

# Optimizing high-redshift galaxy surveys for environmental information

Tobias J. Looser,<sup>1,2,3★</sup> Simon J. Lilly,<sup>1</sup> Larry P. T. Sin,<sup>1</sup> Bruno M. B. Henriques,<sup>1</sup> Roberto Maiolino<sup>2,3,4</sup> and Michele Cirasuolo<sup>5</sup>

<sup>1</sup>*Department of Physics, ETH Zurich, Wolfgang-Pauli-Strasse 27, CH-8093 Zurich, Switzerland*

<sup>2</sup>*Cavendish Laboratory, University of Cambridge, 19 J. J. Thomson Ave., Cambridge CB3 0HE, UK*

<sup>3</sup>*Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

<sup>4</sup>*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

<sup>5</sup>*European Southern Observatory, Karl-Schwarzschild-Strasse 2, D-85748 Garching bei Muenchen, Germany*

Accepted 2021 April 13. Received 2021 March 15; in original form 2020 December 3

## ABSTRACT

We investigate the performance of group finding algorithms that reconstruct galaxy groups from the positional information of tracer galaxies that are observed in redshift surveys carried out with multiplexed spectrographs. We use mock light-cones produced by the L-Galaxies semi-analytic model of galaxy evolution in which the underlying reality is known. We particularly focus on the performance at high redshift, and how this is affected by choices of the mass of the tracer galaxies (largely equivalent to their co-moving number density) and the (assumed random) sampling rate of these tracers. We first however compare two different approaches to group finding as applied at low redshift, and conclude that these are broadly comparable. For simplicity we adopt just one of these, ‘Friends-of-Friends’ (FoF) as the basis for our study at high redshift. We introduce 12 science metrics that are designed to quantify the performance of the group-finder as relevant for a wide range of science investigations with a group catalogue. These metrics examine the quality of the recovered group catalogue, the median halo masses of different richness structures, the scatter in dark matter halo mass and how successful the group-finder classifies singletons, centrals, and satellites. We analyse how these metrics vary with the limiting stellar mass and random sampling rate of the tracer galaxies, allowing quantification of the various trade-offs between different possible survey designs. Finally, we look at the impact of these same design parameters on the relative ‘costs’ in observation time of the survey using as an example the potential MOONRISE survey using the MOONS instrument.

**Key words:** catalogues – surveys – galaxies: groups: general – galaxies: high-redshift – dark matter – large-scale structure of Universe.

## 1 INTRODUCTION

Large-scale galaxy redshift surveys carried out with efficient multi-object spectrographs (MOS) allow us to investigate the evolution of galaxies over cosmic time. While the earliest of these simply established the broad characteristics of the evolving galaxy population (e.g. Broadhurst, Ellis & Shanks 1988; Colless et al. 1990; Lilly et al. 1995) in terms of the luminosity (or mass) functions, in the last two decades, more extensive redshift surveys have enabled the study of the role of environment in driving this evolution, even for galaxies in quite typical environments, both at low redshift (e.g. Colless et al. 2001) and at much greater look-back times (e.g. Davis et al. 2003; Le Fèvre et al. 2005; Gunn et al. 2006; Lilly et al. 2009). Since the underlying and dominant dark matter structures are largely undetectable directly, these structures must be traced by other methods, such as the presence of extended hot gas (which is impractical on many scales of interest) or the distribution of the galaxies themselves, which is, in principle, accessible through highly multiplexed redshift surveys of large numbers of galaxies.

The identification of galaxy ‘groups’, i.e. a set of galaxies populating the same dark matter halo, from galaxy redshift catalogues enables the environments of galaxies to be characterized down to quite low halo masses. It is widely understood that the mass of the host halo of a galaxy, and whether the galaxy is the ‘central’ or a ‘satellite’ within that halo, are the dominant determinants of the evolution of the galaxy. The determination of these properties as accurately as possible is a pre-requisite for a physical understanding of the properties of galaxies. Even if one is interested in further ‘second-order’ effects, such as the growth history of the halo (e.g. Forbes, Krumholz & Burkert 2013), or the location of the halo within the filamentary structure of the cosmic web (e.g. Alonso, Eardley & Peacock 2015), or effects such as ‘galactic conformity’ (e.g. Kauffmann et al. 2013; Knobel et al. 2015; Sin, Lilly & Henriques 2017), one must understand as well as possible the immediate halo environment in order to remove the effects of this from any more detailed analysis.

Classifying the group environment of galaxies allows many physical investigations, including for instance the relationships between stellar mass, environment and star formation rate (SFR). Previous analyses of the SFR in galaxies in relation to their (group) environment have led to the isolation of distinct processes of ‘mass

\* E-mail: [tjl54@cam.ac.uk](mailto:tjl54@cam.ac.uk)

quenching’, and ‘environment quenching’ (Peng et al. 2010), the latter particularly affecting satellite galaxies. Differentiating centrals (which we operationally define to be the most massive member of the group, independent of its actual spatial position within the group) from satellites has established that the SFR of satellites is effected by their environment: When former central galaxies fall into larger dark matter haloes and become satellites, they are likely to quench their star formation and become quiescent, a process known as ‘satellite quenching’ (e.g. van den Bosch et al. 2008; Peng et al. 2012; Wetzel et al. 2012).

In planning an observational campaign to construct a large redshift catalogue of distant galaxies using an MOS, there are a number of considerations in the design of the survey. These can have a major impact on the amount of observing time that is required. These design decisions include, straightforwardly, the total number of galaxies to be observed and the amount of observing time required to secure a redshift for each, which will itself depend on parameters such as the stellar masses, SFR or other selection criteria. The projected number density of target galaxies on the sky, which will also be determined by the selection criteria of the galaxies and by the desired redshift range, will also determine how many individual ‘configurations’ of the MOS are required in each field. It may also affect how efficiently the multi-object capability of the instrument can be used, since this will generally degrade as the choice of targets within each field becomes more and more constrained. This may depend in detail on the design of the instrument. Another key parameter is the required sampling rate, i.e. what fraction of the available tracer targets are to be observed (including what fraction of those should have a successful redshift measurement). A high required sampling rate will generally produce a more constrained set-up and this will lower the efficiency. Finally, how the set of tracer galaxies is defined, and in particular their number density, will certainly affect, along with the sampling rate, the range of halo masses that is probed, and the accuracy with which the halo environment of each galaxies can be reconstructed. While it is obvious that these design decisions for the survey all have a potential impact on the resulting group catalogue, their quantitative impact is not trivially assessed.

The aim of this paper is therefore to quantitatively examine some of these issues. This investigation was carried out in the context of the design of the future MOONRISE (Maiolino et al. 2020) survey that will be undertaken using the MOONS MOS (Cirasuolo et al. 2020) that will soon be commissioned at the VLT.<sup>1</sup> Our approach however is to understand the main effects by idealizing the problem. MOONS-specific issues are as far as possible minimized or isolated, and the main results should be relevant for any similar survey. The point of this paper is not to design a particular survey with a particular instrument, but rather to explore the trade-offs in survey design that should be relevant for almost any multi-object spectroscopic survey. Of course, for any individual future survey, a more detailed survey strategy considering the real data, the real instrument, and the real fibre-positioning or other target selection software needs to be defined.

Semi-analytic models (SAMs) of the evolving galaxy population are now good enough that the basic properties of the galaxy population produced by the models, such as the mass or luminosity functions, well match the real galaxy population at different redshifts. This is especially true for those SAMs that use Markov chain Monte Carlo (MCMC) methods to explicitly tune the main parameters of

the SAM so as to best match the main distribution functions of the galaxy population. Likewise, the success of the standard model of cosmology means that the properties of the underlying population of dark matter haloes in such models are likely to closely approximate reality. The mock ‘light-cones’ from these models should therefore be quite a good match to the observed sky in the most relevant aspects. Such mock catalogues therefore enable us to apply exactly the same selection criteria as in a potential survey and then compare the properties of the reconstructed group catalogue to the underlying known ‘reality’. This enables us to assess the actual performance of the group-finder in the proposed survey using quantifiable metrics. At the very least, the differential performance of different survey designs can be quite reliably examined.

No group-finder based on galaxy positions can hope to perfectly assign galaxies to haloes because of the limited information available to the astronomer (essentially two projected spatial positions and a third radial velocity measurement). It should also be appreciated that even within the simulated Universe the association of galaxies to parent haloes is not without some ambiguity. The isolation of individual dark matter haloes is not perfectly defined as they continually accrete new material. Further difficulties can arise in unambiguously associating galaxies to be satellites within larger haloes, especially during the early phases of virialization. A gravitationally bound galaxy (subhalo) may still be found at some distance from the halo even after a first passage through it – the so-called ‘splash-back’ effect.

An important part of the current paper will be to construct quantitative metrics for assessing the performance of a group finding algorithm. These metrics go beyond the simple statistical concepts of purity and completeness and aim to be relevant for a wide range of different scientific applications of the output from the group-finder.

We will concentrate on the relative performance of the group-finder as the design of the survey is varied. We will focus especially on the two major design criteria for the selection of the tracer galaxies. For simplicity, we will assume that the primary selection criteria for the survey is the limiting stellar mass of the target galaxies,  $M_{\text{stellar}}$ . While previous deep spectroscopic surveys (e.g. CFRS, VVDS, zCOSMOS, etc.) used flux-limited samples, photometric redshifts are now good enough in well-studied fields like COSMOS that the stellar mass uncertainties are quite modest, allowing mass-selection of targets in future deep spectroscopic survey designs. As an example, the scatter in the observed  $H$ -band mass-light ratio for fixed rest-frame colour at high redshift  $z \approx 1.5$  is about 0.1 dex ( $1\sigma$ ). Additionally, the ability to vary the exposure times for different targets allow an efficient observation of mass-selected samples. For instance, the planned MOONS GTO survey MOONRISE will be selecting targets based on their photometrically estimated stellar masses as well as on their (AB)  $H$  magnitudes; see Maiolino et al. (2020).

The limiting stellar mass of the target galaxies translates directly into their volume number density in the Universe. Clearly, the volume number density of the tracers is one of the most important parameters in their ability to define structure in the Universe. To first order, any set of galaxies selected by some other criteria (e.g. the luminosity in some band) could be converted, via their volume number density, to an ‘equivalent’  $M_{\text{stellar}}$  limit. Of course, any particular selection of the tracers (including the straight stellar mass) will unavoidably introduce some biases into the galaxy content of the resulting identified groups. These biases in galaxy-content, arising from using other tracers than stellar mass, will not be considered in this paper.

<sup>1</sup>Very Large Telescope; <https://www.eso.org/public/teles-instr/paranal-observatory/vlt/?lang>.

The other major design parameter aside from the mass (or number density) of the tracers is the required sampling rate,  $s$ . This is defined as the ratio of the number of tracer galaxies with a usable redshift to the total number of such tracers. Incomplete sampling may arise from either not observing a galaxy at all, or in failing to secure a spectroscopic redshift – we do not distinguish between these in this investigation. Choice of lower  $s$  may reflect a simple desire to maximize the area of sky covered by the survey, or to maximize the multiplexing efficiency of the instrument due to technical constraints from the instrument design.

Non-random incompleteness, in the sense that the observations may fail to secure a redshift for some targets with particular properties, may be especially a problem at high redshifts. While this can be mitigated by adaptively varying the exposure time until a redshift is obtained, there will inevitably be some remaining incompleteness. Similarly, there will likely be a range of reliabilities of the redshifts, which may again introduce biases into the content of recovered structures.

Furthermore, the incompleteness may not be completely spatially random if local projected overdensities or particular geometries of nearby objects, lead to a reduction in the fraction of targets that can be observed. However, if the survey covers a large enough range in radial distance, the projected surface density of tracers is not strongly correlated with the volume number density associated with individual structures due to the high number of foreground and background objects (we examine this question later in this paper). Multiplex spectrographs like MOONS allow variable fibre-placement densities by using overlapping patrol fields of individual fibres, and the effects of target geometry can be reduced by multiple passes over each part of the sky. None the less, technical constraints might potentially prevent the proper sampling of clustered targets. Even though the effect is small, this issue must be addressed in any final survey design. In the current work, we will assume for simplicity that the incompleteness is completely random across the adopted set of tracers. The fraction of randomly observed objects is therefore described by a random sampling rate  $s$ .

There are several approaches in the literature to constructing group catalogues from galaxy redshift catalogues. These span a range of philosophies, depending on how much astrophysical information is assumed. On the one extreme are algorithms such as ‘Friends-of-Friends’ (FoF; e.g. Huchra & Geller 1982; Merchán & Zandivarez 2002, 2005; Eke et al. 2004; Gerke et al. 2005; Berlind et al. 2006; Knobel et al. 2009; Robotham et al. 2011) and similar approaches such as that based on the Voronoi–Delaunay tessellation (Marinoni et al. 2002). These rely only on the spatial location of galaxies in projected space and in velocity to associate galaxies together. In order to optimize the group finding, the FoF method uses three free parameters that control the galaxy–galaxy ‘linking-lengths’ perpendicular to, and along, the line of sight. These three free parameters can be optimized by applying the group-finder to simulations, for which we know the underlying dark matter haloes and the ‘true’ groups populating them, and optimizing the parameters to recover the ‘true’ catalogue as well as possible. Once structures are identified in this way, their dark matter masses must be estimated through the application of suitable algorithms, e.g. using the richness, the integrated stellar mass, estimates of the size or velocity dispersion, or some combination of these. These mass estimators may be calibrated against the mock catalogues, or by using other approaches such as abundance matching. FoF and related techniques have been extensively used at high redshift (e.g. Gerke et al. 2005; Knobel et al. 2009).

Particularly at low redshifts, e.g. for the SDSS,<sup>2</sup> some other, more refined, approaches have been developed. These assign membership of galaxies to groups in an iterative scheme that builds up the population of dark matter haloes around galaxies based on assumptions about the sizes of haloes (in projected space and velocity). These then assign membership of galaxies to individual groups based on a probabilistic approach. These approaches generally have estimates of the dark matter mass built in to the algorithm. This approach has been used by e.g. Yang et al. (2007) and Tinker et al. (2011). Again, such algorithms generally still have tunable parameters that should be optimized.

Of course, in principle, galaxies for which only photometric (but no spectroscopic) redshifts are available could also be included in any sample of galaxies. This was done for example in Kovac et al. (2009) to define the density field out to  $z = 1$  or in Wang et al. (2020) for group finding. For simplicity, we do not include these galaxies in the group catalogues used in this work.

In this work, we will first review in Section 2 the basic approaches to group finding used in this paper, review the most basic ideas of purity and completeness, and how these may be used to optimize the tunable parameters of the group finding algorithms. We also introduce the mock light-cones used throughout this paper. We describe a new implementation of an FoF group-finder (in this paper) and of what we call a ‘halo-based’ group-finder that has also been re-implemented in Sin et al. (in preparation).

Then in Section 3 we construct a set of 12 metrics to quantitatively assess the performance of group finding algorithms in a scientific context. In Section 4 we use these metrics to first examine the relative performance of the FoF and ‘halo-based’ algorithms when they are applied to an SDSS-like mock catalogue at low redshift. We will use for this a fixed stellar mass cut of  $10^{9.0}$  solar masses for the tracers, but examine a wide range of sampling rates,  $s$ , so as to better understand the strengths and short-comings of these two approaches as the sampling rate is reduced, concluding that their performance is comparable.

We will then apply, in Section 5, the FoF group-finder (alone) to three high-redshift ranges within the overall redshift range  $0.9 < z < 2.6$ , and examine how these metrics change as we vary both the stellar mass cut in the range  $10^{9.1-10.5}$  solar masses and the sampling rate  $s$ . This leads to a more quantitative understanding of the impact of the choices of  $M_{\text{stellar}}$  and  $s$  on the size and quality of the recovered group catalogues and on their usability for various scientific investigations.

In order to assess the recovered FoF catalogues, and how they change with  $M_{\text{stellar}}$ ,  $s$ , and  $z$ , we present numerous quality metrics and introduce various science metrics designed to capture a wide range of environmental information about the recovered galaxy group catalogues.

As a realistic science example, we also explore the degree to which imperfections in the recovered group catalogue perturb a simple science measurement: the quenched fraction of central and satellite galaxies as a function of their host halo mass. We explore how the metrics can be used to construct simple corrections to the observed quenched fractions to best recover the correct values.

Finally, we present an analysis of the ‘costs’ of possible survey designs in terms of observing time. This highlights the trade-offs between the selections in stellar mass  $M_{\text{stellar}}$  and the choice of sampling rate  $s$  in terms of the total telescope observing time and the efficiency with which the multiplexing capability of the instrument is utilized. As part of this, we present a general formalism for

<sup>2</sup>Sloan Digital Sky Survey; <https://www.sdss.org/>.



consideration of the survey costs, in which the instrument-specific issues are all concentrated in just one of the four terms, enabling a relatively straightforward consideration of these trade-offs. We then summarize the conclusions of this paper in Section 7. Throughout this paper, we assume a  $\Lambda$ CDM cosmology with the following parameters:  $H_0 = 69.6 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ,  $\Omega_M = 0.286$ , and  $\Omega_\Lambda = 0.714$ .

## 2 GALAXY GROUP FINDING


Galaxy groups are the set of galaxies populating the same dark matter halo. However, dark matter haloes are not directly observable and so, in practical terms, galaxy groups must be recovered from catalogues of the observable positions (RA, Dec., and redshift) and possibly other quantities (e.g. stellar mass) of the galaxies observed in any large-scale galaxy redshift survey. Fortunately, mock catalogues from ‘light-cones’ based on realistic models of galaxy formation and evolution are now available in which the group membership of all galaxies are, at least in principle, known from the underlying model. These mock catalogues can therefore provide an underlying ‘truth’ against which a recovered or reconstructed group catalogue may be compared.

In this section, we first briefly describe the L-Galaxies Munich Galaxy Model that was used as the (simulated) mock reality throughout this work. This mock reality is used first to optimize the parameters of different group finding algorithms and also then to quantitatively compare their performance, both relative to each other and also, for a given group-finder, as the tracer selection (in terms of stellar mass) and sampling rate are varied.

We then briefly review the basic concepts of how real and recovered groups can be matched and the quantitative definitions of completeness and purity of the recovered group catalogue. We also briefly review the operation of the two representative group-finders examined in this work, which, although described in previous papers are here slightly modified: The first is the Sin et al. re-implementation (Sin et al., in preparation) of a group-finder based on the work of the Tinker et al. (2011). The second is a multi-run implementation of an FoF group-finder (e.g. Huchra & Geller 1982; Gerke et al. 2005; Knobel et al. 2009). We again briefly show how the parameters of these algorithms can be optimized (generally by balancing purity and completeness). Finally, we construct an empirical mass estimator for the FoF algorithm.

