

The completed SDSS-IV extended baryon oscillation spectroscopic survey: pairwise-inverse probability and angular correction for fibre collisions in clustering measurements

Faizan G. Mohammad,^{1,2★} Will J. Percival,^{1,2,3} Hee-Jong Seo,^{1,2,3} Michael J. Chapman,^{1,2} D. Bianchi,⁵ Ashley J. Ross,⁶ Cheng Zhao,⁷ Dustin Lang,³ Julian Bautista,⁸ Jonathan Brinkmann,⁹ Joel R. Brownstein,¹⁰ Etienne Burtin,¹¹ Chia-Hsun Chuang,¹² Kyle S. Dawson,¹⁰ Sylvain de la Torre,¹³ Arnaud de Mattia,¹¹ Sarah Eftekharzadeh,¹⁰ Sebastien Fromenteau,¹⁴ Héctor Gil-Marín,^{5,15} Jiamin Hou,¹⁶ Eva-Maria Mueller,^{17,8} Richard Neveux,¹¹ Romain Paviot,¹³ Anand Raichoor,⁷ Graziano Rossi,¹⁸ Donald P. Schneider,^{19,20} Amélie Tamone,⁷ Jeremy L. Tinker,²¹ Rita Tojeiro,²² Mariana Vargas Magaña²³ and Gong-Bo Zhao^{24,25}

Affiliations are listed at the end of the paper

Accepted 2020 July 20. Received 2020 July 17; in original form 2020 May 26

ABSTRACT

The completed extended Baryon Oscillation Spectroscopic Survey (eBOSS) catalogues contain redshifts of 344 080 quasars at $0.8 < z < 2.2$, 174 816 luminous red galaxies between $0.6 < z < 1.0$, and 173 736 emission-line galaxies over $0.6 < z < 1.1$ in order to constrain the expansion history of the Universe and the growth rate of structure through clustering measurements. Mechanical limitations of the fibre-fed spectrograph on the Sloan telescope prevent two fibres being placed closer than 62 arcsec in a single pass of the instrument. These ‘fibre collisions’ strongly correlate with the intrinsic clustering of targets and can bias measurements of the two-point correlation function resulting in a systematic error on the inferred values of the cosmological parameters. We combine the new techniques of pairwise-inverse probability and the angular upweighting (PIP+ANG) to correct the clustering measurements for the effect of fibre collisions. Using mock catalogues, we show that our corrections provide unbiased measurements, within data precision, of both the projected $w_p(r_p)$ and the redshift-space multipole $\xi^{(\ell=0,2,4)}(s)$ correlation functions down to $0.1 h^{-1}\text{Mpc}$, regardless of the tracer type. We apply the corrections to the eBOSS DR16 catalogues. We find that, on scales $s \gtrsim 20 h^{-1}\text{Mpc}$ for ξ^ℓ , as used to make baryon acoustic oscillation and large-scale redshift-space distortion measurements, approximate methods such as nearest-neighbour upweighting are sufficiently accurate given the statistical errors of the data. Using the PIP method, for the first time for a spectroscopic program of the Sloan Digital Sky Survey, we are able to successfully access the one-halo term in the clustering measurements down to $\sim 0.1 h^{-1}\text{Mpc}$ scales. Our results will therefore allow studies that use the small-scale clustering to strengthen the constraints on both cosmological parameters and the halo occupation distribution models.

Key words: galaxies: distances and redshifts – large-scale structure of Universe – cosmology: observations.

1 INTRODUCTION

The 3D distribution of galaxies in the Universe contains a wealth of information about its composition and dynamical evolution. One of the most efficient ways to access this information is provided by the spectroscopic galaxy redshift surveys. Spectroscopic redshift surveys over the past two decades (York et al. 2000; Drinkwater et al. 2010; Eisenstein et al. 2011; Guzzo et al. 2014; Blanton et al. 2017) have thus played a crucial role in both constraining the standard cosmological model (e.g. Percival et al. 2001; Cole et al. 2005; Eisenstein et al. 2005; Blake et al. 2012; Alam et al. 2017) and probing the galaxy formation and evolution at different

cosmic epochs (e.g. Lewis et al. 2002; Chen et al. 2013; Krywult et al. 2017).

The advent of multi-object spectrographs (MOS) has allowed an unprecedented increase in the surveyed volume and catalogue size of spectroscopic redshift surveys allowing us to measure the cosmological observables with increasing precision. This requires an equivalent effort to correct for systematic issues that can potentially bias the measurements of the observables degrading the accuracy on estimates of physical parameters (Ross et al. 2012; de la Torre et al. 2013). In spectroscopic redshift surveys, based on multislit or multifibre spectrographs, one of these effects is caused by missing observations that result from the limitations of the hardware set-up. In particular, the finite size of fibres (or slits) prevents any pair of targets closer than the fibre diameter to be observed simultaneously, usually referred to as the *fibre collision* or *close pair* problem

★ E-mail: faizan37@gmail.com

(Jing, Mo & Börner 1998). The effect is strongly correlated with the intrinsic clustering of targets as collisions occur in regions of high target density. The resulting systematic offset in the clustering measurements such as the two-point correlation function (2PCF) can mimic the effect of physical parameters such as galaxy bias b or growth rate of structure $f\sigma_8$ introducing a systematic bias in the inferred values of these parameters (Pezzotta et al. 2017). Different techniques have been proposed in the past (see e.g. Guo, Zehavi & Zheng 2012; Reid et al. 2014; Hahn et al. 2017; Zarrouk et al. 2018; Yang et al. 2019; Sunayama et al. 2020) to mitigate for missing observations in clustering measurements. Spectroscopic programs of the Sloan Digital Sky Survey have often previously used variations of the ‘nearest-neighbour’ (NN) technique to correct for the fibre collision problem. Here, either the redshift and classification of the nearest observed target is assigned to an unobserved one (Zehavi et al. 2002, 2005) or the nearest observed target is given the additional weight of the unobserved target. The Baryon Oscillation Spectroscopic Survey (BOSS) part of SDSS-III analyses (Reid et al. 2016) adopted the latter variant, where the weight of an unobserved target was transferred to the closest observed target. The Vimos Public Extragalactic Redshift Survey (VIPERS) employed the ‘target sampling rate’ (TSR) to quantify the selection probabilities. TSR is defined as the ratio between the number of observed and input targets in a rectangular region around an observed target where the size of the region was calibrated using realistic survey mock samples (de la Torre et al. 2013). The common factor of these techniques is that they do not correctly allow for the correlation between the observed and unobserved targets. For example, the upweighting variant of the NN method ignores pairs between observed and unobserved targets, while the redshift assignment variant assumes that these pairs are all transverse to the line of sight.

In order to correctly allow for the correlation between the observed and unobserved targets, Bianchi & Percival (2017) proposed a ‘pairwise-inverse probability’ (PIP) weighting scheme. The technique calculates selection probabilities for each observed pair by generating multiple random survey realizations statistically equivalent to the actual observations, and then upweighting each observed pair of targets by the reciprocal of this probability. The method is statistically unbiased under the assumption that no pair of targets has zero probability of being observed. Percival & Bianchi (2017) extended the PIP technique under the assumption that the upweighted pairs of a particular separation are statistically equivalent to those in the full sample. In this case, the PIP upweighted angular pair counts can simply be scaled to match those in the full target sample. For Sloan telescope observations, this allows the full clustering signal to be recovered on all pair separations even if some regions are only covered by one observation. These techniques have been successfully tested for the multifibre spectroscopy adopted by the Dark Energy Spectroscopic Instrument (DESI) using mock samples (Bianchi et al. 2018; Smith et al. 2019). Mohammad et al. (2018) applied the PIP corrections to the final VIPERS data set based on observations using the multislit visible multi-object spectrograph (VIMOS; Le Fèvre et al. 2003).

