experiments such as the Dark Energy Survey, and significantly larger than the precision reached by current halo mass function emulators (e.g. McClintock et al. 2019). Evidently, while we can make very precise predictions for the halo mass function given a halo-finding algorithm, it is clear that the choice of percolation introduces a significant amount

of systematic uncertainty in our predictions. Moreover, this level of systematic uncertainty is irreducible so long as haloes in simulations are percolated in a way that is different from the way clusters are percolated in real data sets. Our results demonstrate that implementing identical percolation algorithms across both simulated haloes and real clusters is necessary for stage III and IV dark energy experiments.

4.2 Halo–mass correlation function

Fig. 3 shows the halo–mass correlation function measured for each of our halo catalogues (top row), and the relative difference in the halo–mass correlation functions relative to that measured in our fiducial catalogue (bottom row). The left and right columns correspond to haloes with masses in the range $[3 \times 10^{12}, 5 \times 10^{12}] h^{-1} M_{\odot}$ and $[2 \times 10^{14}, 5 \times 10^{14}] h^{-1} M_{\odot}$, respectively. We plot data for these same mass bins through the rest of this paper. We find that aggressive halo exclusion criteria lead to suppression of the halo–mass correlation function ξ_{hm} . There is an obvious large feature on translinear scales ($\sim 1 h^{-1}$ Mpc to a few h^{-1} Mpc), along with a constant change in the clustering amplitude at large scales. The large (up to 30 per cent difference) feature at translinear scales makes sense: the more aggressive exclusion criteria remove low-mass haloes in the vicinity of high-mass haloes. Consequently, the amount of mass in the immediate neighbourhood of the remaining low-mass haloes is suppressed, leading to a large decrease in the halo–mass correlation function. The typical length scale associated with this effect is the exclusion radius of the largest dark matter haloes.

The fact that there is an overall offset in the clustering amplitude of haloes at large scales may seem surprising at first sight. However, this too is easily explained. A stronger halo exclusion region removes more small haloes from the vicinity of large haloes. Since large haloes live in high-density regions, the surviving low-mass haloes must necessarily be less clustered. As for the halo mass function, these trends are more pronounced for low-mass haloes than they are for high-mass haloes, and for the same reason: high-mass haloes dominate their environment, and are therefore rarely removed through percolation.

Figure 3. Halo–mass correlation function for different percolation schemes and mass bins. *Top*: Halo–mass correlation function. *Bottom*: Fractional differences with respect to the fiducial percolation. *Left*: Low-mass haloes. *Right*: High-mass haloes. Error bars are jackknife.

We characterize the large-scale clustering amplitude in terms of the halo bias, defined as the ratio of the halo–matter cross-correlation function and the matter autocorrelation function

$$b(r|M) = \frac{\xi_{hm}(r|M)}{\xi_{mm}(r)}, \quad (4)$$

where M is the mass of the halo. On large scales ($r \geq 10 h^{-1} \text{ Mpc}$), the halo bias is approximately constant, as expected. We fit a constant bias model to the data in the radial range $10 < r < 80 h^{-1} \text{ Mpc}$ to arrive at our final value for the large-scale bias $b(M)$ in each of our four halo catalogues. Interestingly, we find that the ratio ξ_{hm}/ξ_{lin} , where ξ_{lin} is the linear correlation function, is *not* constant over the same scales. That is, the linear-bias approximation is valid to significantly lower scales provided the reference clustering function is the matter correlation function rather than the linear correlation function.

The left-hand panel in Fig. 4 shows the fractional difference of the large-scale bias between our different halo catalogues and the fiducial ROCKSTAR catalogue. The data points show the relative bias as measured using the halo–halo correlation function (see the next section for details), whereas the coloured bands represent our measurements using the halo–mass correlation function. The width of the band is set by the error in our measurement. As expected, the bias of the high-mass haloes is largely insensitive to percolation effects, whereas the bias of low-mass ($M \sim 10^{13} h^{-1} \text{ M}_\odot$) haloes changes by as much as ≈ 6 per cent.

The large-scale clustering amplitude of cosmological objects is often used as a way to estimate the mass of the haloes hosting those objects (e.g. Robertson 2010; Mountrichas et al. 2016). The dependence of the clustering amplitude on the choice of percolation algorithm demonstrates that these type of estimates can be subject to large systematic uncertainties. To illustrate this, we consider a class of cosmological objects hosted in haloes of mass M as

Figure 4. Fractional difference between the large-scale bias measured by each percolation scheme and the fiducial scheme. Points with error bars represent the bias measured using the halo–halo correlation functions. Coloured regions show the bias measured using the halo–mass correlation functions. It is reassuring that the two bias measurements are in excellent agreement with each other.

defined using the ROCKSTAR percolation algorithm. We calculate the clustering amplitude of these haloes, and then use the $b(M)$ relation for the haloes in each of our four halo catalogues to infer the corresponding halo mass. Fig. 5 shows the bias in the inferred halo masses for each of our four halo catalogues. We see that the choice of percolation algorithm can bias the inferred halo masses by as much as 30 per cent for haloes of mass $M \approx 10^{13} h^{-1} \text{ M}_\odot$.