### 2.1 The L-Galaxies SAM

The L-Galaxies SAM<sup>3</sup> (Henriques et al. 2015) of galaxy formation and evolution is used throughout this paper as the ‘mock’ reality. L-Galaxies is built on to the Millennium (Springel et al. 2005) and Millennium II (Boylan-Kolchin, Ma & Quataert 2008) dark matter simulations, which are  $\Lambda$ CDM cosmological  $N$ -body simulations, using  $10^{10}$  particles, predicting the evolution of dark matter haloes over cosmic time. Such simulations retain the merger history of individual dark matter haloes and the evolution of individual ‘sub-haloes’ within the larger haloes. The L-Galaxies SAM then models the evolution of the baryonic matter component within these dark matter haloes, based on prescriptions for the physical processes that affect galaxy formation and evolution such as gas cooling, star formation, supernova feedback, the formation and growth of black



**Figure 1.** The distribution of dark matter halo mass as function of true group richness; a basic property of the L-Galaxies SAM illustrative true group catalogue for  $\log_{10}(M_{\text{stellar}}/M_{\odot}) > 9.5$  within the redshift range  $1.2 < z < 1.7$ . An artificial random scatter is added to the richness values for readability.

holes and the feedback from these, as well as galaxy interactions and mergers.

In this paper, we use the mock ‘pencil-beam’ ‘light-cones’ produced by the L-Galaxies SAM using M05 stellar populations. These mimic the distant (model) universe as observed by us. These light-cones include information about the mass and location of dark matter haloes with redshift, and of the simulated galaxies within them. In what follows, these simulated galaxy groups are regarded to represent ‘reality’, and we will refer to them as the ‘true galaxy groups’. Each true galaxy group is defined to be the set of galaxies populating the same dark matter halo in the simulation. The ensemble set of true galaxy groups forms a true galaxy group catalogue. The Richness  $N$  is defined as the number of member tracer galaxies. Some galaxies may be in a group of just one galaxy, and these will be referred to as singletons.

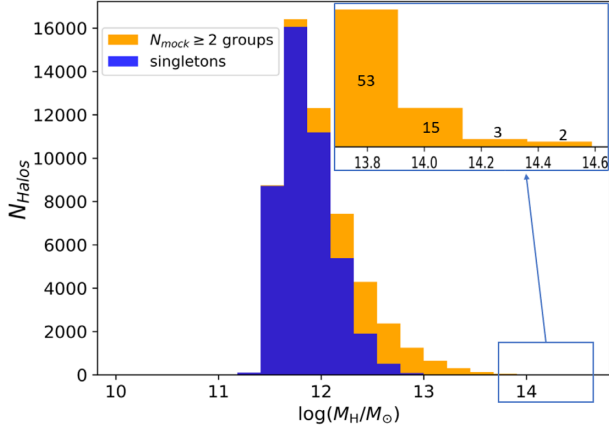
For orientation, it is worthwhile to briefly review the true group catalogue from the L-Galaxies model. As an example, we consider the catalogue that is obtained using tracers of stellar mass  $\log_{10}(M_{\text{stellar}}/M_{\odot}) > 9.5$  within a light-cone of  $2 \times 2 \text{ deg}^2$  on the sky and in the redshift range  $1.2 < z < 1.7$ . This volume contains 75 838 galaxies above the stellar mass limit. Two basic properties, the halo mass  $M_H$  versus richness  $N$  distribution and the number density of groups as function of halo mass, are shown in Figs 1 and 2, respectively.

The effects of varying the stellar mass selection limit  $M_{\text{stellar}}$  on the characteristic properties of the group catalogue are shown in Fig. 3. With a more restrictive (higher) mass cut, the multiplicity  $m$ , which is the average number of members per group, decreases, as does the fraction of galaxies that are satellites  $F_{\text{satellites}}$ . The fraction of galaxies that are centrals (including the singletons)  $F_{\text{centrals}}$  and the fraction of singleton galaxies themselves  $F_{\text{singletons}}$ , on the other hand, both increase. The choice of limiting stellar mass selections systematically shifts the basic properties of the resulting galaxy group catalogue.


In this analysis, we use several of the 24 available L-Galaxies pencil-beam  $2 \times 2 \text{ deg}^2$  light-cones (always using the M05 stellar models). We use one light-cone to train the algorithms by optimizing the free parameters (e.g.  $b$ ,  $R$  for the FoF group-finding algorithm, see Section 2.3 and to fit the polynomial coefficients in our halo mass estimator, see Section 2.5). We then apply the algorithms to a different light-cone in order to obtain the results in the Sections 3–6.

<sup>3</sup><https://wwwmpa.mpa-garching.mpg.de/galform/virgo/millennium/>.





**Figure 2.** The number of dark matter haloes populated by singletons and richer structures as a function of dark matter halo mass; a basic property of the L-Galaxies SAM illustrative true group catalogue for  $\log_{10}(M_{\text{stellar}}/M_{\odot}) > 9.5$  within the redshift range  $1.2 < z < 1.7$ .



**Figure 3.** The basic characteristics of the true group catalogue for different limiting mass cuts  $M_{\text{stellar}}$  in the redshift regime  $1.2 < z < 1.7$ . The multiplicity  $m$  is the average number of members per group. Note that the fraction of centrals  $F_{\text{centrales}}$  includes the singletons. Most centrals are in fact singletons.

The single individual light-cones contain large enough numbers of galaxies and groups to allow statistically meaningful conclusions. As a detail, the light-cones have sharp edges and hence cut through some galaxy groups. The galaxies in these truncated groups have not been excluded in this analysis. This might slightly bias the sample of galaxy groups, but the fraction of such groups is extremely small and hence the biases introduced by this effect are negligible.

## 2.2 The fidelity of recovered galaxy groups

In this section, we review the most basic statistical quantities that may be used to assess the quality of the recovered group catalogue produced by a given group-finder applied to a given set of tracer galaxies.

### 2.2.1 Matching of real and recovered groups

The performance of any group-finder can be assessed by comparing the recovered groups (from the group-finder) and those in the true group catalogue (from the mock simulation). Although we will later

define more science-oriented metrics, these most basic quantities are used to optimize the parameters of the group-finder, and also serve as the basis for some of the later metrics.

Following Knobel et al. (2009) and Gerke et al. (2005) the matching procedure is conceptually illustrated in Fig. 4. Each point represents a galaxy and we compare the two orderings (the true group catalogue and the reconstructed group catalogue) of this identical set of points. The precise definitions used are as follows:

(i) Match: a group  $i$  is matched to another group  $j$  if group  $j$  contains a fraction  $f > f_M$  of the members of group  $i$ . For this association to be unique, it must hold that  $f_M \geq 0.5$ . Throughout this paper, we use this minimum  $f_M = 0.5$ , as in Knobel et al. (2009) and Gerke et al. (2005). Hence, for a match, more than 50 per cent of the members of group  $i$  have to be found in group  $j$ . For example, if group  $i$  is a  $N = 4$  member group, group  $j$  must contain at least three of those four galaxies as members, in order to be matched. Similarly, both members of an  $N = 2$  group must be present for a match.

(ii) One-way match: the case where group  $i$  is associated with group  $j$ , but group  $j$  is not associated with group  $i$  (illustrated by a one-way arrow from group  $i$  to group  $j$ ). One-way matches can occur from real to recovered groups or vice versa.

(iii) Two-way match: the case where group  $i$  is associated with group  $j$  and group  $j$  is also associated back with group  $i$  (illustrated by a double-arrow).

While each group can by definition only have a single, unique associated group (i.e. an arrow pointing away), it might well happen that a certain group is the matched group for two or more other groups (i.e. have two or more arrows pointing inwards to it). We therefore use the following further terminology:

(i) Overmerged group: if more than one real group is matched to the same reconstructed group.

(ii) Fragmented group: if more than one reconstructed group is matched to the same real group.

(iii) Spurious group: a reconstructed group that is not matched to any real group.

(iv) Undetected group: a real group that has no matched reconstructed group.

(v) Singleton: a galaxy that is not associated with any group in the catalogue.


An important concept is that of ‘two-way matching’ between real and recovered groups. For these groups, at least 50 per cent of the members of a single real group are found within a single recovered group, and at least 50 per cent of the members of that single recovered group are also members of the original real group. Two-way matches are by construction unique, in the sense that each real or recovered group can have at most a single two-way match. They are the most reliably reconstructed groups.

### 2.2.2 Completeness and purity

In order to quantify the basic fidelity of the group-finder, we again follow Knobel et al. (2009) in defining the completeness  $c$  and purity  $p$  for assessing the overall performance of the group-finder. The one-way completeness is defined as

$$c_1(N) = \frac{\mathcal{A}[N_{\text{gr}}^{\text{mock}}(N) \rightarrow N_{\text{gr}}^{\text{rec}}(> N/2)]}{N_{\text{gr}}^{\text{mock}}(N)}, \quad (1)$$

where  $\mathcal{A}[N_{\text{gr}}^{\text{mock}}(N) \rightarrow N_{\text{gr}}^{\text{rec}}(> N/2)]$  counts the number of true groups  $N_{\text{gr}}^{\text{mock}}(N)$  with a population of  $N$  members, which can be



**Figure 4.** Diagrammatic illustration (Knobel et al. 2009) of the matching between the ‘real’ groups (from the simulation) and the reconstructed groups. See the text for details.

matched to a recovered group, by either a one- or two-way match.  $c_1(N)$  therefore evaluates the fraction of true  $N$ -member groups that are one- or two-way matched. In the same way,

$$c_2(N) = \frac{\mathcal{A}[N_{\text{gr}}^{\text{mock}}(N) \leftrightarrow N_{\text{gr}}^{\text{rec}}(> N/2)]}{N_{\text{gr}}^{\text{mock}}(N)}, \quad (2)$$

evaluates the fraction of groups that are matched using the more restrictive two-way matches only. Analogously, we define the one-way purity,

$$p_1(N) = \frac{\mathcal{A}[N_{\text{gr}}^{\text{rec}}(N) \rightarrow N_{\text{gr}}^{\text{mock}}(> N/2)]}{N_{\text{gr}}^{\text{rec}}(N)}, \quad (3)$$

and the two-way purity,

$$p_2(N) = \frac{\mathcal{A}[N_{\text{gr}}^{\text{rec}}(N) \leftrightarrow N_{\text{gr}}^{\text{mock}}(> N/2)]}{N_{\text{gr}}^{\text{rec}}(N)}, \quad (4)$$

of the recovered group catalogue. The purity evaluates the fraction of recovered groups with a population of  $N$  members that have a match to a true group.

The quality of a group-finder can then be assessed by these two performance metrics, completeness and purity, which respectively measure the fraction of true groups that are recovered and the fraction of recovered groups that have an associated true group. Ideally, we would want the group-finder to exhibit both high completeness and purity, but these are not independent and, generally, one expects a trade-off between the two.

If the group-finder is set up to easily find structures it will generally exhibit high completeness at the cost of low purity, due to overmerging and the formation of spurious groups. Likewise, if the group-finder is set up so that there are restrictions in identifying structures, then the purity will be high at the cost of completeness, as true groups will be in danger of being fragmented and there will be undetected groups.

Depending on the scientific goals, one might wish to value purity over completeness, or vice versa. However, usually a group-finder should represent a balance between completeness and purity. Hence, again following Knobel et al. (2009), we quantify the ‘quality’ of the

recovered group catalogue by a parameter

$$g_1 = \sqrt{(1 - c_1)^2 + (1 - p_1)^2}, \quad (5)$$

which values completeness and purity equally. The parameter  $g_1$  gauges the deviation from a ‘perfect’ group catalogue that would have  $(c_1, p_1) = (1, 1)$ . Note that  $g_1$  should therefore be as small as possible.

### 2.3 FoF group-finder

Our implementation of the FoF group-finder is closely related to the FoF algorithm used in Knobel et al. (2009), which itself was based on the FoF implementation used for the DEEP2 group catalogue by Gerke et al. Gerke et al. (2005). The group finding approach of the FoF method is to ‘link’ nearby galaxies, and subsequently form groups of all those galaxies that are linked together, either directly or indirectly via other galaxies. The linking of galaxies depends on their transverse (spatial) and radial (velocity) separations. The maximum allowed transverse separation, in order to be linked, is governed by the perpendicular linking length parameter  $l_{\perp}$ , and the maximum allowed radial velocity separation is governed by the parallel linking length  $l_{\parallel}$ . These two parameters  $l_{\perp}$  and  $l_{\parallel}$  are controlled by the three free parameters of the FoF group-finder  $b$ ,  $R$ , and  $L_{\text{max}}$ , as follows:

(i) The basic linking length parameter  $b$  is the main parameter controlling the perpendicular linking length  $l_{\perp}$ .  $b$  links  $l_{\perp}$  to the mean observed spatial galaxy number density  $\bar{n}$ , which may (slowly) vary as a function of redshift if required. The purpose of relating  $l_{\perp}$  to  $\bar{n}$  is to allow an increase of the linking length as the number density of tracers decreases.

(ii) As  $\bar{n}$  decreases, however,  $l_{\perp}$  may increase so much as to be larger than the physical scale of any structures of interest. Hence,  $L_{\text{max}}$  is introduced to limit the maximum linking length. The perpendicular linking length is therefore given by

$$l_{\perp} = \min \left[ \frac{b}{\bar{n}^{1/3}}, L_{\text{max}}(1+z) \right]. \quad (6)$$

(iii) The transverse and radial linking lengths are related by  $R$ . The need for this arises due to the peculiar velocities of the galaxies. These

local velocities shift the measured redshifts, and hence considerably increase the apparent comoving radial distances between galaxies based on redshift alone, the famous ‘Fingers of God’ effect.  $R$  is therefore introduced in order to allow a larger radial linking tolerance.  $R$  is defined as the ratio between  $l_{\perp}$  and  $l_{\parallel}$ ,

$$l_{\parallel} = R l_{\perp}. \quad (7)$$

$l_{\perp}$  and  $l_{\parallel}$  then govern the linking of galaxies in the following way: Two galaxies  $i$  and  $j$  with apparent comoving radial distances from the observer  $d_i$  and  $d_j$ , as observed from their redshifts, and an intermediate angle  $\theta_{ij}$  are considered linked if the difference of their apparent comoving radial distances fulfills the inequality

$$|d_i - d_j| \leq \frac{l_{\parallel,i} + l_{\parallel,j}}{2}, \quad (8)$$

and, at the same time, their intermediate angle satisfies

$$\theta_{ij} \leq \frac{1}{2} \left( \frac{l_{\perp,i}}{d_i} + \frac{l_{\perp,j}}{d_j} \right). \quad (9)$$

These two inequalities are designed to impose adequate limits, dependent on the observed redshifts, on the maximal transverse and radial separations of two linked galaxies.

The choice of the free parameters  $b$ ,  $R$ , and  $L_{\max}$  in the FoF group finding algorithm involves a trade off between completeness and purity: A large linking length improves the completeness at the cost of purity, and vice versa.


In order to obtain an optimal balance between completeness and purity, we define the optimal free parameters as those one which minimize  $g_1$ , defined above in equation (5), in the resulting recovered group catalogue.

Within a large enough volume of the Universe, there will be a wide range of dark matter halo masses and richnesses, from singletons up to groups of fifty or more members. Rich groups and small groups may have quite different densities in projected transverse and radial velocity space. As a result, the optimal parameters for identifying rich groups may well differ from those for smaller groups.


Again following Knobel et al. (2009), we therefore apply a ‘multi-run’ scheme in the FoF group finding algorithm. The optimal linking parameters are determined across the range of richnesses of the real groups. The full recovered group catalogue is then constructed by first applying the group-finder with parameters optimized for some high bin of (true) richness. All recovered groups within this same richness bin are then selected, and all their member galaxies are removed from the tracer population. The next run uses the optimal linking parameters for the next lower richness bin, more groups are identified and the members are again removed. This is repeated all the way down to the final two-member bin. The galaxies that then still remain are classified as singletons.<sup>4</sup>

For illustration, Fig. 5 shows the parameter optimization in the  $N = 6-7$  richness bin in the  $1.2 < z < 1.7$  redshift range for a stellar mass cut of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.5$  and a sampling rate of  $s = 0.8$ . As the spatial galaxy density  $\bar{n}$  is reasonably high for these survey parameters, we can omit  $L_{\max}$  and optimize for  $b$  and  $R$  only. To optimize the FoF group finding method, we apply a  $(b, R)$  grid-search for every combination of limiting stellar mass and random sampling rate ( $M_{\text{stellar}}, s$ ) in each of the three redshift ranges.

<sup>4</sup>As an aside, we investigated a variation of this multi-run method, in which the linking parameters for each richness bin are re-optimized on the remaining galaxies at each step. It was found that this approach performed only slightly better on the test samples and slightly worse on average on the validation samples, and was therefore discarded.



**Figure 5.** FoF  $(b, R)$ -parameter optimization with respect to  $g_1$  in the  $1.2 < z < 1.7$  redshift range for a stellar mass cut of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.5$  and a sampling rate of  $s = 0.8$  in the  $N = 6-7$  richness bin. The deviation of the optimal recovered FoF group catalogue, with corresponding optimized parameters  $b = 0.06$  and  $R = 40$ , from the true group catalogue is  $g_1 = 0.198$ , as illustrated by the red dot. The expected trade-off between completeness and purity as the FoF parameters are adjusted is clearly seen.



**Figure 6.** FoF one- and two-way completeness and purity, and  $g_1$  and  $g_0$  as function of group richness  $N$  in the  $1.2 < z < 1.7$  redshift range and for a stellar mass cut of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.5$  and a sampling rate of  $s = 0.8$ . It should be noted that the closer  $p_1$ ,  $p_2$ ,  $c_1$ , and  $c_2$  are to 1, and thus  $g_0$  and  $g_1$  to 0, the better the quality of the recovered group catalogue.

We observe that while the main parameter  $b$  is rather stable over all richness’s,  $R$  alters substantially, presumably because of the changing effects of peculiar velocities. However,  $g_1$  is more responsive to changing  $b$  than it is to changing  $R$ . When re-running the multi-run procedure multiple times, for the same mass-selection and sampling rates, but with a different random galaxy sampling each time, we observe the same: Over all richness ranges, the  $b$  parameter is stable, while  $R$  varies a good deal. Despite this,  $g_1$  is largely stable for an extended range of  $R$ , meaning that the precise value of  $R$  is not decisive for the overall performance of the multi-run group-finder.