In this paper, we apply the correction schemes presented in Bianchi & Percival (2017) and Percival & Bianchi (2017) to correct for missing observations affecting clustering measurements from the extended Baryon Oscillation Spectroscopic Survey (eBOSS) data (Dawson et al. 2016), one of the components of the SDSS-IV project (Blanton et al. 2017). eBOSS targeted three different samples: emission-line galaxies (ELGs) between $0.6 < z < 1.1$ over 1170 deg^2 , luminous red galaxies (LRGs) between $0.6 < z < 1.0$ over 4242 deg^2 , and quasars (QSOs) split into ‘clustering’ QSOs between $0.8 < z$

< 2.2 over 4808 deg^2 and high-redshift QSOs for Lyman α forest analyses. A number of studies have used clustering measurements from the early eBOSS data at relatively small scales (see e.g. Laurent et al. 2017; Guo et al. 2019; Alam et al. 2020a). However, none of these works have extended the analysis below $\sim 1 h^{-1} \text{ Mpc}$, where fibre collisions strongly affect the measured correlation function.

This study is part of a coordinated release of papers based on the final eBOSS catalogues. These include final eBOSS measurements of the baryon acoustic oscillations (BAOs) and redshift-space distortions in the clustering of LRGs ($0.6 < z < 1.0$; Bautista et al. 2020; Gil-Marín et al. 2020), ELGs ($0.6 < z < 1.1$; de Mattia et al. 2020; Raichoor et al. 2020; Tamone et al. 2020), and QSOs ($0.8 < z < 2.2$; Hou et al. 2020; Neveux et al. 2020). Also part of this joint release are a set of papers describing the data catalogues (Lyke et al. 2020; Ross et al. 2020), mock catalogues (Lin et al. 2020; Zhao et al. 2020), and N -body simulations for assessing systematic errors (Alam et al. 2020b; Avila et al. 2020; Rossi et al. 2020; Smith et al. 2020). At the highest redshifts ($z > 2.1$), the coordinated release of final eBOSS measurements includes measurements of BAO in the Lyman α forest (du Mas des Bourboux et al. 2020). eBOSS Collaboration (2020) presents the cosmological interpretation of these results combined with the final BOSS results and other probes.¹

In Section 2, we describe the main features of eBOSS catalogues, spectroscopic observations, details of the fibre assignment algorithm, and the veto masks. In Section 3, we describe the random catalogues used to perform clustering measurements. We present the PIP and angular upweighting schemes in Section 4. Tests on mock catalogues are described in Section 5 along with details on the survey mocks. We present the measurements from the final eBOSS catalogues in Section 6. Finally, we summarize our results and draw conclusions in Section 7.

Throughout this paper, we adopt the same Lambda cold dark matter (Λ CDM) fiducial cosmology adopted in other eBOSS DR16 papers with parameters: $\Omega_m = 0.31$, $\Omega_\Lambda = 0.69$, $\Omega_b h^2 = 0.022$, $h = 0.676$, $\sigma_8 = 0.8$, and $n_s = 0.97$.

2 SURVEY

The eBOSS was conceived to build upon the remarkable achievements of BOSS, and design a complementary survey to higher redshift. As well as extending the sample of high-redshift QSOs for Lyman α forest analyses and the well-known sample of LRGs from BOSS, it was designed to explore two new tracers: ELGs and QSOs used as direct tracers of the large-scale structure (LSS). The survey design and target selection was driven by the goal of achieving a ~ 1 per cent level precision on the measurements of the BAO. The target selection of different tracers is beyond the purpose of this paper and for the details of target selection, we refer the reader to Raichoor et al. (2017) for the ELG sample, Prakash et al. (2016) for LRG catalogue, and Myers et al. (2015) for QSOs. Detailed description of the survey design is provided in Dawson et al. (2016), while LSS catalogue creation is outlined in Ross et al. (2020) for LRG and QSO targets and in Raichoor et al. (2020) for the ELG sample.

eBOSS collected spectra using the double-armed fibre-fed BOSS spectrograph (Smee et al. 2013) at the 2.5m Sloan Telescope (Gunn et al. 2006). For each exposure, light is collected using optical fibres positioned in a pre-drilled aluminium plate placed at the focal plane of the telescope and transferred to the spectrographs. Each plate (or tile)

¹A description of eBOSS and links to all associated publications can be found here: <https://www.sdss.org/surveys/eboss/>.




Figure 1. Sky coverage of the eBOSS DR16 QSO catalogue. Blue dots are targets included in the LSS catalogue used for the cosmological analyses. Red dots represent targets not included in the LSS catalogues due to a number of issues such as redshift failure, fibre collisions, veto masks, low survey completeness. Black thick (thin) ellipses show the positions of eBOSS (SEQUELS) plates tiled for subsequent observation. Solid ellipses show tiles that were observed, and dashed ellipses those that were not observed. eBOSS LRGs have the same large-scale window function as the QSOs and hence are not shown.

covers $\sim 7 \text{ deg}^2$ and can accommodate 1000 fibres. Among these, 100 fibres are reserved for calibration targets with the remaining 900 dedicated for science targets. Each fibre and associated ferrule has a diameter of 62 arcsec, the ‘fibre-collision’ angle θ^{fc} . This puts a mechanical limitation on the minimum distances between two fibres on a single plate and prevents any pair of targets at smaller separation to be targeted simultaneously. However, multiple plates are allowed to partially overlap and the collision rate is therefore reduced in areas that are covered by more than one plate.

LRGs and QSOs shared the same plates as they were targeted simultaneously over $\sim 5000 \text{ deg}^2$ (Fig. 1) along with targets from the eBOSS subprograms, namely the Spectroscopic Identification of ERosita Sources (SPIDERS) and the Time-Domain Spectroscopic Survey (TDSS). A number of targets in the Legacy QSO target class, identified as having secure redshifts already measured from BOSS and previous surveys, were not reobserved. ELGs were observed using 300 dedicated plates over a limited area of $\sim 1200 \text{ deg}^2$ (Fig. 2) with a small fraction of fibres assigned to TDSS targets. The survey footprint is split into two regions denoted as the north (NGC) and the south galactic cap (SGC) (see Figs 1 and 2). As part of SDSS, eBOSS uses bitmasks to store information about targets.² In eBOSS, the EBOSS_TARGET1 and EBOSS_TARGET2 bitmasks are used to store the target types. The EBOSS_TARGET1 bits that identify eBOSS ‘clustering’ targets are LRG1_WISE for LRGs and QSO1_EBOSS_CORE for clustering QSOs. EBOSS_TARGET2 bits ELG1_NGC and ELG1_SGC are used to identify ELGs.

In this paper, we use PIP weights for galaxies within the eBOSS DR16 LSS catalogues (Raichoor et al. 2020; Ross et al. 2020). These catalogues include only targets with good redshifts in the redshift range of interest and are restricted over the footprint where the survey

completeness is greater than 50 per cent. The survey completeness, denoted COMP_BOSS, is defined as the ratio between the number of observed targets plus targets missed due to fibre collisions and the total number of target candidates within a ‘sector’. A sector is defined as the region covered by a unique set of plates. The redshift range of the LSS catalogues is limited to $0.6 < z < 1.0$ for LRGs, $0.8 < z < 2.2$ for QSOs, and to $0.6 < z < 1.1$ for ELGs.

eBOSS DR16 catalogues of LRGs and QSOs also include eBOSS targets observed as part of the Sloan extended quasar, ELG, and LRG Survey (SEQUELS), which linked the BOSS and eBOSS experiments. These targets are gathered in two adjacent regions (or chunks, see Section 2.1) denoted `boss214` and `boss217` in NGC. These regions were tiled using the software version and the priority system used for BOSS observations.

2.1 Fibre assignment

In eBOSS, the selection of targets for spectroscopic observation was performed through the survey ‘tiling’ process with the goal of maximizing the number of targets that receive a fibre with a given number of tiles (Blanton et al. 2003).

The fraction of targets of a given type that receive a fibre defines the tiling completeness. Collision groups are defined as a set of targets where each member of the group is in a fibre collision (i.e. with separation below 62 arcsec) with at least one other member, such that they cannot all be observed within a single exposure. The tiling is performed to maximize the tiling efficiency (i.e. the fraction of science fibres assigned to targets) while reaching the desired densities and decollided completeness³ of various targets.