The translinear regime of the halo–correlation function has long been difficult to model, requiring ad hoc parametrizations that are

Figure 5. Bias in the inferred halo masses of our four catalogues by using large-scale clustering amplitude measurements to infer halo masses. In each case, the clustering amplitude is set by the clustering of the fiducial ROCKSTAR haloes. The observed amplitude is then mapped to a new halo mass using each of the alternative halo catalogues in turn. It is clear that the choice of percolation algorithm can severely impact the inferred halo mass.

then calibrated in numerical simulations (e.g. van den Bosch et al. 2013). Such an approach is likely sufficient within the context of modelling galaxy–galaxy clustering for the cosmological purposes, though we caution that verifying robustness of the modelling to simulations populated based on a different halo definition would be worthwhile. At the very least, inferences about how galaxy populate haloes will necessarily be impacted by the differences highlighted above. The sensitivity to halo definition will be even more problematic for cosmological studies that rely on the halo statistics directly, e.g. cluster abundance studies. We emphasize again that the work here is not meant to *calibrate* this effect, but rather to demonstrate its existence. Calibrations for observational studies must be specifically tailored to the observational methodologies employed.

Finally, our results bear some impact on the location of the splashback radius as found both in simulations and data. Specifically, the sharp steepening in the halo–mass correlation that occurs at the translinear regime has been identified with the splashback radius, the distance to the apocentre of dark matter substructures falling into a dark matter halo after their first pericentre pass (Adhikari, Dalal & Chamberlain 2014). This splashback radius has been proposed as a physical definition for the halo boundary (Diemer & Kravtsov 2014; More, Diemer & Kravtsov 2015). As demonstrated in this work, the steepening feature of the stacked halo profiles for haloes of a given mass can be moved by a change in the choice of halo percolation. This is not in itself problematic: in changing the population of haloes being stacked, the distribution of mass accretion rates of the resulting haloes will likely change, which in turn will move the average splashback radius (More et al. 2015). It does demonstrate, however, that calibration of the splashback radius via halo stacking is prone to systematics arising from the choice of halo percolation. This is particularly true for measurements of the splashback radius based on observationally selected cluster samples (e.g. More et al. 2016; Baxter et al. 2017; Busch & White 2017; Umetsu & Diemer 2017; Chang et al. 2018; Shin et al. 2018; Zürcher & More 2019; Contigiani, Hoekstra & Bahé 2019). Splashback measurements based on the analysis of particle orbits are, of course, free of such systematics (Diemer 2017; Diemer et al. 2017).

4.3 Halo–halo correlation function

We computed the halo autocorrelation functions for the same mass bins for which we calculated the halo–mass correlation function. As before, the autocorrelation functions exhibit an overall decrease in clustering amplitude for more aggressive halo exclusion criteria. As in the case for the halo–mass correlation function, we see the appearance of features in the translinear regime, though these features are less apparent than for the case of the halo–mass correlation function: even the translinear feature in the autocorrelation of our lowest mass bins has an amplitude of only ~ 10 per cent. Sample halo–halo correlation function plots are shown in Appendix A.

In a way analogous to the halo–mass correlation function, we can define the large-scale halo bias via

$$b^2(r|M) = \frac{\xi_{hh}(r|M)}{\xi_{mm}(r)}. \quad (5)$$

We fit a constant halo bias model over the same radial range as employed in our analysis of the halo–mass correlation function ($r \in [10, 80] h^{-1} \text{Mpc}$). The data points in the left-hand panel of Fig. 4 show the change in the clustering bias relative to our fiducial measurement for each of our four halo catalogues. We see that the change in the clustering bias amplitude is consistent across the halo–mass and halo–halo correlation function measurements.

5 SUMMARY AND CONCLUSIONS

We have shown that halo exclusion criteria impact halo statistics. Specifically, we modified the percolation of the ROCKSTAR halo-finding algorithm to generate four different halo catalogues with four different halo exclusion criteria, but identical mass definitions. We then measured the halo mass functions, halo–matter and halo–halo correlation functions, and large-scale clustering bias of the resulting catalogues. We compared these statistics to the halo statistics of the fiducial halo catalogue to quantify the level of uncertainty that the thus-far arbitrary choice of halo percolation scheme introduces in halo clustering statistics.

We find the following:

(i) The choice of halo exclusion criteria introduces a significant amount of systematic uncertainty on the halo mass function. The largest difference observed in this work was ≈ 6 per cent at a halo mass scale $M \sim 10^{13} h^{-1} M_{\odot}$. This value corresponds to the difference between the fiducial and hard-sphere percolation schemes. Consequently, it is of critical importance for future work on cluster abundances to implement identical exclusion criteria in both theory and simulations, particularly as the low-mass threshold for cluster detection gets progressively lower. Notably, *none* of the current choices of halo percolation can be implemented observationally.