In Fig. 6, the quality of the recovered group catalogue, in terms of one- and two-way completeness and purity, as function of richness is shown. Over all richness bins, we observe a total one-way completeness of  $c_1 = 0.82$  and a one-way purity of  $p_1 = 0.79$ ,



with an overall quality of  $g_1 = 0.28$  and  $g_0 = 0.37$ , where

$$g_0 = \sqrt{(1 - c_2)^2 + (1 - p_2)^2} \quad (10)$$

assesses the relative frequency of failures in finding two-way matches for true and recovered groups.

## 2.4 ‘Halo-based’ galaxy group-finder

The group finding algorithm that we call the ‘halo-based’ method is heavily based on the methodology developed by Tinker et al. (2011) as implemented in Sin et.al. (in preparation) and we refer to that paper for details of the implementation.

The ‘halo-based’ methodology works as follows:

(i) Initially, every galaxy is defined to be the sole member of its group and a dark matter halo mass is assigned to each of these ‘groups’ via its stellar mass, using a calibration based on subhalo abundance matching (or, conceptually, any other method).

(ii) Subsequently, the local matter density  $P_M$  (defined below) is evaluated and, starting from the highest  $P_M$ , galaxies with  $P_M > B$  are put together as members of a single group and masked from further group assignment (for this iteration). The parameter  $B$  is tunable.

(iii) Dark matter halo masses are then re-assigned to these new groups again using the total stellar masses of the recovered groups.

(iv) The steps (ii) and (iii) are repeated iteratively until the group membership converges, i.e. remains unchanged between two iterations.

The local matter density  $P_M$  is defined as

$$P_M = \frac{c}{H_0} \Sigma(R_{\text{proj.}}, \alpha r_s) p(\Delta z, \beta \sigma_v), \quad (11)$$


where  $\Sigma$  is the projected NFW profile of a halo with scale radius  $r_s$  evaluated at projected radius  $R_{\text{proj.}}$ , and  $p$  is the normalized Gaussian function of a halo with velocity distribution  $\sigma_v$  evaluated at redshift-space separation  $\Delta z$ , relative to the respective centres of that halo. The formulas for these two terms are given in Tinker et al. (2011), with the difference that in the Sin et.al. (in preparation) implementation, two additional tunable parameters  $\alpha$  and  $\beta$  are introduced, allowing further degrees of freedom in optimizing the group-finder.

As a detail, the ‘halo-based’ method used in this work is optimized for the balanced pairwise accuracy metric  $\Pi$  introduced in Sin et.al. (in preparation),

$$\Pi = \sqrt{f_{\text{correct}}^2 - (f_{\text{frag.}}^2 + f_{\text{merge}}^2)}, \quad (12)$$

where  $f_{\text{correct}}$ ,  $f_{\text{frag.}}$ , and  $f_{\text{merge}}$  are, respectively, the fraction of pairs that are correctly classified, fragmented (i.e. same-halo pairs that are misclassified as different-halo), and merged (i.e. different-halo pairs that are misclassified as same-halo), when one considers all pairs that are separated by less than 1 Mpc in projected separation, and 500 km s<sup>-1</sup> in apparent velocity difference.

Relative to the FoF algorithm, the ‘halo-based’ approach is more refined, in the sense that it uses more astrophysical and cosmological ‘knowledge’. Of course, if that knowledge is correct, then this will likely produce a more accurate result. The trade-off is the risk of inaccuracies if that knowledge is incomplete or not correct. One of the goals of the current work was to investigate in particular how the relative performance of the refined ‘halo-based’ approach compared with the more basic FoF approach changes as the information on the tracer population is degraded, especially through a reduction in the tracer sampling rate  $s$ .



**Figure 7.** The correlation of halo mass  $M_H$  (plotted in log stellar mass) to  $\sum M_{\text{stellar}}$  and  $\sigma_v$  in the  $1.2 < z < 1.7$  redshift regime, for  $s = 0.8$  and  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.5$ . The red line shows an  $M_H$ -estimator based on  $\sum M_{\text{stellar}}$  only.

## 2.5 $M_H$ -estimator

In this section, we introduce a dark matter halo mass  $M_H$ -estimator for the recovered galaxy groups. An estimate of the dark matter mass is required for many science applications and the scatter in recovered halo mass will be one of the science-based metrics that we introduce in Section 3. Whereas the ‘halo-based’ approach yields estimated masses directly, by construction, the masses of FoF groups are estimated separately, *post facto*. One could of course adopt abundance matching methods, as in the ‘halo-based’ approach, but one can also try to calibrate directly from the mock universe. We here construct an ‘empirical’ mass estimator for FoF groups based on the total stellar mass of the members of the galaxy groups  $\sum M_{\text{stellar}}$ , the rms radial velocity dispersion of the groups  $\sigma_v$ , and the projected rms sizes of the groups  $\sigma_r$ . This is constructed by examining the true groups in the mock ‘reality’ introduced earlier in this paper.

For each recovered galaxy group, we define a  $\sigma_v$  as


$$\sigma_v = \sqrt{\frac{1}{N_{\text{gal}} - 1} \sum_{i=0}^{N_{\text{gal}}} \left( c \frac{z_i - \bar{z}}{1 + \bar{z}} \right)^2}, \quad (13)$$

where  $N_{\text{gal}}$  is the number of galaxies within the group. The factor  $(N_{\text{gal}} - 1)^{-1}$  compensates for the fact that 1 degree of freedom is used in order to estimate  $\bar{z}$ . Analogously, we define a  $\sigma_r$  as

$$\sigma_r = d_{\text{comdist}}(\bar{z}) \sqrt{\frac{1}{N_{\text{gal}} - 1} \sum_{i=0}^{N_{\text{gal}}} \left( (\theta_i - \bar{\theta})^2 + \cos^2 \bar{\theta} (\phi_i - \bar{\phi})^2 \right)}, \quad (14)$$

where  $d_{\text{comdist}}(\bar{z})$  is the comoving distance corresponding to the average redshift of the group  $\bar{z}$ .

Not surprisingly, the dominant determinant of  $M_H$  is  $\sum M_{\text{stellar}} \cdot \sigma_v$  can be used to improve the  $M_H$ -estimator further: In Fig. 7, we show that galaxies with high  $\sigma_v$  tend to lie above the red line, representing a halo mass estimate based on the total stellar mass only. For  $\sigma_r$ , there is little empirical evidence that it improves the halo mass estimate further over  $\sigma_v$  or  $\sum M_{\text{stellar}}$ . We therefore implemented an  $M_H$ -estimator calculating  $M_{H, \text{est}}$  as a polynomial function of  $\sum M_{\text{stellar}}$  and  $\sigma_v$  only. Because it returns empirically the best results, we adopted



**Figure 8.**  $M_H$ -estimator for singletons in the  $1.2 < z < 1.7$  redshift regime, for  $s = 0.8$  and  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.5$ .

a third-degree polynomial function,

$$\log_{10}(M_{H,\text{est}}/M_{\odot}) = \sum_{\alpha+\beta \leq 3} c_{\alpha\beta} \times \left( \log_{10} \left( \sum M_{\text{stellar}}/M_{\odot} \right) \right)^{\alpha} \left( \log_{10}((\sigma_v)/(\text{km/s})) \right)^{\beta}, \quad (15)$$

where  $c_{\alpha\beta}$  are the coefficients of the polynomial function. Note that the coefficients  $c_{\alpha\beta}$  are dependent on the considered redshift regime, the limiting stellar mass and the sampling rate, and hence have to be calibrated for each individual  $(z, M_{\text{stellar}}, s)$ -bin separately. Further, the galaxy-based and the group based coefficients are slightly different; while in the galaxy-based analysis, we consider each galaxy in each halo individually (hence the estimator is weighted by the number of galaxies in each halo), in the group-based analysis, we consider haloes as a whole.<sup>5</sup>

For singletons, the rms velocity dispersion and the rms size are not defined, hence we implement for them a third-degree polynomial  $M_H$ -estimator based on  $\sum M_{\text{stellar}}$ , i.e.  $M_{\text{stellar}}$  only, as illustrated in Fig. 8.

The halo mass estimator is a tool that will be used below to assign masses to recovered groups in order to calculate the scatter in the real and recovered  $M_H$ . It is important to appreciate that this scatter will therefore include both the intrinsic scatter in the  $M_H(\sum M_{\text{stellar}}, \sigma_v)$  relation in the real (mock) universe, plus the additional effects arising from any infidelities of the group catalogue.

### 3 SCIENCE METRICS FOR PERFORMANCE ESTIMATION

The purity and completeness reviewed above (see Section 2.2.2) are operationally well defined and, as explained there, are ideal for optimizing the values of the adjustable parameters of any group finding algorithm. However, it is not immediately clear how the completeness and purity actually map over to a quantitative

<sup>5</sup>To give an example, in the  $1.2 < z < 1.7$  redshift regime with  $\log_{10}(M_{\text{stellar}}/M_{\odot}) > 9.5$  and  $s = 0.8$ , the coefficients of the galaxy-based estimator are  $c_{00} = 9.61899242e + 00$ ,  $c_{10} = 2.97574209e - 15$ ,  $c_{20} = 8.52193554e - 15$ ,  $c_{30} = 2.27243612e - 03$ ,  $c_{01} = 5.46838695e - 23$ ,  $c_{02} = 2.77947535e - 25$ ,  $c_{03} = 2.38635562e - 02$ ,  $c_{11} = 4.39220440e - 23$ ,  $c_{21} = 1.43856458e - 23$ ,  $c_{12} = 6.03824430e - 30$ .

degradation of the scientifically useful information that is in principle contained within a group catalogue.

Therefore, in this section, we will construct several ‘science metrics’. Each of these is designed to assess quantitatively the performance of a given galaxy group-finder (with optimized parameters) from the perspective of an end-user scientific investigator interested in one or other scientific questions.

We stress that the purpose of constructing these different metrics is not to try to identify one or two metrics that are somehow ‘the best’, but rather to explore quite a large number of different metrics that are relevant for the wide and diverse range of scientific investigations that can be based on a (recovered) group catalogue from a given redshift survey. Different end-users may be interested in different combinations of these metrics, according to their particular scientific goals.


For definiteness of illustration, we will also present in this section the quantitative values of each science metric for the particular case of a recovered FoF group catalogue in the  $1.2 < z < 1.7$  redshift range, obtained with a tracer mass cut of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) > 9.5$  and a random sampling rate of  $s = 0.8$ . Later in this paper, we will look at how these metrics differ between FoF and ‘halo-based’ group-finders when applied at low redshift, as we vary the sampling rate (alone). We will then examine how the FoF performance varies at high redshifts as the limiting stellar mass cut and sampling rate are varied, and how this changes with redshift.

#### 3.1 Masses of the recovered groups on a galaxy basis

One defining property of a galaxy group is the mass of the associated dark matter halo. The mass is a fundamental quantity of any gravitationally bound structure, and halo mass may well also be one of the most important drivers of galaxy evolution.

We first distinguish between metrics that are constructed from the properties of the parent haloes that are assigned to each individual galaxy, and those that refer to the groups themselves. Imperfect fidelity of group reconstruction may affect these two in quite different ways. The former may be most relevant, for example, in studies of galaxy evolution because one needs to know the underlying total parent halo mass of each individual galaxy. The latter will be more relevant for studies of the X-ray emitting gas and the like. In this subsection, we will construct metrics describing the performance of the group-finder in accurately describing the parent halo of each individual galaxy. We refer to these as ‘galaxy-basis’ metrics, as opposed to the ‘group-basis metrics’ considered later.

Fig. 9 shows the distribution of the true parent halo mass for every galaxy in the sample, as a function of the richness of the recovered group that contains that galaxy. In other words, it looks at the range of real parent halo masses for the galaxies that are in the reconstructed groups of a particular richness. Each dot therefore represents a single galaxy. At low richnesses the (integer) values of  $N_{\text{rec}}$  have been artificially horizontally broadened for clarity. It should be noted that the L-Galaxies simulation will already have an intrinsic scatter between the dark matter halo mass  $M_H$  and the galaxy group richness  $N_{\text{true}}$ . This will be further increased in Fig. 9 by the imperfect assignment of galaxies to groups bringing into the (recovered) group interlopers who have quite different (true) halo masses. The colour-coding in this figure represents the ratio between the assigned (recovered) richness of the group that contains the galaxy in question,  $N_{\text{rec}}$ , and the true richness  $N_{\text{true}}$  of the group that in reality contains that galaxy. Comparison of these two therefore reflects the amount of fragmentation or merging of groups in the recovered group catalogue. The galaxies shown in purple are



**Figure 9.** True parent halo masses of galaxies in recovered groups as a function of the richness of the recovered group, plotted on a galaxy basis. Every coloured dot represents a galaxy. The black dots stand for the median halo masses  $M_{H,m}$  in each (recovered) richness bin. The error bars indicate 16 and 84 per cent in true halo mass dispersion.

members of (real) richer groups that have been fragmented, while those in yellow are ‘interlopers’ that have been incorrectly merged into larger structures in the recovered group catalogue.

Table 1 gives the median true halo mass, the 16 and 84 per cent scatter in halo mass, and the rms of the ratio between  $N_{\text{rec}}$  and  $N_{\text{true}}$  for bins of different  $N_{\text{rec}}$ . We will discuss how these are affected by different mass-selection  $M_{\text{stellar}}$  and sampling rates  $s$  in detail later in this paper.

We therefore introduce the first two metrics of the recovered group catalogue as follows (based on Fig. 9):

(i) Metric 1: the median true halo mass  $^1M_{H,m}$  of galaxies recovered as singletons (by definition on a galaxy basis). This may be thought of as giving the minimum halo masses for which there is any information about galaxy evolution.

(ii) Metric 2: the median true halo mass  $^2M_{H,m}$  of galaxies in recovered two-member groups, again on a galaxy basis. This may be thought of as giving the halo mass for which there is some non-trivial environmental information. The overall  $M_{H,m}(N_{\text{rec}})$  relation for richer groups scales quite closely with this value.

### 3.2 The recovered two-way matched groups

In addition to looking at the statistics of the dark matter haloes assigned on a galaxy-basis, we can also look at the groups directly. We have to restrict this analysis to recovered groups that have a two-way match because only for these is the allocation of a true dark matter halo to a given recovered group unambiguous. This comparison also really only makes sense for the two-way matched groups since it is only for these that there is a reasonable correspondence (to within at most a factor of two) between the true and recovered membership. Unfortunately, it is not possible for an observer to know (without having the ‘mock’ reality) which of the recovered groups are two-way matched and which are not, and so we also need to consider what fraction of the recovered groups are two-way matched.

It is important to note that, unlike most of the metrics in this section, two-way matching only makes sense in terms of a comparison with a ‘resampled’ reality. Because it involves a comparison with the number of real members in the real group, any galaxies that are excluded by random sampling must not be included. The real group membership is therefore defined here in terms of the real galaxy

population after any random sampling of the galaxies (i.e. as given by our parameter  $s$ ) has taken place.

Fig. 10 shows the distribution of true halo mass for the two-way matched recovered groups, including, with  $N_{\text{rec}} = 1$ , those galaxies that are singletons in both the (‘mock’) true and recovered catalogues. Hence, each dot now represents a single recovered group (or singleton). The points in the integer richness bins are again dispersed horizontally for clarity. In the representative recovered catalogue considered in this section, 73 per cent of all recovered groups in the catalogue (with a richness of at least 2) have a two-way match, while 89 per cent of all recovered singletons are truly singletons (and thus two-way matched). The colour-coding shows the product  $f_{c_2} f_{p_2}$ , which indicates how ‘good’ the two-way match is. It gives the product of the fractions of true group members that are found in the recovered group and visa versa. Straightforwardly from the definition of two-way matched groups, this must be at least 0.25, i.e.  $0.5 \times 0.5$ .

Table 2 gives the median true halo mass, the 16 and 84 per cent scatter in true halo mass and the mean of  $f_{c_2} f_{p_2}$  for recovered groups of different richness.

Based on Fig. 10 and Table 2, we can construct the following three further metrics:

(i) Metric 3: the median true halo mass  $^2_{2w}M_{H,m}$  of recovered two-way matched groups with two members. This metric is quite similar to Metric 2. Again, this sets a scaling for the overall  $^2_{2w}M_{H,m}(N_{\text{rec}})$  relation.

(ii) Metric 4: the two-way purity  $p_2$  of all recovered groups of two or more members. This metric simply tells us the fraction of recovered groups that have a two-way match to a real group (of at least two members, by construction).

(iii) Metric 5: the mean product of  $(f_{c_2} f_{p_2})$  for the two-way matched groups, averaged over all two-way matched recovered groups. This metric, which must, by definition, take values between 0.25 and 1.0, assesses the ‘quality’ of the two-way matches. A value of 1.0 would represent perfection, a value of 0.25 would indicate that every two-way matched group had only just scraped into that class.


### 3.3 Scatter in $M_H$

The  $M_H$ -estimator based on  $\sum M_{\text{stellar}}$  and  $\sigma_v$ , presented in Section 2.5, is superior to the richness  $N_{\text{rec}}$  as a halo mass estimator. Figs 11 and 12 show the difference between the recovered (estimated) mass and the true mass on the galaxy-basis and the group-basis, respectively. As throughout this section, this is for the recovered group catalogue obtained with a mass cut of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) > 9.5$  and a sampling rate of  $s = 0.8$  in the  $1.2 < z < 1.7$  redshift range. While in Fig. 11 all galaxies (also singletons) are included, in Fig. 12, singletons are excluded and furthermore only the two-way matched groups are considered, as in the previous subsection. In both panels, the black uncertainty bars show the 16 and 84 per cent percentile scatter in bins of recovered (estimated) halo mass. For the galaxy basis, the prominent locus extending up diagonally to the right arises from very low mass (true) groups that (mostly singletons) are spuriously overmerged into much larger haloes. The discrepancy between apparent and true halo masses therefore increases towards larger (apparent) halo masses, causing a noticeable increase in the scatter at high masses. In order to capture this effect, we will define the overall scatter to be the average of the individual error bars (spaced in log apparent halo mass) in Fig. 11, rather than simply averaging over all galaxies, which would have been dominated by the lowest mass haloes that have the smallest dispersions. The average



**Table 1.** As function of  $N_{\text{rec}}$ : the median halo masses  $M_{H,m}$ , the rms scatter in halo masses and the quotients of  $N_{\text{rec}}$  and  $N_{\text{true}}$ , on a galaxy-basis (see also Fig. 9).