³The decollided set of targets consists of targets not in collision groups combined with colliding targets that can be assigned a fibre on a single plate (see Dawson et al. 2016).

²<https://www.sdss.org/dr16/algorithms/bitmasks/>




Figure 2. Same as in Fig. 1 but here for the sample of ELGs.

As described in Blanton et al. (2003) in detail, the tiling was performed in three steps:

(i) first, the full survey footprint was divided into multiple chunks. In eBOSS, different chunks were observed independently even if they partially overlap, i.e. targets in the chunks overlap regions can potentially be targeted in more than one chunk (see Ross et al. 2020 for details on how these areas are treated in the LSS catalogue creation);

(ii) given the angular distribution of targets, tile centres in each chunk were initially drawn from a uniform covering of the celestial sphere which were then perturbed with respect to these initial positions to maximize the tiling completeness;

(iii) finally, given the tiling solution from (ii) fibres were assigned to targets, eventually solving collisions between targets.

Steps (i)–(iii) were adopted for the sufficiently large chunks. For chunks with narrow or small geometry, step (ii) in the procedure above was not able to produce an optimal configuration of tiles to reach the desired fibre efficiency and target density. For these chunks, we manually placed evenly located tiles for step (ii).

Given that different tracers were observed simultaneously, eBOSS observations adopted a tiered-priority system for assigning fibres to targets. A collision between two targets with different priorities is referred to as ‘knockout’ as opposed to fibre collisions that occur between targets of the same type.

For the chunks with LRG and QSO targets, fibres were assigned in two rounds: (i) all non-LRG targets get maximum priority and receive fibres first requiring a 100 per cent tiling completeness for their decollided set; (ii) in the second round, remaining fibres are assigned to LRG targets at a lower priority. The requirement of 100 per cent decollided completeness for LRG targets was lifted since they were selected with a higher number density than the available fibres and

they collide with targets with higher priority.⁴ In the first round, collisions between non-LRG targets were resolved with following decreasing priorities: SPIDERS, TDSS, known QSOs selected for reobservations, clustering QSOs, variability-selected QSOs, QSOs from the Faint Images of the Radio Sky at Twenty-Centimeters (FIRST) survey, and white dwarfs. The tiling density is on average one tile per 5 deg² for chunks observed as part of eBOSS while it is approximately one tile per 4 deg² for the two chunks observed as part of SEQUELS.

For the chunks dedicated to ELGs, the target densities were dominated by ELG targets in the first round with a small number of TDSS Few-Epoch Spectroscopy (FES) objects targeted at a density of $\sim 1 \text{ deg}^{-2}$ at the same priority as ELG targets.⁵ The tiling density of these chunks is on average one tile per around 4 deg².

2.2 Veto masks

The portion of the survey area where spectroscopic observations are impossible is accounted for by different veto masks. Under the assumption that veto masks do not correlate with the target sample, their effect is to change the survey window function. As such these regions are masked from both data and random catalogues used for clustering measurements (Raichoor et al. 2020; Ross et al. 2020).

In particular, for LRG and QSO samples, areas contaminated by bright stars are masked using the ‘bright-star’ veto mask. Other bright objects such as bright local galaxies and bright stars missed by the bright-star mask were also visually identified and masked using the ‘bright-object’ veto mask (Reid et al. 2016). Regions

⁴In fact, the tile centres of these LRG/QSO chunks were chosen using a downsampled set of the main LRG samples (by 15 per cent) in order to decrease the density of tiles; once the tile centres were decided, the fibre assignment pipeline was applied again to the entire sample.

⁵Additional TDSS targets were targeted on the same plates at lower priorities. The total density of TDSS targets that share plates with ELGs is $\sim 15 \text{ deg}^{-2}$.

with bad or no photometric observations are removed using the ‘bad-field’ veto mask. No fibre can be placed within 92 arcsec of the centre of each plate where a hole is drilled for the centrepost. These regions are part of the ‘centrepost’ veto mask. A mask around infrared bright stars was also applied to the LRG target file before tiling. Knockouts are also assumed to be uncorrelated with the target sample and are accounted for by applying the ‘collision-priority’ veto mask. In regions covered by a single tile, the collision-priority mask removes any QSO within 62 arcsec of a TDSS or SPIDERS target. The LRG collision-priority mask removes any LRG within 62 arcsec from a non-LRG target, regardless of whether it lies in single-pass or multipass regions. This is a conservative approach motivated by the fact that, being targeted at the lowest priority, not all knocked-out LRGs are targeted in areas covered by more than one tile. For further details on the construction and properties of different veto masks, we refer the reader to Ross et al. (2020).

The sample of ELGs presents similar effects. However, veto masks for ELG sample were more complicated and built in form of pixelized masks. They account for issues related to different bright stars, systematics in the target selection, and defects in the photometry. As anticipated in Section 2.1, ELG targets share plates with some TDSS targets. In particular, TDSS FES-type targets have the same targeting priority as ELGs. To account for this a veto mask, similar to the collision-priority mask for LRG and QSO targets, is created around each TDSS-FES target. We refer the reader to Raichoor et al. (2020) for a detailed description of the veto masks for ELGs.

2.3 Weights

In redshift surveys, the target selection is affected by different systematic effects that can alter the observed target density with respect to the true one (Ross et al. 2012; de la Torre et al. 2013; Bautista et al. 2018; Ross et al. 2020). Target selection from photometric data, fibre assignment, and inaccurate redshift estimates are typical systematic effects that affect the observed target density in spectroscopic redshift surveys. Each systematic effect is corrected by applying a weight to each target. In eBOSS, targets in the LSS catalogues are assigned the following weights to correct for systematics or optimize the clustering measurements:

- (i) w_{sys} corrects for spurious fluctuations in the photometric target selection;
- (ii) w_{cp} is the standard correction to fibre collisions adopted in eBOSS cosmological analyses. In the following part of this paper, we will refer to the w_{cp} weighting as the ‘CP’ correction (or weighting) method;
- (iii) w_{noz} accounts for redshift failures;
- (iv) w_{FKP} are the standard FKP weights (Feldman, Kaiser & Peacock 1994) used to minimize the variance of the measurement.

The CP correction is a variant of the standard NN method (see Ross et al. 2020). In particular, w_{CP} weights are computed for collision groups where the weight of each target missed due to a fibre collision is equally distributed among the observed members of the group rather than assigning it only to its nearest neighbour as in the standard NN method. In this work, we use the PIP weights as an alternative to w_{cp} keeping all other weights as in equation (1).

The overall standard weight is then

$$w_{\text{tot}} = w_{\text{sys}} \times w_{\text{noz}} \times w_{\text{FKP}} \times w_{\text{cp}}. \quad (1)$$

The same weights listed in this section are also assigned to the objects in the random catalogues used to perform clustering measurements. However, the task of these weights for random points is to match

the radial selection function of eBOSS targets rather than correcting for the related systematic effects. An exception is for w_{sys} weights for random points that are used to correct for survey completeness in each sector when the CP correction is adopted to correct for fibre collisions. We report, in Section 3, the procedure used in Ross et al. (2020) and Raichoor et al. (2020) to assign weights to objects in the random catalogues.

3 RANDOM CATALOGUES

For clustering measurements, we need to compare the galaxy distribution to the expected distribution, or window function, in order to determine the overdensity. Given the complexities in the window function this is usually quantified using random catalogues, matching the angular and radial selection functions of data targets, coupled with a set of weights applied to the galaxies. Here, we outline the main features of the random catalogues used in this paper. We refer the reader to the corresponding ELG (de Mattia et al. 2020; Raichoor et al. 2020) and LRG/QSO (Ross et al. 2020) catalogue papers for further details.

For all eBOSS tracers, random points are distributed homogeneously over the sky area covered by the survey and subsequently masked to match the footprint of the catalogues used for clustering analyses. This includes removing patterns inside vetoed regions and those excluded due to the low survey completeness (regions where COMP_BOSS drops below 50 per cent).