(ii) The choice of percolation impacts the halo–matter correlations in two ways. At intermediate ($\approx 1 h^{-1} \text{Mpc}$) scales, large (≈ 30 per cent) relative differences in the halo–mass correlation function of the different halo catalogues arise. In addition, at large scales we see an offset in the large-scale clustering bias. We demonstrate that halo mass estimates based on the clustering amplitude of a set of cosmological objects can be biased by as much as 30 per cent due to the choice of percolation used when calibrating the bias–mass relation for haloes.

The differences in the predicted halo–mass correlation function of haloes will necessarily propagate into the predicted weak-lensing profiles of the resulting halo population, leading to further sources of systematic uncertainty impacting cluster abundance studies, a systematic which we intend to quantify in future work.

It is worth noting that while these differences are similar in spirit to differences associated with halo mass definition, a ‘right’ answer would naively appear to be more elusive. Within the context of halo mass definitions, the splashback radius (e.g. More et al. 2015) is now a leading contender for the ‘right’ radius at which to define the halo edge, naturally leading one to define halo mass as the mass contained within the splashback radius of a halo. By contrast, no similar leading candidate exists within the context of percolation schemes. We note, however, that the point–mass exclusion criterion adopted in this work is clearly the one closest in spirit to that of a strict spherical-overdensity mass definition. From an observational point of view, however, none of these percolation schemes can be easily implemented when analysing data. Thus, it is possible that halo percolation will need to be defined using data-motivated algorithms that may be theoretically less ‘clean’, but more directly applicable to the analysis of cosmological data sets. We will investigate in future work whether one can define objective quantitative criteria that might lead one to select one exclusion criteria over another within the context of specific science goals.

ACKNOWLEDGEMENTS

RG is supported by CONACyT scholarship 710106. ER is supported by DOE grant DE-SC0015975 and the Cottrell Scholar Award program. We would also like to thank Risa Wechsler, Joe de Rose, and Matt Becker for making the simulation used in this study available to us, Tom McClintock for technical help throughout the project, and Surhud More, Matt Becker, and Tom McClintock for comments on an early version of this manuscript.

REFERENCES

Adhikari S., Dalal N., Chamberlain R. T., 2014, *J. Cosmol. Astropart. Phys.*, 2014, 019

- Baxter E. et al., 2017, *ApJ*, 841, 18
 Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013, *ApJ*, 762, 109
 Bhattacharya S., Heitmann K., White M., Lukić Z., Wagner C., Habib S., 2011, *ApJ*, 732, 122
 Bocquet S., Saro A., Dolag K., Mohr J. J., 2016, *MNRAS*, 456, 2361
 Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80
 Busch P., White S. D. M., 2017, *MNRAS*, 470, 4767
 Chang C. et al., 2018, *ApJ*, 864, 83
 Contigiani O., Hoekstra H., Bahé Y. M., 2019, *MNRAS*, 485, 408
 Cooray A., Sheth R., 2002, *Phys. Rep.*, 372, 1
 DeRose J. et al., 2019, *ApJ*, 875, 69
 Diemer B., 2017, *ApJS*, 231, 5
 Diemer B., Kravtsov A. V., 2014, *ApJ*, 789, 1
 Diemer B., Mansfield P., Kravtsov A. V., More S., 2017, *ApJ*, 843, 140
 Knebe A. et al., 2013, *MNRAS*, 435, 1618
 McClintock T. et al., 2019, *ApJ*, 872, 53
 More S., Diemer B., Kravtsov A. V., 2015, *ApJ*, 810, 36
 More S. et al., 2016, *ApJ*, 825, 39
 Mountrichas G., Georgakakis A., Menzel M. L., Fanidakis N., Merloni A., Liu Z., Salvato M., Nandra K., 2016, *MNRAS*, 457, 4195
 Robertson B. E., 2010, *ApJ*, 716, L229
 Shin T. et al., 2019, *MNRAS*, 487, 2900
 Springel V., 2005, *MNRAS*, 364, 1105
 Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, *ApJ*, 688, 709
 Umetsu K., Diemer B., 2017, *ApJ*, 836, 231
 van den Bosch F. C., More S., Cacciato M., Mo H., Yang X., 2013, *MNRAS*, 430, 725
 Zürcher D., More S., 2019, *ApJ*, 874, 184

APPENDIX A: SAMPLE HALO–HALO CORRELATION FUNCTION PLOTS

Fig. A1 shows the halo–halo auto- and cross-correlation functions for the two mass bins used throughout the paper, as labelled. These plots are included here for completeness.

Figure A1. The halo–halo auto- and cross-correlation functions for the two mass bins used throughout the paper, as labelled. The top row shows the correlation functions for each of our four halo catalogues, whereas the bottom row shows the difference relative to the correlation function for the fiducial ROCKSTAR halo catalogue. The left-most column corresponds to our low-mass bin, the right-most column to the high-mass bin, and the central column shows the cross-correlation between the two.

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