$N_{\text{rec}}$	1	2	3	4	5	6–7	8–10	11–14	15–20	21–28	29–39	40–	Mean	$\geq 2$
$\log(M_{H,m}/M_{\odot})$	11.88	12.3	12.62	12.8	12.92	13.11	13.2	13.28	13.58	13.72	13.5	13.79	–	–
$\text{rms}(\log(M_{H,m}/M_{H,m}))$	0.42	0.42	0.45	0.48	0.49	0.52	0.56	0.72	0.71	0.78	0.68	0.85	0.59	0.49
$\text{rms}(\log(N_{\text{rec}}/N_{\text{true}}))$	0.19	0.21	0.24	0.27	0.3	0.31	0.39	0.52	0.51	0.61	0.69	0.86	0.42	0.3



**Figure 10.** The true halo mass of recovered two-way matched groups as a function of the recovered richness, plotted on a group basis. Each coloured dot represents a single two-way matched group. The error bars indicate the median and 16 and 84 per cent halo mass dispersion in each recovered richness bin.  $F_s$  is the fraction of recovered singletons that are truly singletons, only these are plotted at  $N_{\text{rec}} = 1$ .

uncertainty calculated in this way is 0.40 dex. This effect is not present, for obvious reasons, in the group-basis analysis in Fig. 12, but we adopt the same approach to characterizing the scatter in halo masses for consistency.

For the two-way matched groups, the  $M_H$ -scatter averages to 0.19 dex. The large discrepancy between these two dispersions is mainly due to the interloper galaxies falsely merged into larger groups, for which, as discussed above, the true halo masses are significantly overestimated.

The  $M_H$ -estimator and its application to the recovered group catalogue therefore defines two further metrics:

(i) Metric 6: the scatter in halo mass  $\sigma_{\text{gal}}(M_H)$ , as given by average of the 16 and 84 per cent percentiles, averaged over (logarithmic) bins of recovered estimated halo mass ( $M_{H,\text{est}}$ ) and calculated within each bin on a galaxy basis. This metric reflects both infidelities of the group-finder (in terms of overmerging and fragmentation, spurious groups and undetected groups), and also the underlying scatter in the halo mass estimator itself.

(ii) Metric 7: the scatter in halo mass  $\sigma_{\text{gr}}(M_H)$  averaged over all logarithmic  $M_{H,\text{est}}$  bins, given by the 16 and 84 per cent percentiles for each bin, and calculated within each for two-way matched groups, but excluding the singletons. This gives a better estimate of the underlying scatter in halo mass, as the two-way matched avoid, by construction, the worst failures of the group reconstruction.

Uncertainties in the estimated stellar masses of galaxies (e.g. due to intrinsic mass-to-light ratio variations in the Universe, or observational uncertainties) may contribute an uncertainty in the halo mass estimates that use, in part, those stellar masses. We have

looked at this effect by adding scatter randomly to the stellar masses of the galaxies in Fig. 11 (Metric 6). Using a 0.1-dex ( $1\sigma$ ) scatter (see discussion in Section 1), the increase in scatter in halo mass is undetectable. Even using a much larger 0.5-dex (random) scatter in galaxy masses the additional scatter in halo mass remains small: The overall scatter increases from 0.40 to 0.42 dex.

### 3.4 Metrics for central/satellite classification

Whether a galaxy is the central galaxy or a satellite galaxy in its parent halo is thought to be a major driver of its evolution, both in the real universe, and certainly in the L-Galaxies model. Centrals exhibit different physical characteristics than satellites. The performance of the group-finder in correctly identifying centrals and satellites is important in galaxy evolution studies. For both the true groups and the recovered groups, we define for simplicity the central to be the most massive (stellar mass) galaxy in the group and all others to be satellites. Singleton galaxies (whether in the real or recovered catalogues) are therefore always centrals.

It is then a well-defined question whether a recovered central (including singletons) is really the central galaxy in its (true) halo and whether a recovered satellite is truly a satellite in its (true) halo. In short, did the group-finder correctly classify centrals and satellites, accepting singletons as stand-alone centrals. We define an accuracy  $A$  to be the fraction of objects that were correctly classified, i.e. that their classification in the recovered catalogue matches their true classification.


We stress here that the ‘true’ central/satellite classification of the galaxies is done *before* the random sampling of the galaxy sample is done (cf. the issue of two-way matched groups discussed in Section 3.2, where the two-way matching is done by comparing with the ‘re-sampled membership’ of the real group). Galaxies are therefore classified as real centrals or real satellites based on their stellar-mass ranking amongst all the members of the full real group, not just those that were randomly selected for observation.

Fig. 13 shows  $A$  for the recovered group catalogue used in this section as a function of the recovered richness  $N_{\text{rec}}$ . There is a weak overall trend of decreasing  $A_{\text{centrals}}$  with  $N_{\text{rec}}$ . Singletons ( $N_{\text{rec}} = 1$ ) are correctly identified as centrals in 90 per cent of cases; the accuracy of the central classification in groups of at least two members goes down marginally to 88 per cent and is even lower for richer groups. At first sight, both of these numbers could be surprisingly high: One might naively expect an accuracy of at most the sampling rate  $s$  (in this case is 80 per cent) since this defines the chance that a given central is even observed in the first place. However, this logic only works for high  $N_{\text{rec}}$ , since  $s$  will then reflect the chance that the true central was actually observed, and, if it was not, then for sure a true satellite will be wrongly recovered as the central. At low  $N_{\text{rec}}$ , however, this simple logic fails, as shown in Appendix B.


Turning to the satellites, the accuracy in recovering them lies around 80 per cent. Since a recovered satellite by definition has a more massive galaxy nearby, it can only have been a true central

**Table 2.** As function of  $N_{\text{rec}}$ : the median halo masses  $M_{H,m}$ , the rms scatter in halo masses, and the mean of the product  $f_{c_2} f_{p_2}$  for the two-way matched groups and the recovered singletons that are truly singletons (see also Fig. 10).

$N_{\text{rec}}$	1	2	3	4	5	6–7	8–10	11–14	15–20	21–28	29–39	40–	Mean	$\geq 2$
$\log(M_{H,m}/M_{\odot})$	11.84	12.37	12.64	12.87	12.97	13.18	13.29	13.44	13.65	14.01	13.86	14.15	–	–
$\text{rms}(\log(M_H/M_{H,m}))$	0.29	0.31	0.33	0.3	0.32	0.25	0.26	0.31	0.24	0.24	0.32	0.0	0.26	0.31
$\text{mean}(f_{c_2} f_{p_2})$	–	0.97	0.86	0.86	0.81	0.8	0.78	0.71	0.71	0.7	0.75	0.63	0.79	0.92



**Figure 11.** The logarithmic difference between the true dark matter halo mass  $M_H$  and the estimated (recovered) halo mass  $M_{H,\text{est}}$ , plotted for all groups on a galaxy basis as a function of  $M_{H,\text{est}}$ .




**Figure 12.** The logarithmic difference between the true dark matter halo mass  $M_H$  and the estimated (recovered) halo mass  $M_{H,\text{est}}$ , plotted for all two-way matched groups, excluding singletons.

because of infidelities in the group reconstruction, i.e. a real central was wrongly brought into the group through ‘overmerging’.

Because these measurements of  $A$  vary little with richness for  $N_{\text{rec}} \geq 2$ , we simply average across all the galaxies in the sample and define three further metrics as follows:

(i) Metric 8: the accuracy of recovered centrals for  $N_{\text{rec}} = 1$ ,  $A_{\text{centrals}}$ , i.e. for singletons. This states the fraction of recovered singletons that are truly centrals (including singletons) rather than truly satellites.

(ii) Metric 9: the accuracy of centrals in recovered groups with richness  $N_{\text{rec}} \geq 2$ . This states the fraction of recovered centrals (excluding the recovered singletons) that are indeed truly centrals (or



**Figure 13.** The accuracy in classifying centrals and satellites, i.e. the fraction of recovered centrals and satellites that are truly centrals and satellites, as a function of the recovered richness. Note that  $N_{\text{rec}} = 1$  therefore refers to recovered singleton galaxies.

singletons), rather than satellites, averaged across all the recovered centrals.

(iii) Metric 10: the accuracy of satellites  $A_{\text{satellites}}$ . This states the fraction of recovered satellites that are truly satellites (i.e. not truly centrals or singletons), averaged across all recovered satellites.

### 3.5 Number of groups and multiplicity


Finally, we come to two other important characteristics of the recovered group catalogue. The first is simply the number of groups in the recovered catalogue. In our illustrative catalogue ( $1.2 < z < 1.7$  redshift range with a mass cut of  $M_{\text{stellar}} > 9.5$  and a sampling rate of  $s = 0.8$  and covering  $2 \times 2 \text{ deg}^2$ ), there are 7793 recovered groups with at least two members. This may also be defined, as desired, as per area of sky, or per comoving volume.

The second is the average multiplicity, by which we mean the average group richness in the catalogue, including singletons (with  $N_{\text{rec}} = 1$ ). The multiplicity tells us how many galaxies there are on average within a (recovered) individual dark matter halo. It therefore captures the amount of ‘information’ in the group catalogue in the sense that the group catalogue loses its usefulness as the multiplicity reduces towards unity, i.e. as more and more galaxies become singletons. In the illustrative group catalogue considered here, the multiplicity is 1.34.


Hence, we introduce the two last metrics:

(i) Metric 11: the number of groups  $N_{\text{grs}}$  in the recovered group catalogue. This assesses the size of the group catalogue and could usefully be expressed if desired as per comoving volume or surface area of sky.

(ii) Metric 12: the multiplicity of the recovered group catalogue. This measures the average richness (number of members) of the structures (including singletons) in the recovered group catalogue.



**Figure 14.** Comparison of the two-way purity  $p_2$  in the ‘halo-based’ and FoF recovered catalogues as a function of the sampling  $s$ . See the text for details.



**Figure 15.** Two-way completeness  $c_2$  of the ‘halo-based’ and FoF methods.


#### 4 PERFORMANCE OF ‘HALO-BASED’ AND FOF METHODS AT LOW REDSHIFT

We now turn to compare the performance of the ‘halo-based’ and the FoF group finding approaches. We do this at low redshift, using an L-Galaxies SAM light-cone in the redshift range  $0.02 < z < 0.17$  with a flat stellar mass cut of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.0$  but varying the sampling rate in the range of  $s = 0.1-1.0$ . This analysis is conducted at low redshift because the nearby Universe is well known, the SAM is probably most reliable, and the astrophysical knowledge used in the more refined ‘halo-based’ approach is probably better established.


We focus in this first analysis on the most basic performance metrics assessing the quality of the recovered group catalogues: completeness, purity and the fraction of interlopers  $f_i$  in the catalogue, as well as the fraction of galaxies in true groups (of at least two members),  $S_{\text{Gal}}$ , which are successfully put in a recovered group (of at least two members) by the group-finder. In our previous terminology, completeness and purity are ‘group-based’ quantities, meaning that they assess the performance of the group-finder in recovering groups, while  $f_i$  and  $S_{\text{Gal}}$  are ‘galaxy-based’ quantities.

Fig. 14 shows the two-way purity  $p_2$  (Metric 4) of the ‘halo-based’ and FoF recovered group catalogue over a wide range of random sampling rate  $s = 0.1-1.0$ . There is an overall improvement in the purity with  $s$ , reaching up to 78 per cent in the ‘halo-based’ method and 74 per cent in FoF. The dependence is quite weak, and even at very low  $s$ , the purity remains at 62 per cent for the ‘halo-based’ method and 66 per cent for FoF. Overall the two group-finders perform very similarly, with only small differences: As might be expected for a more ‘refined’ approach, the ‘halo-based’ method performs better for complete samples (e.g. 4 per cent better than FoF at  $s = 1.0$ ), while for seriously incomplete samples,  $s \lesssim 0.6$ , the less refined FoF yields better results. The ‘halo-based’ method is the more sophisticated group-finder in the sense that it uses more information (knowledge of the size of a halo as a function of mass), the strength of the FoF group-finder possibly lies in its use of minimal information. Hence, while the ‘halo-based’ method performs better for complete information, FoF performs better when the amount of available information is degraded by incomplete random sampling.

For the two-way completeness  $c_2$ , we observe the same trends, as shown in Fig. 15: Again, both group-finder perform very similarly, but the ‘halo-based’ method performs slightly better at high  $s$  but worse at low  $s$ : At  $s = 1.0$ , the ‘halo-based’ method has  $c_2 = 78$  per cent, while FoF reaches  $c_2 = 77$  per cent but FoF performs better at  $s \lesssim 0.8$ , declining down to 68 per cent at  $s = 0.1$ , while



**Figure 16.** Comparison of the fraction of interlopers  $f_i$  in the ‘halo-based’ and FoF methods, as a function of the sampling  $s$ .



**Figure 17.** Comparison of the fraction of galaxies correctly recovered in groups,  $S_{\text{Gal}}$ , in the ‘halo-based’ and FoF methods, as a function of the sampling  $s$ .

the ‘halo-based’ method reduces more rapidly down to 64 per cent. Overall, the two-way completeness is (perhaps) surprisingly rather stable over this wide range of sampling rate in both group finding approaches.

While  $c_2$  and  $p_2$  assess the quality of the recovered group catalogue on a group basis,  $f_i$  and  $S_{\text{Gal}}$  assess the quality of the recovered group catalogue on a galaxy basis, considering galaxies and their environment individually:  $f_i$  states the fraction of interlopers in groups, i.e. the fraction of galaxies that are put into a group of at least two members by the group finding algorithm, but are truly singletons;  $S_{\text{Gal}}$  states the fraction of galaxies truly populating groups of at least two members, which are recovered in groups of at least two members.


Fig. 16 shows that the fraction of interlopers  $f_i$  increases from 11 per cent in the ‘halo-based’ and FoF methods for  $s = 1.0$  full sampling, up to 22 per cent for FoF and 25 per cent for the ‘halo-based’ method with a  $s = 0.1$  sampling rate. Again, while they perform comparably at high sampling rate, FoF performs slightly better for low sampling rate. Overall, the fraction of interlopers  $f_i$  changes significantly with  $s$ .

The fraction of true galaxies successfully recovered in groups is shown in Fig. 17. Again, the ‘halo-based’ method performs significantly better at  $s = 1.0$  with  $S_{\text{Gal}} = 89$  per cent, compared to  $S_{\text{Gal}} = 83$  per cent for FoF, while at low  $s$  for the ‘halo-based’ method, this reduces to 74 per cent while FoF declines only to 79 per cent. As for the other quality metrics, FoF performs better at a low sampling rate, while the ‘halo-based’ method performs better at a high sampling rate.

The 12 science metrics can be found in Appendix A. We find that most metrics are virtually indistinguishable at the same sampling rate  $s$ , except the Metrics 6 (Fig. A1e) and 7 (Fig. A1f) concerning the scatter in dark matter halo mass metrics on a galaxy-basis and on a group-basis, respectively. While the estimated dark matter halo mass in the ‘halo-based’ catalogues is recovered by the built-in halo mass estimator in the group finding algorithm, for the FoF catalogues, we use the halo mass estimator presented in Section 2.5 based on total stellar mass and velocity dispersion. As in the basic statistical metrics completeness and purity, FoF exhibits the slightly lower dark matter halo mass scatter  $\sigma(M_{\text{H}})$  at low  $s$  than the ‘halo-based’ method, on both, galaxy-basis and group-basis.

As a general conclusion, both of the two group-finders compared here perform rather similarly, despite their quite different conceptual





**Figure 18.** Metric 4: the purity  $p_2$ . Changing  $M_{\text{stellar}}$  and  $s$  does not significantly change  $p_2$ . See the text for discussion. The red lines, reproduced from Fig. 34 later in this paper, show the constant (comoving) number density of available (post-sampling) tracers.

approaches. In detail, the ‘halo-based’ approach, using more external information (e.g. the halo size–mass relation), performs slightly better when the observational information level is high, from high sampling  $s \approx 1$ . But as the completeness of the observational information reduces  $s < 1$ , we find that the resulting degradation of group-finding performance is less for FoF, and it ends up performing slightly better than the ‘halo-based’ approach at low sampling rates.

For the rest of this paper, we will focus on the expected performance of group finding algorithms at high redshift. Since the sampling rate is unlikely to be extremely high, and because of the comparable overall performance of the two approaches established in this section, we will for simplicity henceforth focus on the FoF approach alone.


## 5 PERFORMANCE OF FOF CATALOGUES AT HIGH REDSHIFT IN $s/M_{\text{STELLAR}}$ -SPACE

### 5.1 The redshift range $1.2 < z < 1.7$

In this section, we investigate how the fidelity of the recovered FoF group catalogues obtained from the L-Galaxies mock at high redshift depends on the lower stellar mass cuts, over the range  $9.1 \leq \log_{10}(M_{\text{stellar}}/M_{\odot}) \leq 10.5$ , and on the random sampling rate in the range  $0.1 \leq s \leq 1.0$ . The former produces a wide range in comoving density of the potential tracers, while the latter multiplies this to yield the actual final number density of tracers. In the plots to follow, the thin red curves represent loci of constant tracer number density.

We first focus on the 12 ‘science metrics’ that we introduced in Section 3 that are obtained in the redshift range  $1.2 < z < 1.7$ , before then briefly examining how the results differ with redshift, looking at two other redshift intervals by examining slightly lower  $0.9 < z < 1.1$  and higher  $2.0 < z < 2.6$  redshifts.

Returning to  $1.2 < z < 1.7$ , Fig. 18 first presents the basic two-way purity  $p_2$  (Metric 4), the fraction of recovered groups with a




**Figure 19.** Metric 2: the true median halo mass  $^2M_{H,m}$  of galaxies that are recovered to be members of richness 2 groups depends primarily on the number density of tracers (indicated by the red lines).

two-way match to a real group. It is therefore one of the most basic measures of the overall fidelity of the group catalogue. It should however be remembered that it is also one of the more artificial metrics because it compares the recovered groups to the ‘resampled’ reality obtained *after* the application of the random sampling  $s$  to the full mock sample. Probably for this reason,  $p_2$  is found to be largely independent of both  $M_{\text{stellar}}$  and  $s$ . Despite being quite noisy, values between  $\approx 60$  and 80 per cent are found, with many close to 0.72, and there is no obvious trend across the diagram. Once optimized, the FoF algorithm is evidently able to reconstruct the underlying ‘resampled’ group population more or less independently of how much this has been degraded from the full underlying population by the incomplete sampling.