The next step is to assign redshifts to random points making sure they accurately reproduce the radial selection function of eBOSS targets. This is done in different ways for different tracers:

- (i) In the case of LRG and QSO samples, to ensure that random catalogues match the effective radial distribution of targets, their redshift and weights (w_{sys} , w_{cp} , w_{noz} , w_{FKP}) are drawn from a randomly selected galaxy within the sample.
- (ii) The radial selection function of ELG targets depends on the imaging depth. To account for this effect, survey area is first divided in subregions of approximately equal grz imaging depth. The redshift and w_{FKP} for a random point in a given subregions is then drawn from a randomly selected ELG target in the same subregion with a probability proportional to $w_{\text{sys}} \times w_{\text{cp}} \times w_{\text{noz}}$. A normalization factor is included in the w_{sys} weights of random points to assure that the ratio between weighted sums of random points and ELG targets is constant among different chunks while w_{cp} and w_{noz} are set to unity.

The standard CP correction for the fibre collision adopted in eBOSS cosmological analyses, a variation of the NN weighting (see Section 4), does not account for the sector-to-sector variation in the survey completeness (COMP_BOSS). This requires either upweighting target samples by the survey completeness or downweighting the randoms by the same quantity. The last option is adopted by embedding a COMP_BOSS factor in the w_{sys} weights for random points (Raichoor et al. 2020; Ross et al. 2020). Consequently, a fraction of legacy QSOs in each sector are removed to match the completeness of eBOSS QSOs. We make use of these random catalogues when correcting for fibre collisions with the modified NN weighting scheme adopted in eBOSS.

The PIP weights, on the other hand, are inferred by rerunning the eBOSS fibre assignment algorithm. As such they already account for the survey completeness since this lowers the probability of any target to get a fibre in sectors where the number of fibres is lower than the number of decollided targets. Consequently, no correction is required for the randoms, and in effect the PIP weights leave an isotropic expected galaxy distribution within the survey mask. We thus remove

the completeness factor from the w_{sys} weights for randoms when using the PIP corrections for fibre collisions. Furthermore, this also allows us to bring back into the LSS sample all legacy QSOs that were downsampled to match the survey completeness.

4 METHODOLOGY

In this section, we present the details of our analysis methodology. Before describing the PIP upweighting method, we first recap how standard measurements of the 2PCF are performed as this sets the scene for the PIP upweighting method. We then present the main features of the PIP and angular upweighting schemes and the specific way that target selection probabilities are inferred for eBOSS galaxy samples.

4.1 Measurement of the correlation function

We adopt the widely used least-biased and least-variance Landy–Szalay estimator (Landy & Szalay 1993) to measure the 2PCFs:

$$\xi(s) = \frac{DD(s) - 2DR(s)}{RR(s)} + 1, \quad (2)$$

where DD , DR , and RR are the data–data, data–random, and random–random pair counts normalized to the total number of corresponding weighted pairs and s is the pair separation vector in redshift space, i.e. when distances are inferred through the observed redshifts.

We measure two types of correlation functions:

(i) the projected correlation function $w_p(r_p)$ that is commonly used in literature to constrain the halo occupation distribution (HOD) models,

$$w_p(r_p) = 2 \int_0^{\pi_{\text{max}}} \xi(r_p, \pi) d\pi. \quad (3)$$

(ii) The multipoles $\xi^{(\ell)}$ of the redshift-space 3D correlation function, widely used to measure and model anisotropic 3D clustering in redshift space,

$$\xi^{(\ell)}(s) = (2\ell + 1) \int_0^{+1} \xi(s, \mu) L_\ell(\mu) d\mu. \quad (4)$$

In equation (3), r_p and π are the transverse and parallel to the line-of-sight components of the pair separations s . The integral in equation (3) is truncated at $\pi_{\text{max}} = 80 h^{-1} \text{Mpc}$ as measurements at larger scales are noise dominated. Given the angular coordinates (RA, Dec) and redshifts z , we compute r_p and π following Fisher et al. (1994). In equation (4), the angle-averaged pair separation is $s^2 = r_p^2 + \pi^2$ and the cosine of the angle between pair separation and line of sight is $\mu = \pi/s$.

In order to resolve the clustering at sub- $h^{-1} \text{Mpc}$ scales, we use a logarithmic binning for the pair separation s and its transverse to the line-of-sight component r_p

$$\log s_{i+1} = \log s_i + \Delta s_{\log}, \quad (5)$$

where $\Delta s_{\log} = 0.1$. The logarithmic mean of the bin edges is used as the sampling point (Mohammad et al. 2018). The line-of-sight component π of the pair separation is binned using a $1 h^{-1} \text{Mpc}$ linear binning. When computing the multipoles of the 2PCF, we divide μ between $[0, 1]$ in 200 linear bins.

4.2 Pairwise-inverse probability (PIP) weighting

The PIP weight for a given pair is defined as the inverse of the probability of it being targeted within an ensemble set of possible

realizations of the survey, which includes all possible pairs of targets, and from which the actual realization of the survey undertaken can be considered to be drawn at random. The selection probabilities depend strongly on the particular fibre assignment algorithm adopted to select targets from a parent photometric catalogue for the spectroscopic follow-up, and are therefore difficult to model except by rerunning the actual algorithm adopted. We thus rely on inferring the selection probabilities by generating multiple replicas of the survey target selection, changing the ‘random seed’ for each run such that different choices are made in which target to select for follow-up spectroscopy (see Section 4.4). The inverse probability is then simply estimated as the number of realizations N_{runs} in which a given pair could have been targeted divided by the number of times it was actually targeted (see Bianchi & Verde 2020, for a discussion about inverse-probability estimators; specifically, following the nomenclature introduced in that work, we adopt the zero-truncated estimator). The PIP correction gives unbiased measurements of the 2PCF provided that there are no pairs with zero probability of being targeted in the ensemble of survey realizations.

Following Bianchi & Percival (2017), rather than storing pairwise weights for individual pairs we store what are referred to as bitwise weights $w_i^{(b)}$ for each target. These are simply binary arrays of length N_{runs} where each bit (either one or zero) represents the outcome of the corresponding fibre assignment run for target i (either this target is, or is not included in run b). Bitwise weights are then combined ‘on the fly’ to compute the pairwise weights between target m and target n as

$$w_{mn} = \frac{N_{\text{runs}}}{\text{popcnt} [w_m^{(b)} \& w_n^{(b)}]}. \quad (6)$$

In equation (6), popcnt and $\&$ are standard bitwise operators. In particular, popcnt is the ‘population count’ operator that given an array, in this case a bit sequence of 0 and 1, returns the number of elements different than 0. In equation (6), $\&$ is the bitwise ‘and’ that, given two arrays of equal length, performs the logical ‘AND’ operation on each pair of the corresponding bits and returns the result as an array with length equal to that of input arrays. The weights w_m for individual targets, called individual-inverse-probability (IIP) weights, can be calculated simply by replacing $m = n$ in equation (6). The same random catalogue is valid for all fibre assignment runs, and so the pair counts in equation (2) are now:

$$\begin{aligned} DD(\vec{s}) &= \sum_{\vec{x}_m - \vec{x}_n \approx \vec{s}} w_{mn} w'_{\text{tot},m} w'_{\text{tot},n}, \\ DR(\vec{s}) &= \sum_{\vec{x}_m - \vec{y}_n \approx \vec{s}} w_m w'_{\text{tot},m} w_{\text{tot},n}, \end{aligned} \quad (7)$$

In equation (7), $w'_{\text{tot}} = w_{\text{sys}} \times w_{\text{noz}} \times w_{\text{FKP}}$ and w_{mn} and w_m are PIP and IIP weights, respectively. The RR pairs are computed using the overall weights in equation (1).

4.3 Angular upweighting

The PIP weighting scheme is unbiased only if there are no pairs with zero selection probability. This is not the case for pairs with separations below the fibre-collision scale that fall in the single pass regions of the survey. Indeed these pairs are systematically missed regardless of the number of survey realizations used to infer the selection probabilities since at least one of the two targets cannot be observed. However, a fraction of colliding pairs are targeted in regions where two or more tiles overlap. Under the assumption that the set of unobserved pairs is statistically equivalent to the set of

observed pairs, we can use the angular upweighting scheme proposed in Percival & Bianchi (2017) to recover the small-scale clustering. Although a reasonable assumption, survey designs may produce scenarios where this ansatz is not valid. In eBOSS, to maximize the targeting efficiency, areas where tiles overlap correlate to some extent with the regions of high target density. This can introduce a bias in the measurements from eBOSS DR16 catalogues shown in Figs 15–20 that cannot be corrected using methods available in literature.

The angular upweighting is performed both on the *DD* and *DR* pair counts. Equation (7) then becomes

$$DD(\vec{s}) = \sum_{\substack{\vec{x}_m - \vec{x}_n \approx \vec{s} \\ \vec{u}_m \cdot \vec{u}_n \approx \cos \theta}} w_{mn} w'_{\text{tot},m} w'_{\text{tot},n} \times w_{\text{ang}}^{\text{DD}}(\theta),$$

$$DR(\vec{s}) = \sum_{\substack{\vec{x}_m - \vec{y}_n \approx \vec{s} \\ \vec{u}_m \cdot \vec{v}_n \approx \cos \theta}} w_m w'_{\text{tot},m} w_{\text{tot},n} \times w_{\text{ang}}^{\text{DR}}(\theta), \quad (8)$$

where $\vec{u} = \vec{x}/x$. The angular weights $w_{\text{ang}}^{\text{DD}}(\theta)$ and $w_{\text{ang}}^{\text{DR}}(\theta)$ in equation (9), used to upweight *DD* and *DR* pair counts, respectively, are defined as

$$w_{\text{ang}}^{\text{DD}}(\theta) = \frac{DD^{\text{par}}(\theta)}{DD^{\text{fib}}_{\text{PIP}}(\theta)},$$

$$w_{\text{ang}}^{\text{DR}}(\theta) = \frac{DR^{\text{par}}(\theta)}{DR^{\text{fib}}_{\text{IIP}}(\theta)}. \quad (9)$$

The superscripts *par* and *fib* in equation (9) denote pairs of targets from the reference parent sample and pairs of targets that receive fibres, respectively. The subscript PIP and IIP denote the fact that the pair counts are upweighted using the PIP weights or their counterpart, the IIP weights, for the *DR* pairs. In the following part of this paper, we will use the abbreviation PIP+ANG to refer to the overall weighting outlined in equation (8).

The angular weights derived in Percival & Bianchi (2017) correct for the geometrical selection given by the survey targeting strategy. Because fibre assignment is independent of the properties upon which one normally selects subsamples, equation (9) can therefore be applied to any subsample of the parent catalogue selected e.g. by colour or redshift. This is demonstrated by Mohammad et al. (2018), where angular weights derived using the VIPERS parent sample were successfully applied to correct for ‘slit collisions’ for subsamples selected in two different redshift bins. Angular upweighting in the framework of the generalized inverse-probability weighting is discussed in Bianchi & Verde (2020).

4.4 Survey realizations

To infer realistic selection probabilities, we generate multiple random realizations of the survey that are statistically equivalent to the actual observations. The eBOSS tiling algorithm resolves collisions between targets in a random fashion by means of a random number generator, except for targets in collision groups with more than two objects. For targets within these collision groups, the algorithm performs a procedure designed to optimize the number of fibres allocated to targets (Blanton et al. 2003). As we will see later, this causes some problems for the PIP weights close to the fibre-collision scale, as a result of there being zero probability pairs. However, these effects are significantly below the noise level for any single realization of the survey.

We generate multiple survey realizations by rerunning the eBOSS fibre assignment algorithm many times, changing the random seed used to initiate the random number generator each time. However,




Figure 3. The distribution of the fraction of targets that get a fibre among 1860 fibre assignment runs on the LRG DR16 catalogue. The vertical red dashed line shows the mean of the distribution while the vertical shaded band represents the standard deviation. The vertical blue line shows the fraction of targets that received a fibre for the actual eBOSS observation.




Figure 4. As in Fig. 3, here for the catalogue of QSOs.

this feature is hardcoded in a large IDL-based software package that consists of several codes that optimize the fibre assignment and run sanity tests. Furthermore, the default set-up only allows a single survey realization to be generated at a time. For a processing time of $\sim 1\text{--}2$ h this presents the main limitation for our purpose where thousands of survey realizations are needed for real data sets and for a relatively large number of mock samples. We have thus modified the relevant components of the eBOSS tiling software package in order to generate multiple survey realizations simultaneously by running it in parallel on a multiprocessor computer cluster. In order to ease this task for future surveys, we suggest implementing target selection algorithms in well-documented packages written in a fast and widely used programming language.

In rerunning the fibre assignment, we keep the spatial distribution of tiles across all fibre assignment runs fixed, matching the distribution used in eBOSS observations. The accuracy of the selection probabilities calculated in this way depends on the number of survey realizations used. We use a total of 1860 fibre assignment runs to infer the selection probabilities when using eBOSS DR16 catalogues. The first of these runs corresponds to that used for the actual eBOSS observations. The outcome of these runs is stored in bitwise weights that are then used to compute the PIP weights and correct the pair counts following equations (6)–(8). Figs 3, 4, and 5 show the distribution of the fraction of targets that receive a fibre among these 1860 survey realizations (red histograms) and highlight the specific fibre assignment realization that was used for




Figure 5. As in Fig. 3, here for the catalogue of ELGs.

eBOSS observations (vertical blue lines). The run used for eBOSS spectroscopic observations (vertical blue line) shows a fraction of targets with fibre higher than the mean of the distribution, but within 1σ of the mean. For individual mocks, the number of fibre assignment runs is limited to 310, a fair compromise between the accuracy of selection probabilities and the computation time.

As shown in Fig. 1, a small fraction of plates were not observed even though they were included in the process of fibre assignment. Targets that were assigned a fibre from one of the unobserved plates thus contribute to lowering the survey completeness on the edges of the observed area. To correctly account for this, we run the fibre assignment using the full set of tiles and, once the fibres are assigned, we flag all targets that get a fibre from one of the unobserved plates as missed.

5 VALIDATION USING MOCK CATALOGUES

We use the Effective Zel’dovich mock samples (EZmocks; Zhao et al. 2020) to assess the performance of the PIP and angular upweighting schemes. These mocks are each built from a Gaussian random field in a $5 h^{-3}\text{Gpc}^3$ cubic box assuming an initial power spectrum and geometry matching a flat ΛCDM fiducial cosmology with parameters $\Omega_m = 0.307115$, $\Omega_b = 0.048206$, $h = 0.6777$, $\sigma_8 = 0.8225$, and $n_s = 0.9611$. The matter particles are displaced from their initial to final positions using the Zel’dovich approximation. The tracer number density is calculated using the matter density field and assuming a bias function that accounts for critical density required to form gravitationally bound structures, and for the stochasticity in the halo bias relation. The bias relation is calibrated using the real eBOSS data set to match the clustering on linear and mildly non-linear scales. Redshift-space distortions are added by means of a linear term calculated using the Zel’dovich approximation that accounts for the bulk flows. Non-linear motions are included through an isotropic Gaussian motion added to the linear component. Mocks are then cut according to the survey geometry removing areas not included in the DR16 LSS catalogues. We refer the reader to Zhao et al. (2020) for a detailed description of the mock construction.