It should be noted that this same characteristic value was also found in the SDSS-like comparison of the FoF and ‘halo-based’ approaches (see Section 4) and therefore reflects a fairly robust measure of the performance of optimized algorithms of different sorts. That this  $p_2$  purity Metric 4 nevertheless has a more or less uniform value that lies well below unity (i.e. perfection) is almost certainly a largely unavoidable consequence of the reduced phase-space information on galaxy positions (two spatial positions and one velocity measure) that is available to the astronomer trying to use galaxy positions as tracers of underlying structure.

In Figs 19–21, we show the changes in, respectively, the median halo mass for recovered  $N_{\text{rec}} = 2$  structures on a galaxy basis (Metric 2), for the two-way matched recovered two-member groups (Metric 3), and for haloes populated with only a recovered singleton (Metric 1).

For the first two, the overall trend is for the mass of two-member groups to decrease with both increasing sampling rate and with decreasing limiting stellar mass (towards the upper left-hand side of the figures). These two parameters determine the (comoving) number density of the tracers that are available for the group reconstruction (i.e. after the observational sampling has been applied to the tracer target population). As noted above, the red lines in the figures,




**Figure 20.** Metric 3: the median halo mass  ${}_{2w}^2M_{H,m}$  of two-way matched recovered groups of richness 2 depends mostly on the number density of tracers (indicated by the red lines).



**Figure 21.** Metric 1: The median halo mass  ${}^1M_{H,m}$  of singletons depends mainly on the mass cut  $M_{\text{stellar}}$ . This is due to a dominant fraction of singletons: Singletons after sampling were likely to be singletons in the original full catalogue.

and in all later figures in this section, show the lines of constant (comoving) number density of the available (post-sampling) tracers: Not surprisingly, increasing the number density of these tracers enables lower mass structures to be identified as multi-member groups.

To first order, this gain is independent of whether the number density is increased by lowering  $M_{\text{stellar}}$  or by raising  $s$ . The fact



**Figure 22.** Metric 6: This plot shows the average scatter in halo mass  $\sigma_{\text{gal}}(M_{\text{H}})$  of the galaxy groups (given by the average of the 16 and 84 per cent percentiles, averaged across all the logarithmic bins in mass), considered on a galaxy basis. This depends primarily on  $s$ .


that these metrics largely reflect the resulting final number density of tracers rather than the actual mass limit strongly suggests that these results will also be valid for other tracer selection criteria, e.g. selection by flux in some band. Any ‘fuzziness’ of the mass-selection limit (at constant number density of tracers) due to uncertainties in the individual masses of galaxies is unlikely to have a large effect on these metrics once the number of density of tracers is fixed.

The dependence of these two Metrics 2 and 3 on  $s$  weakens at high  $s$  due to a survivor effect. Most of the recovered two-member groups, after the sampling, are truly two-member groups and are just the ‘lucky survivors’ in which both members were selected in the random sampling.

This ‘survivor-effect’ is even more clearly seen for the singletons in Fig. 21: Nearly all the recovered singletons, even after sampling, are truly singletons in the (full) mock catalogue, rather than being the result of incompletely sampled groups. As  $s$  is reduced, fewer galaxies (including singletons) are observed, but most of the recovered singletons would still have been singletons with higher  $s$ . This is why there is very little change in the median halo mass of singletons with  $s$ . As a result,  ${}^1M_{H,m}$  is primarily affected by the stellar mass cut  $M_{\text{stellar}}$  alone and is found to be roughly proportional to  $M_{\text{stellar}}^{0.5}$ , as expected from the cosmic halo–stellar mass relation in the mock (and real) universe.

Two further metrics focus on the uncertainty (i.e. scatter) in halo mass  $\sigma(M_{\text{H}})$ , averaged over bins in recovered halo mass. This scatter is calculated on a galaxy basis (Metric 6; see Fig. 22), and for two-way matched groups on a group basis (Metric 7; see Fig. 23).

As noted above, there will always be an intrinsic scatter in  $M_{\text{H}}$  in any recovered group catalogue, however good it is, because of underlying astrophysical or cosmological scatter in the galaxy content of haloes. This underlying scatter is of course also present in the true ‘mock’ catalogue. The uncertainty in halo mass in any recovered group catalogue is a result of (a) the almost irre-




**Figure 23.** Metric 7: The average scatter in halo mass  $\sigma_{\text{gr}}(M_{\text{H}})$  of the galaxy groups (given by the average of the 16 per cent and 84 per cent percentiles averaged across all the logarithmic bins in mass) of the two-way matched groups. This also depends primarily on  $s$ .

ducible scatter in the Universe and (b) of failures in the group-finding.

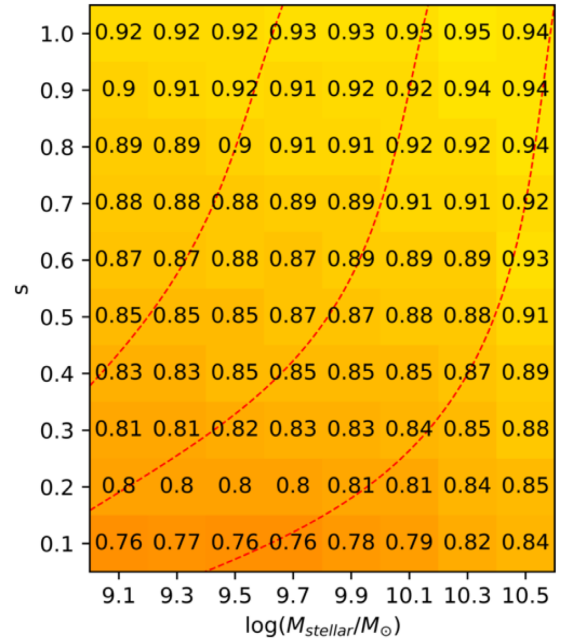
In order to assess the relative importance of these two effects we ran the halo mass estimator presented in Section 2.5 on perfectly sampled true group catalogues. We found that the intrinsic scatter varied from 0.15 dex for a limiting stellar mass of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.1$  to 0.20 dex for  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 10.5$ . By comparing these numbers to the top line of Fig. 22, one can see that the scatter on a galaxy basis is by a factor of  $\approx 2$  higher in the reconstructed catalogues compared to the true catalogues. Accordingly, the variance changes by a factor of  $\approx 4$ . Hence, the failures in group-finding dominate over the intrinsic scatter.

For both metrics, 6 and 7, the scatter in the recovered halo mass shows very little change with  $M_{\text{stellar}}$ , but a significant gradient with  $s$ . As  $s$  increases from 0.5 to 1.0, the scatter in halo mass decreases from  $\approx 0.4$  to 0.3 dex on a galaxy basis, and from  $\approx 0.25$  to 0.18 dex on a group basis. These changes are not negligible: the variance changes by a factor of  $\approx 1.8$ . The variance represents the sum of the contributions from (i) the astrophysical/cosmological scatter, (ii) the irreducible scatter due to the incomplete phase-space information even at full sampling, plus (iii) the additional scatter arising from the variation in the fidelity of the group catalogue with  $s$ . Note that we are here considering the fidelity of the recovered groups relative to the full true mock catalogue, not to the resampled mock catalogue that was discussed in connection with Metric 4. That the variance almost doubles when  $s$  goes from 1.0 to 0.5 emphasizes the relative importance of the last term (iii) as  $s$  is reduced.

In contrast, Fig. 24 shows  $f_{c_2} f_{p_2}$  (Metric 5), which reflects the mean ‘quality’ of the two-way matches. The mean of  $f_{c_2} f_{p_2}$  hardly changes at all with  $s$  and  $M_{\text{stellar}}$  and has a more or less constant high value of  $\approx 0.92$ . As with Metric 4 above (see Fig. 18), this also reflects the robust performance of the FoF group-finder relative to the ‘resampled’ reality, even as that degrades relative to the full



**Figure 24.** Metric 5: the fidelity of the group-finder given as  $f_{c_2} f_{p_2}$ . This is almost completely independent of both  $s$  and the number of tracers. However, it should be noted that this metric compares the recovered catalogue with the resampled true catalogue, i.e. it does not include any degradation that arises because of the random sampling.




**Figure 25.** Metric 8: the accuracy of central classification for singletons in the recovered group catalogue  ${}^1A_{\text{centinals}}$  for  $N_{\text{rec}} = 1$ . The performance of the group-finder in recovering those improves with the sampling rate  $s$  and slightly with a more restrictive (increasing) mass cut  $M_{\text{stellar}}$ .


unsampled reality. Again, the difference with Metrics 6 and 7 is noticeable.

We next look at the performance of the group-finder in accurately classifying singletons (Fig. 25), centrals (Fig. 26) and satellites (Fig. 27) correctly. For the centrals, it will be recalled that we






**Figure 26.** Metric 9: the accuracy of the centrals  $\geq 2 A_{\text{centrals}}$ , i.e. the fraction of recovered centrals (excluding singletons), which are truly centrals. This improves quite significantly with sampling rate  $s$ .



**Figure 27.** Metric 10: the accuracy of the satellites  $A_{\text{satellites}}$  states the fraction of recovered satellites, which are truly satellites. The accuracy of the satellite classification is largely independent of both the mass cut  $M_{\text{stellar}}$  and the sampling rate  $s$ .

distinguish between centrals of recovered groups of at least two members (Metric 9), and the recovered singletons (Metric 8).

We first note that the accuracy of the classification of centrals  $\geq 2 A_{\text{centrals}}$  in the ( $N_{\text{rec}} \geq 2$ )-structures improves very significantly with higher sampling rate  $s$  (and much more weakly with  $M_{\text{stellar}}$ ), from a minimum of 56 per cent at a sampling rate of  $s = 0.1$  up to maximally 94 per cent at  $s = 1.0$ . It is noticeable that the accuracy



**Figure 28.** Metric 11: number of recovered groups  $N_{\text{grs}}$  (in log) of at least two members.


of classifying singletons  $^1 A_{\text{centrals}}$  (i.e. the centrals of ( $N_{\text{rec}} = 1$ )-structures) is better than for ( $N_{\text{rec}} \geq 2$ )-structures. This is because, as noted above, most of the recovered singletons are the ‘lucky survivors’ of the sampling and are true singletons in the original full mock catalogue, as opposed to being degraded two- or more member groups structures, in which there is a danger of incorrect classification (the true unsampled group catalogue is dominated by small richness structures). The parameter  $^1 A_{\text{centrals}}$  goes from a minimum 76 per cent at a sampling rate of  $s = 0.1$  up to a maximum 95 per cent at  $s = 1.0$ .

It can be seen that the accuracy of central classification, both for  $N_{\text{rec}} = 1$  singletons and for the centrals of richer groups, improves slightly as  $M_{\text{stellar}}$  is increased. This is due to the decreasing multiplicity of the overall group catalogue and the associated increase in the ‘survivor-effect’, since more galaxies will be true singletons in the original (true) sample.

In contrast to the centrals, the accuracy of the satellite classification  $A_{\text{satellites}}$  is remarkably constant at  $\sim 80$  per cent. Since recovered satellites are by definition ranked below other more massive galaxies in the recovered groups, errors in satellite classification will only occur due to errors in assigning galaxies to groups, i.e. the basic underlying fidelity of the group reconstruction in terms of completeness and purity. As shown for the Metrics 4 (Fig. 18) and 5 (Fig. 24), this is only very weakly dependent on  $s$  and  $M_{\text{stellar}}$ .

Fig. 28 shows the number of recovered groups of at least two members (Metric 11). This number is quite sensitive to  $s$  and  $M_{\text{stellar}}$  and, as might be expected, follows quite closely the red lines of constant number of tracers, independent on whether this density is achieved through the initial selection of the tracers or their random sampling rate.

Finally, we show in Fig. 29 how the average group multiplicity (Metric 12) changes with  $s$  and  $M_{\text{stellar}}$ . As might be expected, the overall trend is the same as for the number of recovered groups, and largely follows the number density of available tracers after the application random sampling.



**Figure 29.** Metric 12: the multiplicity  $m$ , which is the average richness of recovered galaxy groups (including singletons). The multiplicity decreases with decreasing sampling rate  $s$  and increasing mass cut  $M_{\text{stellar}}$  and largely follows the number density of tracers (given by the red lines).

## 5.2 The variation with redshift

In this section, we simply and concisely compare all 12 performance metrics, and their dependence on  $s$  and  $M_{\text{stellar}}$  over three almost contiguous redshift ranges  $0.9 < z < 1.1$ ,  $1.2 < z < 1.7$ , and  $2.0 < z < 2.6$ . Exactly the same procedures, discussed above in detail for the  $1.2 < z < 1.7$  redshift range, are followed for the other two redshift ranges and all three are then presented together with the same colour scales to allow easy visual comparison. Figs 30–32 present the 12 metrics in each redshift range. In order to compare the 12 metrics over all redshift ranges, the colour-coding is set to be the same over all three figures. The goal here is to allow a simple visual comparison of the three redshift ranges.

The most noticeable change over the three redshift ranges lies in the changing multiplicity of the recovered structures (at fixed  $s$  and  $M_{\text{stellar}}$ ). This is simply due to the overall growth of structure over this redshift range in the mock (and real) universe. Associated with this reduction in multiplicity is an improvement(!) with redshift in the accuracy of central classification,  $A_{\text{centrals}}$  (for  $(N_{\text{rec}} = 1)$ -, as well as  $(N_{\text{rec}} \geq 2)$ -structures) over the three redshift ranges at fixed  $s$  and  $M_{\text{stellar}}$ . This is because the number of potentially contaminating satellites decreases at higher redshift because of the decrease in multiplicity.

It is noticeable how the quality of the recovered group catalogue expressed by the two-way purity  $p_2$  and the fidelity of the two-way matching represented by  $f_{c_2} f_{p_2}$  is again strikingly independent of the redshift range. Also, the median halo masses  $M_{\text{H}, m}$  of singletons, two-member groups as well as two-way matched two-member groups (at fixed  $s$  and  $M_{\text{stellar}}$ ) are largely independent of the redshift range, although this of course also reflects the details of the L-galaxies SAM.

We also see a small decrease in the halo mass scatter  $\sigma(M_{\text{H}})$ , especially on a galaxy basis, with increasing redshift. This effect

is again due to the changing multiplicity with redshift due to the growth of structure. At lower redshift, the average  $N_{\text{rec}}$  is higher, and the recovered group catalogue is not quite so dominated by very low richness structures. Further, the scatter in dark matter halo mass increases with  $N_{\text{rec}}$  (and correlated with  $\sum M_{\text{stellar}}$ ), mostly due to the increasing number of interlopers, when calculated on a galaxy basis. Hence, the average scatter in halo mass over all  $\sum M_{\text{stellar}}$ -bins, on a galaxy basis, is higher at low redshift for fixed  $s$  and  $M_{\text{stellar}}$ .

## 5.3 Science example: quenching fractions

In this section, we discuss a single example of a possible science investigation that can be investigated with a galaxy group catalogue of the type discussed here. Based on a recovered galaxy group catalogue and a dark matter halo mass estimator, the fraction of central and satellite galaxies that are ‘quenched’ (i.e. which have their sSFR reduced by a large amount relative to normal ‘Main Sequence’ galaxies) can be investigated as a function of the parent dark matter halo mass. The true effect will however be perturbed by both errors in classifying centrals and satellites (see Metrics 8–10), in associating galaxies to haloes of a particular mass and in assigning a mass to that group (both of which are combined into Metric 6).

The quenching threshold is, in our example, set at a specific star formation rate (sSFR) of  $0.22 \text{ Gyr}^{-1}$ , taken from the L-Galaxies SAM. The left-hand panel of Fig. 33 shows the true (unsampled) quenched fractions (solid lines), taken from the SAM, and the recovered quenched fraction (dots) that is constructed from the recovered group catalogue, again using for illustration a sampling rate of 80 per cent and a stellar mass cut of 9.5 solar masses (in log). Both are plotted as functions of the dark matter halo mass (true or estimated recovered, respectively). The green dashed line shows the recovered satellite fraction as a function of estimated recovered halo mass. While the recovered group catalogue reproduces the overall trends, significant inaccuracies are introduced due to the incomplete sampling and other failures in the group reconstruction.

Information obtained from our different science metrics (about the accuracy in identifying centrals and satellites, and interloper fractions) can be used to try to correct the observed central and satellite quenched fractions for these shortcomings in the recovered galaxy group catalogue. Of course, one could apply ever more sophisticated corrections to the observed quenching fraction using more and more information from the mock universe, until one ended up with a scheme that simply fully recovered the underlying mock. This would be a meaningless exercise. Rather, our goal here is to see how easily one can recover a good approximation to reality with very simple corrections.

The most straightforward (and traditional) correction would be simply to correct the raw observed quenched fractions of the two populations at a given mass, using the cross-contamination of centrals and satellites given by the accuracy Metrics 8–10. The middle panel of Fig. 33 shows the quenching fractions corrected for the (in)accuracy in identifying centrals and satellites by solving the simultaneous equations for  $f_{\text{Q, corr, cen/sat}}$ :

$$f_{\text{Q, obs, cen/sat}} = f_{\text{Q, corr, cen/sat}} A_{\text{cen/sat}} + f_{\text{Q, obs, sat/cen}} (1 - A_{\text{cen/sat}}). \quad (16)$$

The accuracies in recovering centrals,  $A_{\text{cen}}$  (here including singletons and higher richness centrals) and satellites,  $A_{\text{sat}}$ , are simply taken

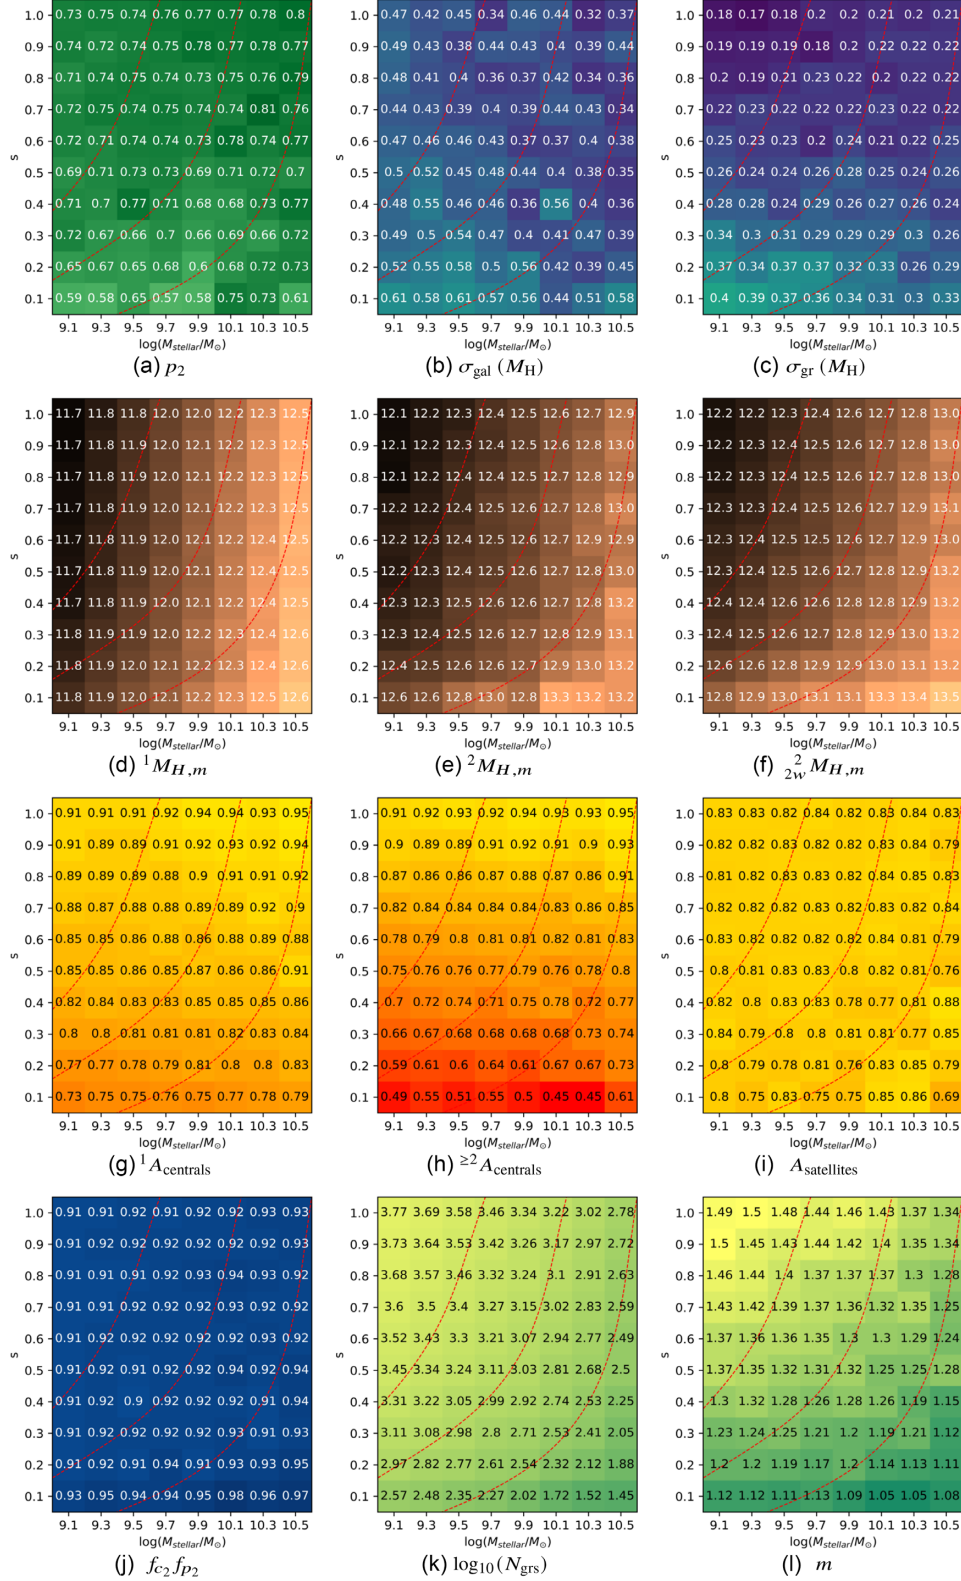
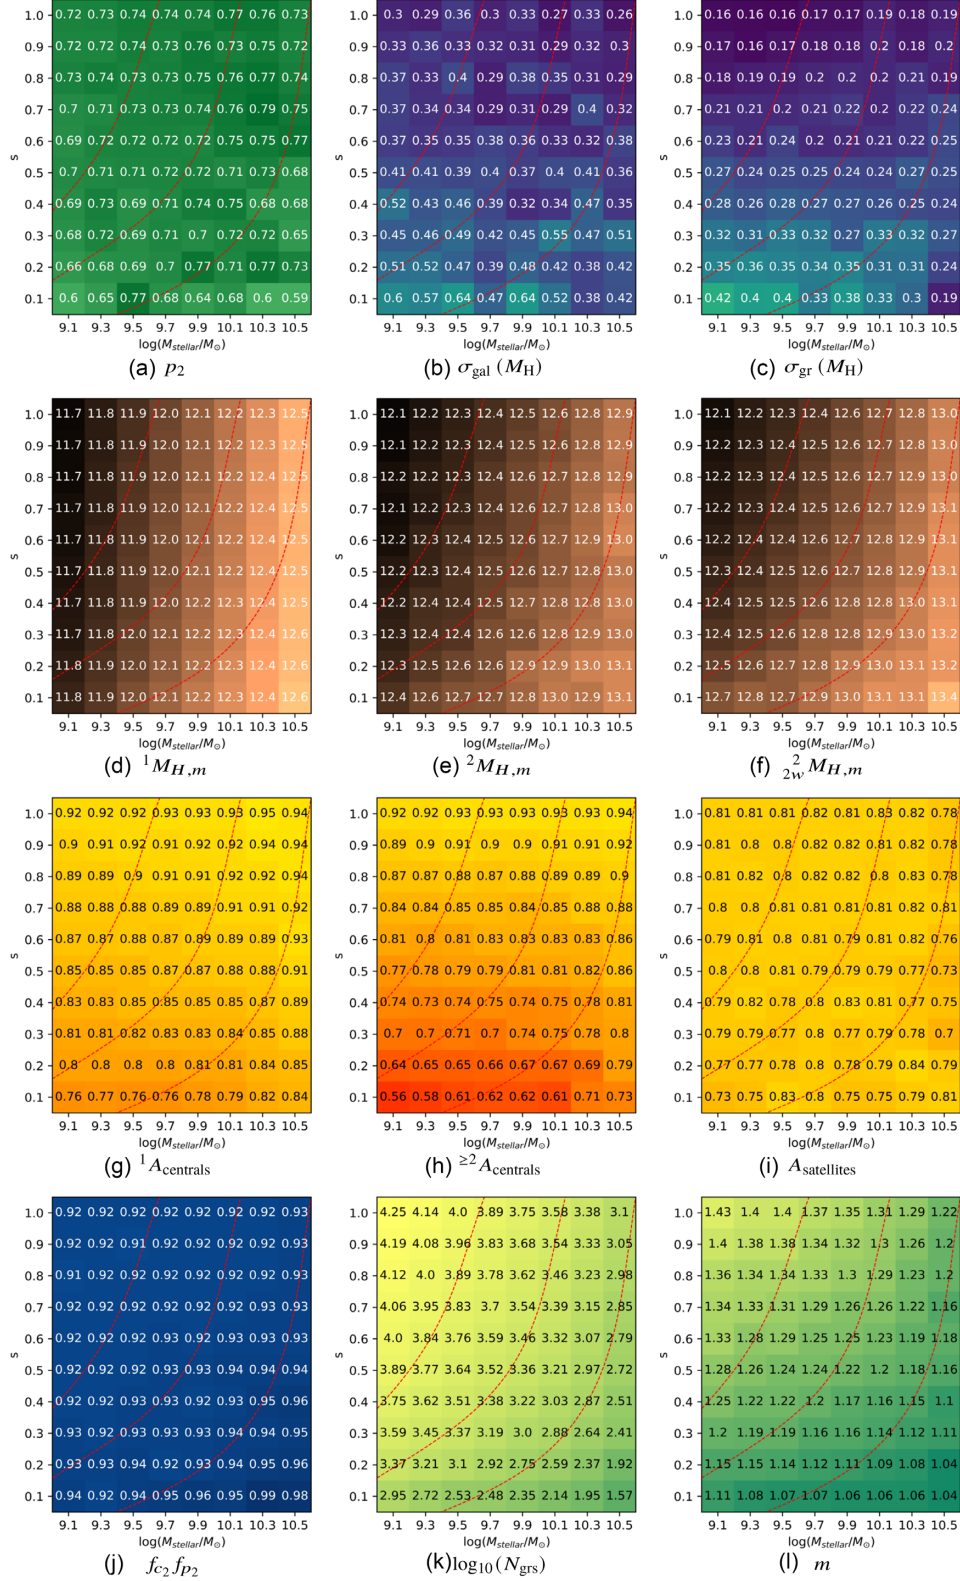


Figure 30. Metrics of the FoF group-finder for  $0.9 < z < 1.1$ .

from Metrics 8–10, i.e. they are the average ‘global’ values across all richnesses, not those determined at that particular recovered dark matter halo mass. This is a justifiable simplification because there is little dependence on mass in Fig. 13. On the other hand, the

$f_{\text{Q, obs, sat/cen}}$  is the uncorrected quenched fraction within that particular dark matter halo bin. We therefore refer to this as a ‘local correction’. It can be seen that already this simple correction is much closer to the real (mock) quenched fractions.



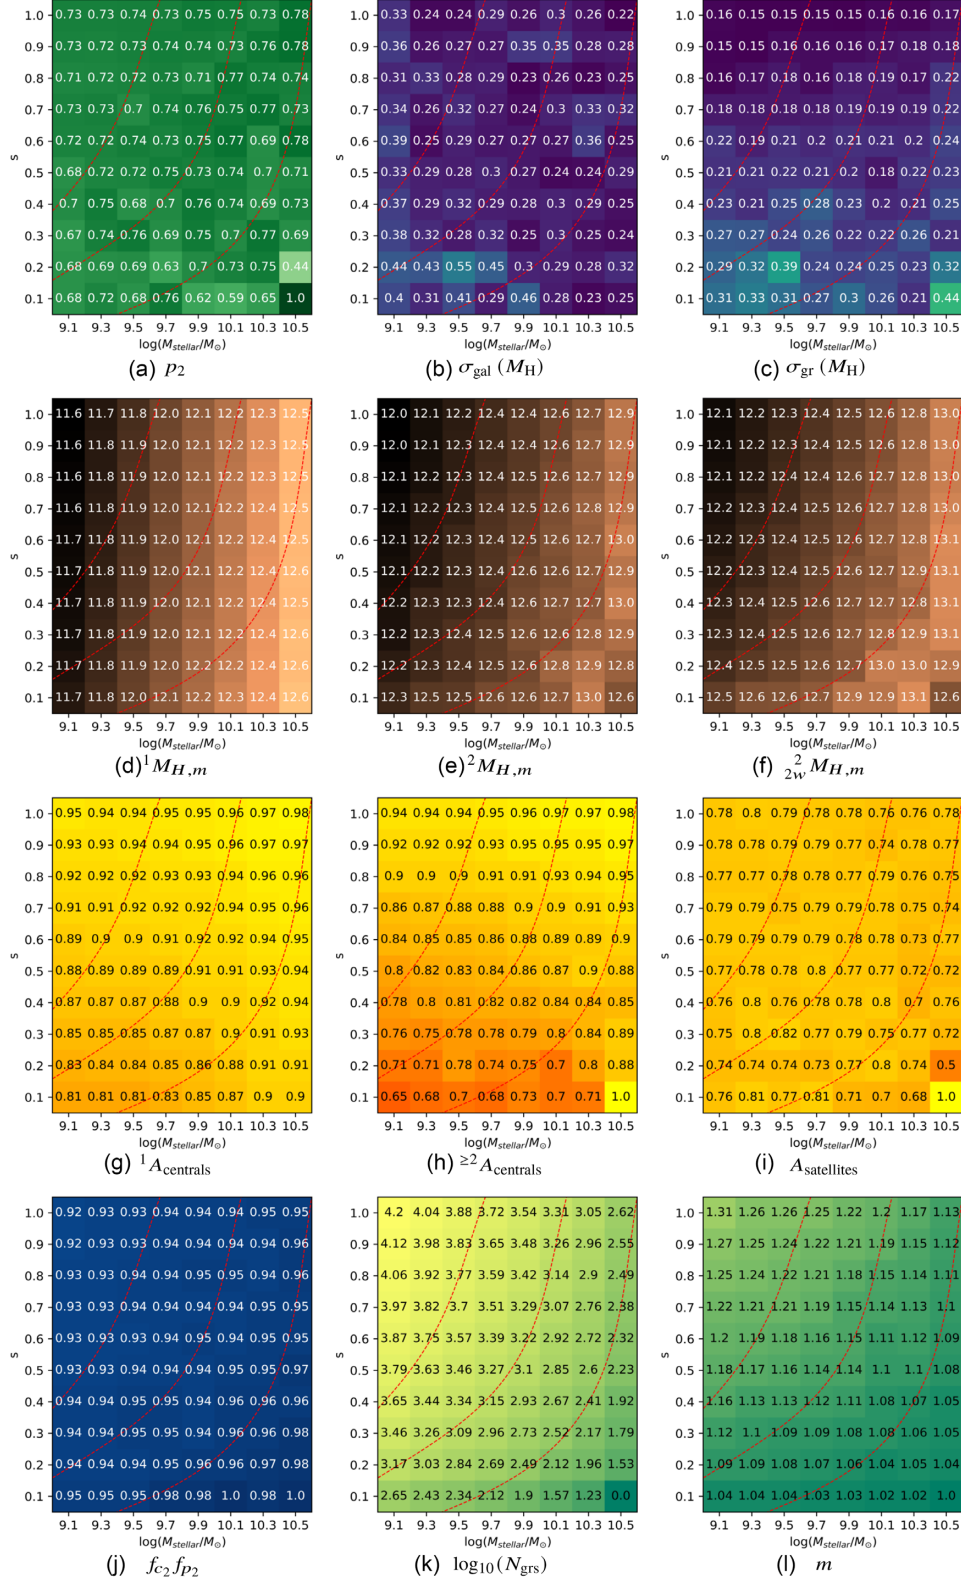


**Figure 31.** Metrics of the FoF group-finder for  $1.2 < z < 1.7$ .

Of course, interlo per centrals or satellites are unlikely to have come from haloes of the same mass, and so another simple correction would be to instead use the global quenching fraction  $f_{Q, \text{obs, sat/cen, glib}}$  of the complement (centrals for satellites and vice versa), the global

interloper fraction  $f_{i, \text{glib}}$  and the global singleton quenched fraction  $f_{Q, s}$ . This makes the simple assumption that the classification failures involve galaxies scattered in from anywhere on the recovered dark matter halo mass axis. We call this a ‘global’ correction and therefore





**Figure 32.** Metrics of the FoF group-finder for  $2.0 < z < 2.6$ .


correct the quenching fractions by solving

$$f_{Q,\text{obs,cen}} = f_{Q,\text{corr,cen}} A_{\text{cen}} (1 - f_{1,\text{glb}}) + f_{Q,\text{s,glb}} A_{\text{cen}} f_{1,\text{glb}} + f_{Q,\text{obs,sat,glb}} (1 - A_{\text{cen}}) \quad (17)$$

and

$$f_{Q,\text{obs,sat}} = f_{Q,\text{corr,sat}} A_{\text{sat}} + f_{Q,\text{obs,cen,glb}} (1 - A_{\text{sat}}) \quad (18)$$

for  $f_{Q,\text{corr,cen/sat}}$ . This correction is presented in the right-hand panel of Fig. 33.



**Figure 33.** The solid lines show the true (unsampled) quenched fractions, the dots the recovered quenched fractions, as a function of the true and recovered dark matter halo mass respectively, for centrals and satellites. The recovered group catalogue was generated for a sampling rate of 80 per cent and a stellar mass cut of 9.5 solar masses (in log) in the  $1.2 < z < 1.7$  redshift bin. Left-hand panel: uncorrected central and satellite quenched fractions. The recovered satellite fraction as a function of estimated dark matter halo mass is given by the green dashed line. Middle panel: Locally corrected central and satellite quenched fractions for global accuracies (see the text for details). Right-hand panel: Global accuracy and interloper corrected central and satellite quenched fractions (see the text for details).

This further improves the recovery of the true quenched fraction. At the very high dark matter halo mass end, however, the corrected quenched fractions of both centrals and satellites are now overcorrected, and are noticeably worse than the ‘local’ correction. This is not too surprising. For false centrals, it is probably unrealistic that satellites or interlopers from much lower dark matter haloes will have been falsely identified as the central of very high mass recovered groups: It is more likely that a satellite in the same halo is misidentified as the central because the true central was not observed. For the satellites, it is likely that using the low global central quenched fraction  $f_{Q, \text{obs, cen, glb}}$  leads to an overcorrection. Accordingly, we adopt the local correction at very high masses. This combined correction, even though very simple, yields quenching fractions very close to the real mock quenched fractions.

In summary, this short science example provides a simple illustration of the recovery of accurate scientific information using only information from high level metrics in a very simple way.

## 6 SURVEY COSTS ANALYSIS IN $s/M_{\text{STELLAR-SPACE}}$

In Section 5 of this paper, we explored how the science return from a group catalogue constructed from a redshift-survey at high redshift depends on two key parameters of the survey design: the limiting mass of the tracer galaxy population (or more generally the number density of tracers) and the sampling rate of those tracers, i.e. what fraction of those tracers are observed and/or yield a usable redshift. In this section, we analyse the ‘costs’ of these different possible survey designs in terms of required observation time in order to better understand the trade-offs involved.

For a given amount of telescope observation time, there are clear trade-offs in the survey design parameters. Going deeper (i.e. lowering the stellar mass limit  $M_{\text{stellar}}$  of the tracers) and achieving a high sampling rate  $s$  are likely to be expensive, in the sense of requiring more observing time to cover a given area, or equivalently permitting only a smaller area to be covered in a given total amount of observing time. The optimum choice of survey design may well be different for different science goals. While these may to a certain degree be accommodated through a wedding-cake approach of different nested surveys of different depths and areas, the final detailed design requires the best possible quantitative understanding of the trade-offs between science performance and cost.

This paper has aimed to provide a general framework for considering these trade-offs in the area of galaxy group science. However, since some of the factors are undoubtedly instrument-specific, we will use for illustration in this section recent simulations of the multiplexing capabilities of the MOONS instrument (using the path analysis algorithm Faststar) kindly provided by co-workers on the MOONS project at IASF-MI.<sup>6</sup> However, since the ultimate performance of the MOONS instrument on the sky will not be known for some time, and since the efficiency of fibre positioning algorithms is still being developed, we will focus in this analysis primarily on relative effects. This also helps to ensure that this analysis will also have validity for other (similar) observational instruments and survey designs.