In this work, we limit our analyses to only 100 mocks since running eBOSS fibre assignment on each mock catalogue requires a significant amount of computing time. The EZmocks made available to the eBOSS team are designed to match the LSS catalogues, in order to facilitate the cosmological analyses (de Mattia et al. 2020; Gil-Marín et al. 2020; Hou et al. 2020; Neveux et al. 2020; Raichoor et al. 2020; Zhao et al. 2020). Among other effects, fibre collisions are emulated in these mocks in an approximate way, removing objects

that do not receive a fibre, to reproduce the effect observed in eBOSS. Therefore, this set of mocks is not suitable for the purpose of this work where we need to assess realistic selection probabilities that requires processing each mock catalogue through the same eBOSS tiling package used for actual observations. We thus use ‘raw’ EZmocks, which have the same angular footprint as that of the eBOSS samples used for the cosmological analyses, and exhibit a flat radial selection function over the redshift range covered by the eBOSS DR16 LSS catalogues. We now describe the manipulations required in order to convert these mocks to match the eBOSS target samples.

5.1 Adding contaminants

Raw mocks were built to reproduce the clustering of the eBOSS DR16 LSS catalogues. However, they significantly differ from real data catalogues, used in the survey tiling, in the radial selection function and number density. We modify raw mocks to make them as realistic as possible and match the features of the data catalogues as following:

- (i) all objects in the raw mocks are flagged as eBOSS clustering targets. In particular, LRGs and QSOs are assigned EBOSS_TARGET1 bits LRG1_WISE and QSO1_EBOSS_CORE, respectively, while ELGs are flagged with EBOSS_TARGET2 bits ELG1_NGC (in NGC) and ELG1_SGC (in SGC);
- (ii) mocks are then downsampled to match the radial distribution of eBOSS clustering targets. In order to replicate the number density of targets, we match the number of mock targets to the weighted number of eBOSS targets in narrow redshift bins;
- (iii) eBOSS data targets with unknown redshifts or redshifts outside the mocks redshift range are added to each mock;
- (iv) spectroscopically confirmed stars initially misidentified as eBOSS targets in data are added as contaminants to the mock catalogues;
- (v) lack of mock targets in a posteriori vetoed regions, i.e. vetoed after survey tiling and spectroscopic observations, is compensated by adding targets from the eBOSS data in these regions to the mocks;
- (vi) a set of ancillary targets were observed simultaneously with eBOSS targets. When running the fibre assignment on mocks, we supplement the mock catalogues resulting from steps (i)–(v) with the same ancillary target sample used in real observations. However, in data catalogues there is a strong correlation between ancillary targets and QSO catalogue. Namely ~ 50 per cent of ancillary targets are within less than 2 arcsec of a QSO. These cases were treated as duplicates in eBOSS observations and they lower the number of fibres required to target QSO and ancillary samples. In mocks, this fraction decreases to below ~ 10 per cent since the correlation occurs mainly between ancillary targets and data targets that are added to mocks as contaminants in (iii)–(v). This has an effect of leaving a small number of fibres for mock LRG targets and therefore degrading the completeness of LRG sample targeted at a lower priority. To improve the completeness of mock LRG samples, we randomly downsample ~ 50 per cent of ancillary targets in chunks used for targeting LRG and QSO;
- (vii) finally, eBOSS ‘clustering’ QSOs that have known redshifts are not reobserved in eBOSS. These constitute ~ 21 per cent (13 per cent) of eBOSS QSO1_EBOSS_CORE sample in NGC (SGC) over the redshift range covered by mocks. We thus flag the same fractions of mock QSOs as known. As such these objects are not candidates to receive a fibre.




Figure 6. Weighted number of eBOSS DR16 LRG (red histogram), QSO (green histogram), and ELG (blue histogram) targets as a function of redshift z . The same quantities from 100 mocks are shown by the grey lines although they are not visually distinguishable from the coloured lines due to the tiny difference between data and mocks $N(z)$. Mock and data $N(z)$ histograms perfectly overlap at redshifts not covered by mock catalogues as all targets at these redshifts are taken from data catalogues (see Section 5).

The redshift distribution of the modified mocks (not accounting for contaminants added in steps iii–v) is shown in Fig. 6 along with the same quantity from eBOSS DR16 catalogues. The total number of targets in mocks differs, on average, from those in data catalogues by less than ~ 5 per cent.

We implement steps (i)–(vii) in order to make the fibre collisions in the mocks as close as possible to the real data. However, we are limited by the fact that EZmock catalogues are built using approximate prescriptions and are not based on the outcome of a full N -body simulation. As such it is inevitable that mock and data will present different small-scale clustering and consequently will have different collision rates. Nevertheless, this does not affect our conclusions as our tests are performed in order to demonstrate that we can recover the correct small-scale clustering without bias from the fibre collision issue. We therefore need to compare the clustering recovered from the mocks to the true clustering of the mocks, rather than to that of the data. It is worth noting that the three tracers show different small-scale clustering from each other as well as from the data, providing breadth to the tests performed.

Finally, we process mocks obtained from steps (i) to (vii) through the eBOSS tiling code to implement the realistic fibre assignment. In running the fibre assignment on mock samples, we follow the same procedure described in Section 2.1. In particular, the full survey area is split into multiple chunks, identical to those used for eBOSS data catalogues and we use the eBOSS tiered-priority system (Dawson et al. 2016) to solve collisions between different types of targets. In the following, we will refer to mocks resulting from step (i) to (ii) as ‘parent’ mocks while we will use the term ‘spectroscopic’ mocks to refer to the corresponding subsamples of mock targets that were assigned a fibre.

We do not include systematics other than fibre collisions in our parent and spectroscopic mocks. This choice is motivated by the fact that issues such as systematics in the photometric target selection and redshift failures are correlated with different galaxy properties (Ross et al. 2012; Scodreggio et al. 2018; Raichoor et al. 2020; Ross et al. 2020) and extremely difficult to accurately reproduce in simulated data sets. We therefore set $w_{\text{sys}} = w_{\text{noz}} = 1$. The FKP weights are computed as $w_{\text{FKP}} = 1/(1 + \bar{n}(z)P_0)$, where $\bar{n}(z)$ is the mean number density of mock targets and $P_0 = 4000 h^{-3} \text{Mpc}^3$ for ELGs (Raichoor et al. 2020), $P_0 = 6000 h^{-3} \text{Mpc}^3$ for QSOs, and

$P_0 = 10000 h^{-3} \text{Mpc}^3$ for LRGs (Ross et al. 2020). We infer the PIP weights through 310 fibre assignment runs on each mock.

The difference in the small-scale clustering between mocks and data also affects the implementation of the angular upweighting for the mock catalogues. The key requirement for the angular upweighting to be unbiased is that the sample of observed targets (labelled fib in equation 8), used for the clustering measurements, is statistically equivalent to the parent catalogue (labelled par in equation 8). In the case of mocks, the ‘spectroscopic’ mock sample is significantly different, in terms of small-scale clustering, than the mock catalogue used for fibre assignment since the latter one contains contaminants from eBOSS target catalogues added in steps (iii)–(v). We therefore use mocks obtained from steps (i) to (ii) as the reference parent samples (discarding any target added from the eBOSS target catalogue) to compute quantities labelled with par and the corresponding spectroscopic mock samples as the targeted samples to compute quantities labelled with fib in equation (8).

The goal of this section is to perform robustness tests for the novel PIP and PIP+ANG upweighting scheme. We present an indirect test of the validity of the standard CP correction for fibre collisions adopted in eBOSS in Section 6 where a comparison between PIP, PIP+ANG, and CP correction schemes is presented on the eBOSS DR16 LSS catalogues.

5.2 Effect of PIP and angular upweighting

We discuss here the impact that the PIP and angular upweighting are expected to have on clustering measurements presented in Sections 5 and 6 for mocks and eBOSS DR16 catalogues, respectively.