### 6.1 Breakdown of survey costs

In this section, we present a generalized breakdown of the costs of a survey into four factors. This should be applicable to any redshift survey carried out with a multiplexed spectrograph that can simultaneously obtain spectra of multiple sources. Our approach is to construct a flexible formalism in which the increasing levels of detail can be easily incorporated as desired.

We will assume for simplicity that the observations are broken down into units of some convenient integration time, during which the configuration of the instrument, e.g. the location of individual slits or fibres, is not changed. It may well be that some sources may require much more integration time than others: an obvious example might be passive galaxies that will require much longer integration times than star-forming galaxies of the same stellar mass. Another possibility would be very faint but rare very high redshift galaxies. We will however assume that these different needs can be accommodated by retaining some fibres or slits on those objects for subsequent configurations, while redeploying the others on to new ‘normal’ targets.

The total telescope observation time  $T_{\text{obs}}$  required to complete the observations in a given area of sky, which we, for convenience, set equal to the patrol field of view (FoV) of the instrument, may then be written in quite general terms, as follows. We define the total number

<sup>6</sup>Andrea Belfiore, Paolo Franzetti, Bianca Garilli, Adriana Gargiulo et al.; <https://www.iasf-milano.inaf.it/>.

of the tracer target population within this given area to be  $N_{\text{obj}}$ , so the number that is to be actually observed is  $s \times N_{\text{obj}}$ . In the context of the earlier sections of this paper,  $N_{\text{obj}}$  will be a function of the limiting stellar mass  $M_{\text{stellar}}$ , but this can be completely generalized for any selection function. We also define the ideal multiplexing capability of the instrument in a particular observing mode to be  $N_{\text{fibre}}$  (though the formalism applies equally well to slits). It is then clear that  $T_{\text{obs}}$  may be written quite generally as

$$T_{\text{obs}} = \bar{T}_{\text{exps}} \times (s \times N_{\text{obj}}(M_{\text{stellar}})) \times \frac{1}{N_{\text{fibre}}} \times \epsilon^{-1}(s, N_{\text{obj}}). \quad (19)$$

The first three terms are straightforward: the average exposure time per object, the total number of objects to be observed, and the number of objects that can, in principle, be observed simultaneously. The final efficiency term is also conceptually simple, as it is just the fraction of the fibres that can in practice actually be usefully allocated to a target, averaged over all of the fibre configurations that are required to complete the observations in this one FoV.


In practice, this efficiency term will likely be quite complicated and will depend on the details of the technical design of the instrument and on the mode of operation of it. While it will generally be possible to use the full multiplexing capability when the number of potential targets is very much larger than  $N_{\text{fibre}}$ , this will not be the case as the choice of targets becomes more constrained, whether through having a lower  $N_{\text{obj}}$ , or by requiring high  $s$ , or through other constraints such as a possible need to locate ‘sky’ fibres to allow nodding. Given such constraints, not all fibres (slits) that are, in principle, available will actually be usable in practice. The efficiency will therefore drop. However, it is important to appreciate that the  $\epsilon$  in equation (19) is the efficiency *averaged* over all the fibre configurations that are needed to complete this field.

Equation (19) is quite flexible and quite general. As a simple example, it might be decided that 10 per cent of the target objects required ten times the exposure time. In practice, this means that 47 per cent of the fibres should be allocated at any point in time to the majority objects needing only the shorter exposure time and 53 per cent to the more difficult minority objects. One could then consider this in equation (19) in one of two ways. One could either increase the average exposure time  $\bar{T}_{\text{exps}}$  by a factor of 1.9, while keeping  $N_{\text{obj}}$  and  $N_{\text{fibre}}$  the same. Alternatively, one could conceptually ‘forget’ about these more difficult objects, keep  $\bar{T}_{\text{exps}}$  the same, reduce  $N_{\text{obj}}$  by 10 per cent, but then also reduce the  $N_{\text{fibre}}$  to 0.47 of its original value. Both approaches give exactly the same total observation time  $T_{\text{obs}}$ .

All of the complications arising from practical limitations of fibre placement, etc., are put into the final efficiency term,  $\epsilon(s, N_{\text{obj}})$ . This has the advantage that it may be calculated with whatever precision or sophistication is desired, as the performance of an instrument, including the fibre positioning etc., is further refined. The breakdown of  $T_{\text{obs}}$  in equation (19) is designed to separate the largely instrument-independent terms from the detailed instrument-specific technical constraints in the last term.

## 6.2 Generic survey design considerations in $s/M_{\text{stellar}}$ -space

In this section, we discuss the first three factors of equation (19) that are largely instrument-independent apart from the ‘ideal’ multiplexing factor  $N_{\text{fibre}}$ . To first order, the average observation time per object  $\bar{T}_{\text{exps}}$  might be expected to vary as  $M_{\text{stellar}}^{-2}$  at a given redshift. This is certainly true for passive objects. For star-forming galaxies, the combination of the weakly declining  $sSFR(M_{\text{stellar}})$  relation and



**Figure 34.** The product ( $s N_{\text{obj}}(M_{\text{stellar}})$ ) indicates the (log) number of targets that need to be observed in a MOONS FoV to achieve the sampling rate  $s$  of the tracer population down to a limiting stellar mass  $M_{\text{stellar}}$ . The red curves, representing the values 2.7, 3.25, and 3.65, indicate the lines of constant number of tracers superposed on numerous other figures in this paper. The equal number density of tracers is of course equivalent to the constant number of targets (up to the appropriate spatial density factor).

the increasing effects of dust obscuration at higher stellar masses means that the luminosity of the emission lines on which a successful redshift determination depends may have a weaker dependence on mass (a point emphasized to us by Emanuele Daddi, private communication). For illustrative purposes, we assume a scaling as  $M_{\text{stellar}}^{-2}$ .

The latest estimates for the MOONS suggest that an integration time of  $2h$  is needed for a star-forming object with a stellar mass of  $10^{9.5} M_{\odot}$ . Hence, we approximate

$$\bar{T}_{\text{exps}} \approx 2h \times \left( \frac{10^{9.5} M_{\odot}}{M_{\text{stellar}}} \right)^2. \quad (20)$$

Apart from this potential penalty, decreasing the limiting mass also has a large (beneficial) effect in increasing the number density of the galaxy tracer population, but this also means that the number of objects to be observed in each field also increases proportionally in the second term in equation (19). Fig. 34 shows the number of objects that need to be observed in one MOONS FoV in order to reach a certain sampling rate  $s$  at a given stellar mass cut  $M_{\text{stellar}}$ , using the numbers from the L-Galaxies SAM light-cone, and summing over the three redshift ranges considered earlier in this paper. The red lines in the figure represent lines of constant number of galaxies observed, and provide the lines that were superposed on the other figures in this paper.

The third term in equation (19) concerns the theoretical multiplexing capacity of the instrument, i.e. how many targets can be simultaneously observed in a given (fibre or slit) configuration. This may depend on the choice between available observing modes. For instance, for a fibre-spectrograph like MOONS, a key decision is how to obtain spectra of blank sky for sky-subtraction of the target spectra. For MOONS, three observing modes have been envisaged.



XSwitch is the most conservative and assigns a half of the fibres to target sources, paired with an equal number of fibres placed on blank regions of sky that are displaced by a fixed distance and direction from each target, enabling an ABBA nodding technique. This extra constraint on the location of the sky fibres will inevitably degrade the fourth efficiency term in equation (19), but we should simply use the maximum number of allocatable fibres in the third term in equation (19). STARE700 allocates 700 fibres to target sources, with the remainder placed on random blank sky regions while STARE900 puts up to 900 fibres simultaneously on targets.

### 6.3 The efficiency factor $\epsilon$ of MOONS in $s/M_{\text{stellar}}$ -space

We now turn to the instrument-specific fourth term in equation (19). As stated above, this term represents the average efficiency of fibre placement when averaged over all of the configurations required to complete the observations in a particular region of sky. We will here apply several simplifying assumptions in this analysis. A full treatment would require detailed simulations of the positioning of fibres on real targets, but these could be carried out as desired with arbitrary realism. We are primarily interested in understanding relative effects.

We will therefore assume the following for simplicity:

(i) We will consider only the fibre placement for the ‘normal’ shorter integration objects and include the effects of any objects requiring longer exposure times (such as passive galaxies or additional rare targets of special interest, such as very high redshift galaxies) by reducing  $N_{\text{fibre}}$ , as in the example given previously in Section 6.2.


(ii) We further assume that all of these normal objects are observed for the same integration time, defined for simplicity by the nominal integration time required for the limiting stellar mass.

(iii) For each successive fibre configuration in a given survey field, we assume that the fraction of ‘available’ fibres that can be successfully placed on a target is a function only of the total number of targets remaining to be observed and the number of available fibres. For each observing mode, we can estimate this function from the performance of the fibre positioner (see below).

(iv) We might expect that the chance that a given target galaxy is observed depends on the surface density of all targets averaged within some solid angle of that galaxy. For example, this solid angle could be set by the individual fibre patrol field. In this paper, we completely neglect any variation of the projected surface density of tracers  $dN/d\Omega$  across the sky that may have an impact on the (local) sampling rate  $s$  via fibre-positioning issues. At these depths the variation in projected surface density is anyway modest. Furthermore, any correlation of the projected surface density with the richness of individual groups is washed out by the high number of foreground and background objects. This is shown in Figs 35 and 36 for three different solid angles that might be relevant for different instruments. Effectively, we have assumed in this paper that the random selection of targets (and hence the local sampling rate  $s$ ) is independent of the local surface number density of tracers. Of course, this issue will be addressed in the final survey design of MOONRISE, which will use actual target lists and the final fibre allocation software, which is not yet available.

As stressed earlier, any or all of these assumptions can be avoided by undertaking completely realistic simulations of the fibre positioner on real catalogues of targets. Our purpose here is only to explore the most general points.

The MOONS instrument is equipped with 1000 science fibres that can in principle be allocated to either a target object or to a blank sky position. In the XSwitch observing mode, the current implementation



**Figure 35.** The distribution of the (local) surface number density of tracers  $dN/d\Omega$  for galaxies in different richness structures. The surface number density for each galaxy is calculated from counting the number of neighboring galaxies within a radius of 1 arcmin.

of the fibre-positioning Faststar algorithm allocates pairs of fibres (one fibre on target, one on the nearby sky) whenever possible, so the best case scenario would be to have 500 fibres on targets, 500 on the sky. In the presence of passive objects and high- $z$  targets, requiring longer integrations, the number of fibres available for normal (star-forming) targets is likely to reduce in practice to of order  $N_{\text{fibre}} \approx 300$ . Furthermore, imposing the constraint that the remaining ‘sky’ fibres must be placed with a fixed displacement from the target sources (to allow nodding), places large constraints the choice of targets. Given all the constraints (including the requirement of a constant offset to a blank region of sky), a general feature of such algorithms is that full fibre allocation is only achieved when the number of available (remaining) targets greatly exceeds the number of available fibres. When the ratio of remaining targets to fibres reduces, the fraction of fibres that are usable in a given fibre configuration drops considerably. While further optimization of the fibre positioning algorithm are no doubt likely, this drop off in efficiency as the number of targets is reduced is likely to remain a general feature of MOONS and similar instruments.


Hence, in good first-order approximation, the number of allocatable science fibres in each pointing to one instrument field of view depends only the number of targets available for selection, i.e. still not observed, and the number of available fibres. More precisely, the fraction  $f = N_{\text{fibre, all}}/N_{\text{fibre}}$  of allocatable fibres are in good approximation a function of the ratio of the number of targets available and the number of available fibres,

$$f = f(N_{\text{tar}}/N_{\text{fibre}}). \quad (21)$$

The function  $f$  we used in this work is based on the current (May 2020) Faststar algorithm for the MOONS instrument in the XSwitch mode that places a sky fibre at a fixed offset from each target fibre. As mentioned above, about  $N_{\text{fibre}} \approx 300$  science fibres are available for normal objects in this instrument mode.

For a given number density of tracers  $N_{\text{obj}}$  and a desired final sampling rate  $s$ , we can then progressively use the fraction  $f$  of available fibres that are allocatable in the next pointing as a function of the ‘remaining targets’ to compute the number of configurations  $N_{\text{config}}$  that are required to finally reach the desired  $s$  for a particular  $N_{\text{obj}}$ . It is then straightforward to average the  $f$  of these successive configurations to yield the efficiency factor  $\epsilon^{-1}$  needed in equation (19).





**Figure 36.** Richness  $R$  of groups versus surface number density of tracers  $dN/d\Omega$ , plotted on a galaxy basis. Each dot represents a galaxy within one of the three MOONS redshift ranges with  $\log_{10}(M_{\text{stellar}}/M_{\odot}) > 9.5$  and shows its true richness (same in all three panels) and the projected number density within three different radii (shown in the three panels). In each panel, the black dots with error bars show the median  $dN/d\Omega$  within a given richness bin, with the 16 and 84 per cent percentiles. The red and green dashed lines show the same for the overall sample.

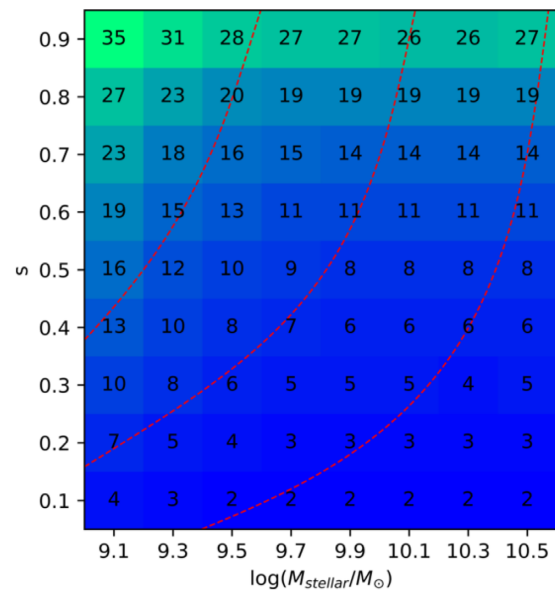
Using for the sake of illustration the particular  $f$  for the latest Faststar simulation, we show  $N_{\text{config}}$  as a function of  $M_{\text{stellar}}$  and  $s$  in Fig. 37. Full sampling,  $s = 1.0$  is not considered because the efficiency in allocating fibres at very low target densities is the most uncertain and will be most dependent on the detailed spatial distribution of objects. Hence, we do not give an estimate for the costs of reaching  $s = 1.0$  in Fig. 37.

It should be noted that for typical survey designs in the middle of the ranges considered,  $N_{\text{config}}$  is large, of order 10 configurations or more. This makes it practical to repeatedly allocate some fibres to those objects requiring much longer exposure times, while giving ‘typical’ objects many opportunities to be allocated a fibre, ensuring that constraints arising from the individual ‘patrol-fields’ of individual fibres (which will be present in each individual configuration) are washed out in the aggregate sampling map.

At fixed  $M_{\text{stellar}}$ , i.e. fixed  $N_{\text{obj}}$ , the required number of configurations at first increases linearly with  $s$  but then increases more steeply as more and more configurations are required in which the number of remaining targets falls below the optimum (large) multiple of the number of fibres.

A general point is that higher  $s$  is more efficiently attained with a high target density of tracers  $N_{\text{obj}}$ . This is because the fraction of available targets to fibres,  $N_{\text{tar}}/N_{\text{fibre}}$  stays high for longer as  $s$  increases when the initial  $N_{\text{obj}}$  is larger. For the particular example of the current implementation of Faststar shown in Fig. 37, we show the efficiency  $\epsilon$  across  $s/M_{\text{stellar}}$ -space in Fig. 38. When going deeper, i.e. starting with a higher number density of targets,  $\epsilon$  is overall high and drops only slightly with  $s$  until a very high  $s$  is desired. This is because many configurations are needed in the high-density target-rich regime. Hence, the allocation of fibres stays efficient to quite high sampling rates. In contrast, for a lower  $N_{\text{obj}}$ , i.e. a higher stellar mass cut, the efficiency is overall lower and furthermore drops steeply with sampling rate at relatively low  $s$ . The bottom line is that higher  $s$  is overall much more expensive for a lower density of targets  $N_{\text{obj}}$  than for a high density of targets that require a large number of configurations per field.

The total observation time  $T_{\text{obs}}$  that is required to complete the observations in one instrument FoV can then be computed as a function of the sampling rate  $s$  and the stellar mass cut  $M_{\text{stellar}}$  by multiplying all of the terms in equation (19). This is presented in Fig. 39. It can be seen that  $M_{\text{stellar}}$ , both through the (assumed) exposure time factor,  $M_{\text{stellar}}^{-2}$ , and through the number of targets  $N_{\text{obj}}$  being proportional to  $M_{\text{stellar}}^{-3/4}$ , is the dominant quantity in determining




**Figure 37.** The required number of configurations required to complete the observations in a given field,  $N_{\text{config}}$ , as a function of  $s$  and  $M_{\text{stellar}}$ .

the total observing time  $T_{\text{obs}}$  required to complete the observation in one MOONS FoV, and thus, in a given amount of observing time, the total number of such FoV that can be surveyed.


## 7 SUMMARY AND CONCLUSION

In this paper, we have presented an in-depth study of the performance of group-finders, based on positional information of individual tracer galaxies, when they are applied to redshift surveys, especially at high redshift. This has been based on comparing the recovered group catalogues (including singleton galaxies) that are obtained by applying the group finding algorithms to realistic mock light-cones with the underlying ‘reality’ from the L-galaxies SAM, in which the full halo information is known.

We have examined how the group-finding performance depends on the main design parameters of the survey, in particular on the limiting stellar mass (or to a certain degree equivalently, to the number density) of the galaxy tracers, and on the (assumed) random sampling rate of those tracers. We also examined how the observational cost



**Figure 38.** The average fibre allocation efficiency  $\epsilon$  as a function of the sampling rate  $s$  and stellar mass cut  $M_{\text{stellar}}$ .



**Figure 39.** The total observation time  $T_{\text{obs}}$  (in log hours) required to complete the observation in one MOONS FoV as a function of sampling rate  $s$  and stellar mass cut  $M_{\text{stellar}}$ .

of such a survey depends on those same parameters, choosing the future MOONS spectrograph on the VLT as an example. This study should therefore be of use in designing future surveys with this and other instruments.