The PIP upweighting provides unbiased pair counts only if there are no pairs with zero probability of being targeted in a random realization of the survey. In eBOSS, zero-probability pairs do originate from fibre collisions in single-pass regions. These zero-probability pairs, however, are confined at angular separations below the fibre-collision angle $\theta^{\text{fc}} = 62$ arcsec. At these separations, the PIP weighting properly upweights pairs in the overlaps between multiple tiles but misses out those in areas covered by a single tile. At angular separations $\theta < \theta^{\text{fc}}$, we therefore expect the PIP upweighting to underestimate the pair counts inferred from the spectroscopic sample with respect to those from the parent catalogue. PIP upweighting provide virtually unbiased pair count at separations $\theta > \theta^{\text{fc}}$ where no or a very small number of zero-probability pairs (those resulting from an optimization of fibre assignment in collision groups with more than two objects) are expected. The effect is quantified by means of the angular weights $w_{\text{ang}}^{\text{DD}}(\theta)$ in equation (8) used to correct the DD pair counts and shown in Fig. 7 for LRGs (red), QSOs (green), and ELGs (blue). Indeed, in all cases $w_{\text{ang}}^{\text{DD}}(\theta)$ is greater than unity for separations below θ^{fc} and sharply reaches 1 for larger values of angular separation θ . A difference is noticeable between NGC (left-hand panel in Fig. 7) and SGC (right-hand panel in Fig. 7) due to the different tiling density between the two caps. For LRGs and QSOs for example, tiles in the SEQUELS chunks in the NGC are more tightly packed with respect to those in eBOSS chunks. This increases the fraction of the area covered by more than one tile decreasing the fraction of zero-probability pairs, which in turn decreases $w_{\text{ang}}^{\text{DD}}(\theta)$. The angular weights $w_{\text{ang}}^{\text{DD}}(\theta)$ for LRGs tend to have significantly lower amplitudes in the EZmocks (markers with error bars) compared to the eBOSS DR16 catalogue. This results from a lower intrinsic clustering of targets in the mocks with respect to eBOSS data combined with the fact that LRGs are targeted at the




Figure 7. Angular weights used to upweight DD pair counts (equation 8) in the clustering measurements for LRGs (red), QSOs (green), and ELGs (blue). Left-hand and right-hand panels display the weights for NGC and SGC. Markers show the mean estimate from 100 EZmocks with error bars being the error on the mean. Lines show the counterparts from DR16 catalogues with shaded bands showing the related errors on a single realization obtained from 100 EZmocks.

lowest priority. For QSOs and ELGs, the difference between mock and data in Fig. 7 is significantly smaller.

In order to assess how the PIP and PIP+ANG corrections impact the measurements of the projected correlation function or the multipole moments $\xi^{(l)}$ at different scales below the fibre-collision scale, it is useful to work with the anisotropic 2PCF $\xi^s(r_p, \pi)$, measured as a function of the parallel π and transverse to the line-of-sight r_p components of the pair separation, in terms of its natural estimator,

$$\xi(s) = \frac{DD(s)}{RR(s)} - 1. \quad (10)$$

With respect to the angular separations, the transverse scale that corresponds to the fibre-collision angle θ^{fc} in the fiducial cosmology varies with redshift. We denote with r_p^{fc} the transverse scale spanned by θ^{fc} at the maximum redshift of the sample. The PIP-corrected $DD(r_p, \pi)$ pair counts are then expected to be negatively biased at any π for $r_p < r_p^{fc}$ that in turn results in an underestimation of the ratio DD/RR and thus of the 2PCF with respect to the reference one. At scales $r_p > r_p^{fc}$, all pairs are at angular separations $\theta > \theta^{fc}$, a regime where PIP-corrections are unbiased, we expect $DD(r_p, \pi)$ corrected using PIP weights, and thus also the anisotropic 2PCF $\xi^s(r_p, \pi)$, to match its value from the reference parent sample. This is illustrated in Fig. 8 for LRGs using 100 EZmocks where we can compare the PIP upweighted measurements with the ones from parent samples. In Fig. 8, we show the ratio $DD(r_p, \pi)/RR(r_p, \pi)$ from equation (10) at two transverse scales smaller (top panels) and two larger than $r_p^{fc} \sim 0.7 h^{-1} \text{Mpc}$. As expected, the PIP correction strongly underestimates the DD/RR ratio, or equivalently the DD pair counts, for $\bar{r}_p < 0.7 h^{-1} \text{Mpc}$ where the angular upweighting is needed to properly account for zero-probability pairs. At transverse scales $\bar{r}_p > 0.7 h^{-1} \text{Mpc}$, on the other hand, angular upweighting has negligible effect as $w_{ang}^{DD}(\theta) \sim 1$ and the PIP and PIP+ANG upweighting are both unbiased.

5.3 Projected correlation function

Mean estimates of the projected correlation function $w_p(r_p)$ and the corresponding statistical errors from a set of 100 mocks are shown in the top panels of Figs 9–11 for the three tracers. The bottom panels of Figs 9–11 show the difference between measurements from the spectroscopic (see Section 5.1) and the parent mocks.




Figure 8. Ratio between DD and RR pair counts, averaged over 100 LRG EZmocks in SGC. Each panel displays the ratio as a function of the radial component π of the pair separation, in a single bin in the transverse to the line-of-sight component r_p of the pair separation. The DD/RR ratio at two r_p scales below the fibre-collision scale is shown in the top panels while the same quantity for two values of r_p larger than the fibre-collision scale are plotted in the bottom panels. Blue lines show the mean estimate from 100 parent mocks while blue filled and red empty markers show the result of PIP and PIP+ANG correction applied to the corresponding mocks affected by fibre collisions. Shaded bands and error bars show the errors on the mean estimates.




Figure 9. Projected correlation function of LRG EZmocks built as described in Section 5 in the two caps, NGC (left-hand panels) and SGC (right-hand panels). Top panels: mean estimate from 100 parent mocks (continuous lines), from catalogues affected by fibre collisions corrected using the PIP-only (empty markers with dashed error bars) and corrected using the combined PIP and angular upweighting (PIP+ANG, filled points with continuous error bars). The blue shaded bands and error bars show the error on the mean. The vertical red shaded bands show the transverse scales in the EZmocks fiducial cosmology corresponding to the fibre-collision angle between the minimum and maximum redshift of the sample. For separations r_p larger than the fibre-collision scale, PIP-only and the joint PIP+ANG corrections provide almost identical results. Therefore, empty (PIP) and filled (PIP+ANG) markers are these scales that cannot be easily distinguished. Empty markers at scales smaller than the fibre collision are not visible in the plot because they are well below the minimum limit set on the y-axis. Bottom panels: mean of the differences between the corrected measurements from mocks affected by fibre collisions and the corresponding parent mock. To reduce the range of variation, each quantity in the bottom panel is multiplied by r_p . For comparison, in the bottom panels that show the differences, the red continuous lines and hatch regions (RAW) show the mean measurements from spectroscopic mocks and related errors in the case where no correction for fibre collisions is applied.




Figure 10. Same as in Fig. 9 here for the QSO EZmocks.




Figure 11. Same as in Fig. 9 here for the ELG EZmocks.

In the bottom panels of Figs 9–11, we report the case where no correction for fibre collisions or survey completeness is applied (red continuous lines with hatch error regions) to the measurement from the spectroscopic mocks. Given the sparsity of the sample and being targeted at higher priority, missing observations have negligible effect on the clustering of QSOs (see Fig. 10) on scales larger than the fibre-collision scale where missing targets resemble the effect of a random depletion. Raw estimate of the correlation function of LRGs (see Fig. 9), on the other hand, shows negligible offset with respect to the reference one on scales $\sim 1\text{--}10 h^{-1}\text{Mpc}$ and scales $\gtrsim 100 h^{-1}\text{Mpc}$, while an offset at $\sim 2\sigma$ level is clearly visible on scales between ~ 30 and $100 h^{-1}\text{Mpc}$. Raw measurements of the projected correlation function of ELGs mocks in Fig. 11 show an overall agreement with the clustering of the parent samples with a marginal offset at scales of $\sim 10 h^{-1}\text{Mpc}$. These raw measurements show that although the fibre collisions have a limited impact compared to the statistical errors, its strength strongly depends on the intrinsic clustering of the tracers and the features of the particular fibre assignment algorithm adopted.