We began with a comparison of two different group finding approaches, which we call the ‘halo-based’ and the FoF group finding approaches, at low redshift, i.e. to SDSS-like surveys. This led to the conclusion that many of the performance and science metrics are almost indistinguishable between the group catalogues recovered by the two methods. Overall, the ‘halo-based’ approach works slightly better at high  $s$ , while the FoF algorithm works slightly better at low  $s$ . This can for example be seen in the most basic statistical quantities of purity  $p_2$  and completeness  $c_2$ . That the ‘halo-based’

approach performs better at high  $s$  might be expected, as it is a more ‘refined’ approach in the sense that it assumes more information (about the sizes of dark matter haloes). The strength of FoF, on the other hand, lies in its minimal dependence on assumptions and it seems to be more robust as the fraction of objects for which positional information is known drops away. However, the ‘halo-based’ and the FoF group-finders perform overall very similarly, and in the remainder of this paper, we considered for simplicity only the FoF approach.

A total of 12 science metrics were then constructed that were designed to capture a range of physical information that might be relevant for a wide variety of science investigations. We then examined how these different science metrics varied across  $(s/M_{\text{stellar}})$ -space when the FoF algorithm is applied at high redshift and identified the following main conclusions:

(i) The number of groups  $N_{\text{grs}}$  scales closely as the number of tracers observed  $N_{\text{tr}}$ , given by the product of the sampling rate  $s$  and the density of the underlying tracer population  $N_{\text{obj}}$ , which is given in our analysis by the limiting stellar mass of the tracers  $M_{\text{stellar}}$ . There is a weak preference for  $s$  over  $M_{\text{stellar}}$  (or  $N_{\text{obj}}$ ) to achieve a given  $N_{\text{tr}}$ .

(ii) As would be expected, the median dark matter halo mass  $M_{\text{H},m}$  for singleton galaxies scales as  $M_{\text{stellar}}^{0.5}$  and is more or less independent of  $s$ . On the other hand, the median  $M_{\text{H},m}$  for groups (containing more than one galaxy) scales quite closely with the number density of tracers  $N_{\text{tr}}$ .

(iii) The scatter in the estimated halo masses of recovered groups,  $\sigma(M_{\text{H}})$  reduces with  $s$ , more or less independent of  $M_{\text{stellar}}$ , regardless of whether this scatter is calculated on what we call a galaxy-basis (considering the difference between the estimated parent halo mass and the true parent halo mass for each individual galaxy in the sample) or on what we called a group-basis, in which we compare the estimated and true masses of those recovered groups that are (uniquely) ‘two-way matched’ with real groups.

(iv) The accuracy in correctly classifying galaxies as centrals or satellites is different for centrals and for satellites. For centrals there is a dependency on  $s$ , especially for richer groups. For singletons, the accuracy is high and depends only weakly on  $s$ . The accuracy of classifying satellites is very largely independent of  $s$  (and also  $M_{\text{stellar}}$ ). This is because the vast majority of satellites are always seen as satellites. Nevertheless, this classification accuracy is only about 80 per cent, because of infidelities in the group reconstruction.

(v) Looking at how these results change with redshift, there is rather little change in almost all the metrics (at a given  $(M_{\text{stellar}}, s)$ ) over the redshift range  $0.9 < z < 2.6$ . The main effects are all related to the unavoidable reduction in the multiplicity with redshift, which is ultimately due to the cosmic growth of structure. This reduction in multiplicity actually has some nominally beneficial effects, most notably an increase in the accuracy of central classification.

We then constructed a quite general and flexible formalism for considering the required cost in the observation time to complete the survey within each instrumental field of view, which translates (inversely) to the number of such fields that can be observed in a given amount of observing time. This formalism puts all of the complexities of instrument operation, such as constraints on fibre positioning, into a single efficiency term,  $\epsilon$ . Our analysis leads to the following main conclusions:

(i) The choice of the mass limit  $M_{\text{stellar}}$  has a strong effect: Not only does this likely increase the typical observation time needed for each configuration (as  $M_{\text{stellar}}^{-2}$  for continuum objects, likely less for emission line objects), it also increases the number of targets that

must be observed in each field so as to reach a given  $s$ , increasing as roughly  $M_{\text{stellar}}^{-3/4}$ .

(ii) Having a high density of target galaxies (e.g. by having a low  $M_{\text{stellar}}$ ) however produces high efficiency. This is because the fraction of available fibres that can successfully be placed on a target is a strong function of the ratio of available targets to available fibres. A survey at a low  $N_{\text{obj}}$  (i.e. high  $M_{\text{stellar}}$ ) is inefficient due to the low number of targets that are available in this regime. For the same reason, there is a cost associated with pursuing high  $s$ , although the decrease in average efficiency  $\epsilon$  with  $s$  is modest at high  $N_{\text{obj}}$ , because more time is spent with high target density, but is much larger when one starts with a low  $N_{\text{obj}}$ .

## ACKNOWLEDGEMENTS

We thank the anonymous referee for their highly insightful and constructive comments and suggestions, which have greatly benefited this paper. This work is based on an MSc thesis that was presented by T.J.L. at the ETH Zürich. R.M. acknowledges support from the ERC Advanced Grant 695671 ‘QUENCH’ and by the Science and Technology Facilities Council (STFC).

## DATA AVAILABILITY

The L-Galaxies light-cone data used in this paper are available on the L-Galaxies Munich Galaxy Formation Model website: <http://galformod.mpa-garching.mpg.de/public/LGalaxies/>.

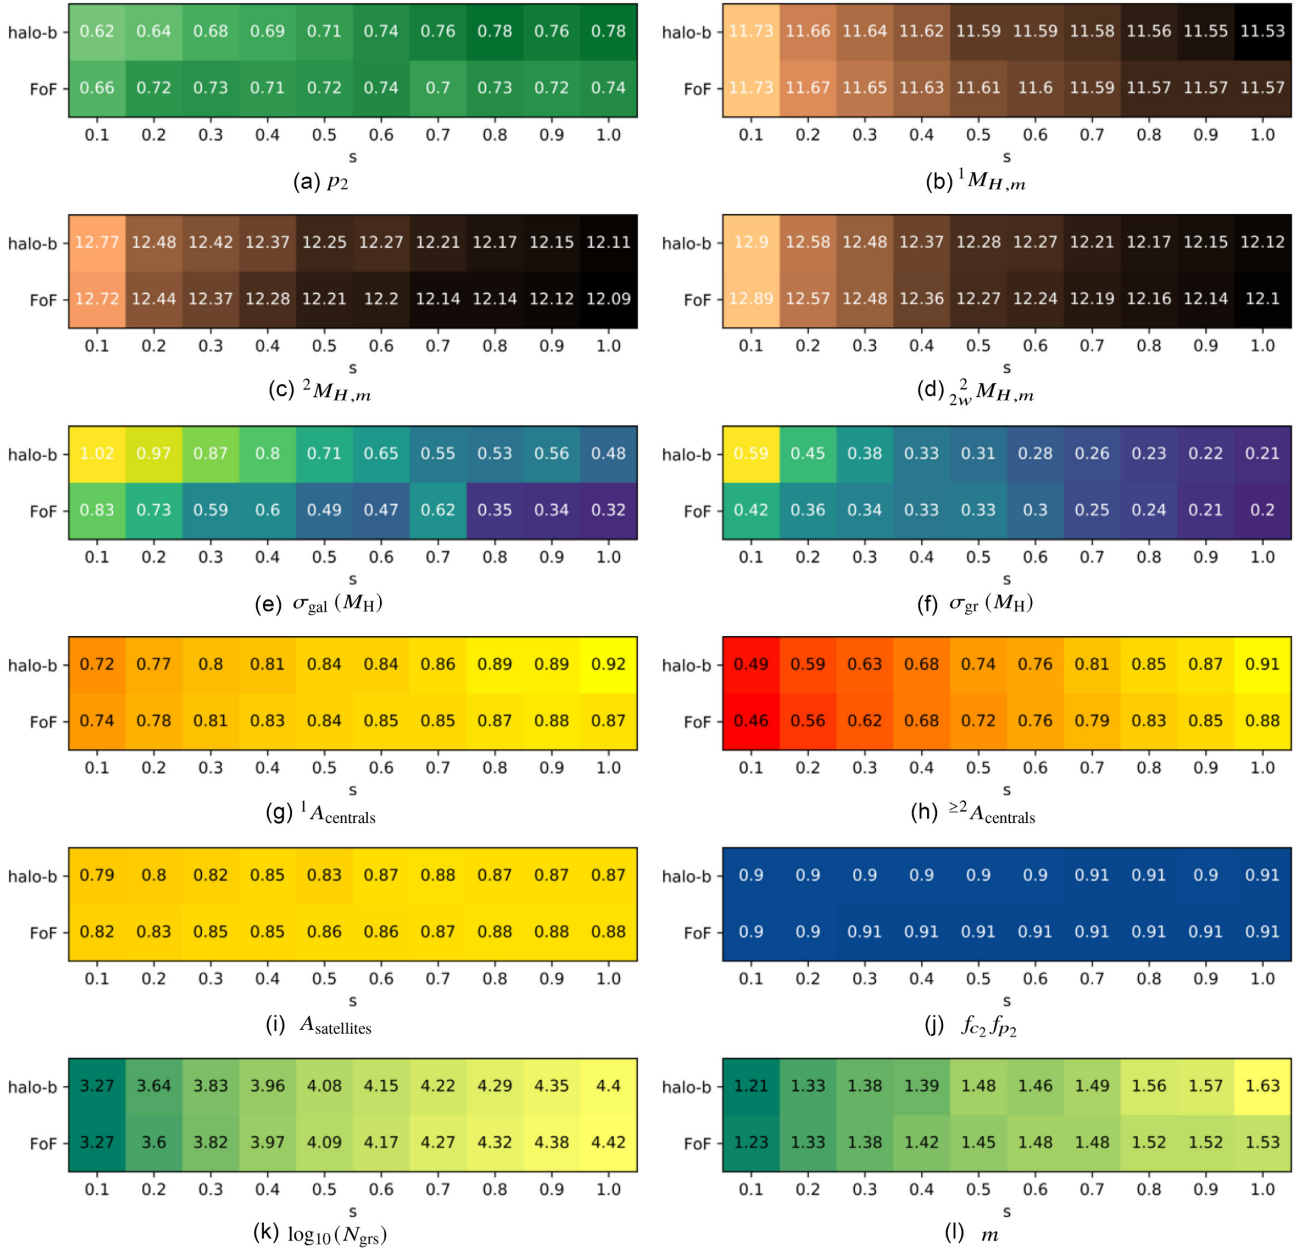
## REFERENCES

- Alonso D., Eardley E., Peacock J. A., 2015, *MNRAS*, 447, 2683  
 Berlind A. A. et al., 2006, *ApJS*, 167, 1  
 Boylan-Kolchin M., Ma C.-P., Quataert E., 2008, *MNRAS*, 383, 93  
 Broadhurst T. J., Ellis R. S., Shanks T., 1988, *MNRAS*, 235, 827  
 Cirasuolo M., et al., 2020, *Msngr*, 180, 10  
 Colless M. et al., 2001, *MNRAS*, 328, 1039  
 Colless M., Ellis R. S., Taylor K., Hook R. N., 1990, *MNRAS*, 244, 408  
 Davis M. et al., 2003, in Guhathakurta P., ed., *SPIE Conf. Ser. Vol. 4834, Discoveries and Research Prospects from 6- to 10-Meter-Class Telescopes II*. SPIE, Bellingham, p. 161  
 Eke V. R. et al., 2004, *MNRAS*, 348, 866

- Forbes J., Krumholz M., Burkert A., Dekel A., 2013, *MNRAS*, 438, 1552  
 Gerke B. F. et al., 2005, *ApJ*, 625, 6  
 Gunn J. E. et al., 2006, *AJ*, 131, 2332  
 Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R., Guo Q., Lemson G., Springel V., Overzier R., 2015, *MNRAS*, 451, 2663  
 Huchra J. P., Geller M. J., 1982, *ApJ*, 257, 423  
 Kauffmann G., Li C., Zhang W., Weinmann S., 2013, *MNRAS*, 430, 1447  
 Knobel C. et al., 2009, *ApJ*, 697, 1842  
 Knobel C., Lilly S. J., Woo J., Kovač K., 2015, *ApJ*, 800, 24  
 Kovac K., Lilly S., Porciani C., Cucciati O., Tasca L., Bolzonella M., Knobel C., Iovino A., 2009, *AAS Meeting Abstracts #213*, p. 323.03  
 Le Fèvre O. et al., 2005, *A&A*, 439, 845  
 Lilly S. J. et al., 2009, *ApJS*, 184, 218  
 Lilly S. J., Le Fèvre O., Crampton D., Hammer F., Tresse L., 1995, *ApJ*, 455, 50  
 Maiolino R. et al., 2020, *Msngr*, 180, 24  
 Marinoni C., Davis M., Newman J. A., Coil A. L., 2002, *ApJ*, 580, 122  
 Merchán M., Zandivarez A., 2002, *MNRAS*, 335, 216  
 Merchán M. E., Zandivarez A., 2005, *ApJ*, 630, 759  
 Peng Y.-J. et al., 2010, *ApJ*, 721, 193  
 Peng Y.-J., Lilly S. J., Renzini A., Carollo M., 2012, *ApJ*, 757, 4  
 Robotham A. S. G. et al., 2011, *MNRAS*, 416, 2640  
 Sin L. P. T., Lilly S. J., Henriques B. M. B., 2017, *MNRAS*, 471, 1192  
 Springel V. et al., 2005, *Nature*, 435, 629  
 Tinker J., Wetzel A., Conroy C., 2011, preprint ([arXiv:1107.5046](https://arxiv.org/abs/1107.5046))  
 van den Bosch F. C., Aquino D., Yang X., Mo H. J., Pasquali A., McIntosh D. H., Weinmann S. M., Kang X., 2008, *MNRAS*, 387, 79  
 Wang K., Mo H. J., Li C., Meng J., Chen Y., 2020, *MNRAS*, 499, 89  
 Wetzel A. R., Tinker J. L., Conroy C., 2012, *MNRAS*, 424, 232  
 Yang X., Mo H. J., van den Bosch F. C., Pasquali A., Li C., Barden M., 2007, *ApJ*, 671, 153

## APPENDIX A: THE SCIENCE METRICS FOR ‘HALO-BASED’ AND FOF METHODS

The science metrics designed to capture a wide range of environmental information introduced in Section 3 are here presented for the ‘halo-based’ and FoF group finding algorithms in the  $0.02 < z < 0.17$  redshift range, taking a constant stellar mass cut of  $\log_{10}(M_{\text{stellar}}/M_{\odot}) = 9.0$  but exploring the full sampling rate range  $s = 0.1-1.0$ . As truth, we use light-cones produced by the L-Galaxies SAM.



**Figure A1.** The science metrics designed to capture a wide range of environmental information for the ‘halo-based’ and FoF group catalogues at low redshift in SDSS.

## APPENDIX B: HOW THE SAMPLING RATE AFFECTS THE ACCURACY OF CENTRALS

We statistically analyse the effect of the sampling rate  $s$  on the success in recovering centrals, i.e.  $A_{\text{centrals}}$ . Recovering centrals incorrectly, i.e. claiming that a galaxy is a central in the recovered group catalogue, even though it is a satellite in truth, can have two causes: This can happen (a) due to a failure of the group finding algorithm (fragmentation of the real group), making a central of a false group, or (b) due to the incomplete sampling of galaxies in the recovered group catalogue, the true central was not, in fact, observed and so a satellite was wrongly identified as the most massive group member.

We present here a statistical analysis of case (b). For simplicity, we first consider a 50 per cent sampling rate: When applying a 50 per cent sampling on an underlying true group catalogue, half

of the true singletons in the true group catalogue survive. Hence, obviously, there are singletons after the sampling that are truly singletons. For a group catalogue dominated by small richness structures – as it is expected to be the case – even a dominant fraction of them are. Hence,  $A_{\text{centrals}}$  of singletons is expected to be close to 1.0 (in fact this is the case for wide range of sampling rates, even to quite low ones, due to the same reasoning). Let us next consider centrals of groups with at least two members. When applying a 50 per cent sampling on a true two member group, there are four possible outcomes (which all have the same probability): (1) Both galaxies are sampled out (not observed), (2) the central or (3) the satellite is sampled out, or (4) both survive the sampling (and are observed). If small richness structures are dominant, most of the recovered two member groups are truly two member groups and hence must (case 4) have a correct detection of the central (unless



there is a failure of the group-finder). Considering true groups of three members, there are in total eight possible outcomes, four of them are groups of at least two members, and in three of them the central is not sampled out. Hence, as long as there is no failure of the group-finder, the central is identified correctly three out of four times. Generalizing this to groups of richness  $N$ , there are  $\sum_{i=0}^{N-2} \binom{N}{i}$  possible outcomes that exhibit a group of at least two members after the sampling, out of which in  $\sum_{i=0}^{N-2} \binom{N-1}{i}$  cases, the central is not sampled out. The ratio of these two sums goes to the sampling rate 0.5 with increasing  $N$ , as it obviously has to. In summary, if the group catalogue is dominated by small richness structures, the success of recovering centrals  $A_{\text{centrals}}$  is expected to lie well above the sampling rate of 50 per cent.

When generalizing this statistical analysis to any sampling rate  $s$ , we have to take into account that the possible outcomes of the sampling do not happen with the same probability any more. For any sampling rate  $s$  the recovered two member groups that are truly two member groups (which is the dominant fraction down to low sampling rate) have – as long as there is no other failure of the group-finder – a correct central identification. For the true three member

groups, the four possible and three successful outcomes (with a group of at least two members) have to be weighted by their relative probabilities; For example when applying a sampling of 80 per cent, in 86 per cent of the cases the central is correctly identified.

Generalizing this to any  $N$  and  $s$ , the fraction of correctly recovered centrals (when the sampled group has at least two members) is

$$\frac{\sum_{i=0}^{N-2} \binom{N-1}{i} s^{N-i} (1-s)^i}{\sum_{i=0}^{N-2} \binom{N}{i} s^{N-i} (1-s)^i}. \quad (\text{B1})$$

This fraction obviously converges to the sampling rate  $s$  with increasing  $N$ . In conclusion, the accuracy of recovered centrals  $A_{\text{centrals}}$  (for singletons and richer structures) is expected to be always higher than the sampling rate  $s$ , provided (a) that the performance of the group-finder is otherwise good enough and (b) that the underlying true group catalogue is dominated by small richness structures, as it will be in any hierarchical universe such as our own.

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