Measurements from spectroscopic mocks, corrected using the PIP technique and averaged over 100 mocks, are shown as empty markers in the top panels of Figs 9–11 along with the reference measurements from parent mocks (continuous lines). For all three tracers, the agreement between the PIP-only corrected and reference measurements

is remarkable for transverse scales larger than $\sim 1 h^{-1}\text{Mpc}$ (bottom panels in Figs 9–11). However, PIP weighting fails to recover the input clustering at transverse scales r_p below the fibre-collision scale (vertical red shaded bands in Figs 9–11). The interpretation of the systematic offset at these scales, when PIP-only upweighting is applied, follows from the discussion in Section 5.2. In particular, in the limit of small r_p and small π , where targets are strongly clustered, the anisotropic correlation function $\xi^s(r_p, \pi)$ is well approximated by DD/RR and the PIP-corrected $\xi^s(r_p, \pi)$ results are offset by a factor of $w_{\text{ang}}^{\text{DD}}(\theta)$ with respect to the reference. At small r_p and large π , specifically between $\pi \sim 50 h^{-1}\text{Mpc}$ and the upper limit in the integral in equation (3) fixed at $80 h^{-1}\text{Mpc}$, i.e. in the regime of weak intrinsic clustering, DD/RR approaches unity and $\xi^s(r_p, \pi)$ tends to 0. At these scales, the underestimation in the PIP upweighted DD pair counts reduces the corresponding $\xi^s(r_p, \pi)$ to very small or negative values and the scaling by $w_{\text{ang}}^{\text{DD}}(\theta)$ between the PIP-corrected and reference $\xi^s(r_p, \pi)$ is not valid anymore. This drives the PIP-corrected projected correlation function $w_p(r_p)$, obtained integrating $\xi^s(r_p, \pi)$ in equation (3), to values below the lower limit on y-axis shown in the top panels of Figs 9–11 enhancing the relative difference between PIP upweighted and reference measurements well above the factor of $w_{\text{ang}}^{\text{DD}}(\theta)$.

The angular upweighting uses the fraction of close pairs lost in single-pass regions to restore the DD/RR ratio at small transverse scales r_p , to its expectation values. As a result, when PIP weights are supplemented with the angular upweighting (filled markers with error bars in Figs 9–11), we are able to successfully recover the clustering signal down to very small scales $\sim 0.1 h^{-1}\text{Mpc}$ without altering the large-scale measurements.

It is important to stress that, as opposed to the raw measurements, the performance of both the PIP and joint PIP and angular upweighting does not vary with the type of tracers. As anticipated in Section 4.4, the optimization performed by the fibre assignment algorithm within the collision groups can give rise to the ‘zero-probability’ pairs at the scale of fibre collisions. This is likely to be the source of the small deviation in the PIP+ANG corrected measurements at the close-pair scales (vertical red shaded bands) seen in Figs 9–11 with respect to the reference. The effect, stronger for the LRGs in Fig. 9, is well within the statistical error for a single realization and becomes evident only when averaged over a high number of samples. We therefore do not consider this further: as demonstrated in the plot it is small and limited to a narrow range of scales close to the collision scale.

5.4 Multipoles

We now test the corrections for the multipoles of the redshift-space 2PCF. We limit the tests to the first three even multipoles, namely the monopole, quadrupole, and hexadecapole that are mostly used to detect and model the redshift-space distortions and the BAO in large galaxy redshift surveys. Results are shown in Figs 12–14, and match the results presented for the projected correlation functions in Figs 9–11.

Differences between ‘raw’ measurements from spectroscopic mocks not corrected for fibre collisions or survey completeness, and the reference measurements from the parent mocks are shown in the bottom panels of Figs 12–14 (dashed lines with hatch error regions). Raw measurements tend to underestimate the reference clustering up to scales of $\sim 10\text{--}20 h^{-1}\text{Mpc}$ for all tracers although the bias is more severe for LRGs due to their higher clustering and being targeted at the lowest priority. At larger scales, a mild bias appears for LRGs

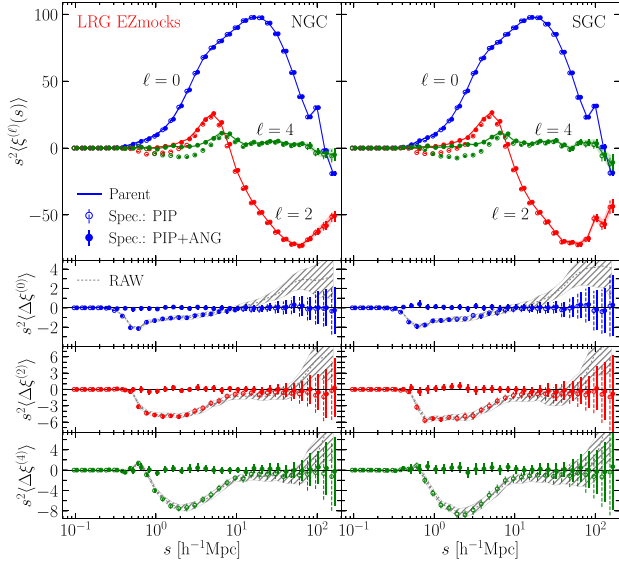


Figure 12. Multipoles of the redshift-space 2PCFs from 100 EZmocks of LRGs. In all panels, measurements are averaged over 100 mocks and the error bars correspond to the error on the mean. Top row: mean measurements from the reference parent mock catalogues (lines with shaded bands) and from the corresponding samples of targeted objects that are corrected using PIP-only (empty markers with dashed error bars) and combined PIP and angular (PIP+ANG, filled markers with thick error bars) weighting schemes. Bottom three panels: mean difference between the raw measurements (grey dotted lines and hatches), PIP (empty markers), PIP+ANG (filled markers) corrections and the reference measurements.

while it is not clearly visible for QSOs and ELGs due to relatively large statistical errors at these scales.

The PIP corrections are unbiased on scales larger than $\sim 10\text{--}20 h^{-1}\text{Mpc}$ recovering the input clustering signal within less than $\sim 1\sigma$ errors. However, they systematically underestimate the clustering at smaller scales. The nature of the systematic bias in the PIP-corrected measurements of the multipole moments $\xi^{(\ell)}$ is the same discussed in Sections 5.2 and 5.3 for the projected correlation function. However, the effect, confined at small transverse scales r_p in the anisotropic $\xi^*(r_p, \pi)$ and projected $w_p(r_p)$ correlation functions, is now spread to all angle-averaged separations s in the multipoles $\xi^{(\ell)}$. At angle-averaged scales smaller than the transverse fibre-collision scale $s < r_p^{\text{fc}}$, PIP weighting underestimates the reference $DD(s, \mu)$ pair count at any value of μ resulting in a strong negative bias in the measured multipoles. At scales $s \gtrsim r_p^{\text{fc}}$ the bias is reduced due to the fact that the PIP-corrected $DD(s, \mu)$ pair counts are underestimated, with respect to the reference, only between $1 < \mu < \mu^{\text{fc}}$ with,

$$\mu^{\text{fc}} = [1 - (r_p^{\text{fc}}/s)]^{1/2}. \quad (11)$$

At scales s significantly larger than r_p^{fc} , μ^{fc} approaches unity and the underestimate in the DD pair counts (see Section 5.2) between $[\mu^{\text{fc}}, 1]$ has negligible effect on the multipole moments $\xi^{(\ell)}$. The systematic bias at a given scale s is also higher for higher order multipoles. This follows from the μ dependence of the Legendre polynomials in equation (4).

As for the projected correlation function $w_p(r_p)$, the angular weights properly upweight the DD pair counts below the fibre-collision scale providing unbiased measurements of the anisotropic correlation function $\xi(s, \mu)$ and its multipoles. The combined PIP+ANG correction results are unbiased at all scales explored in




Figure 13. Same as in Fig. 12, here for the QSO mock samples.




Figure 14. Same as in Fig. 12, here for the mock samples of ELGs.

this work at a level well below the statistical errors, as shown in the bottom panels of Figs 12–14.

6 RESULTS

The tests performed using the EZmock catalogues in Section 5 showed that we can successfully recover the input clustering signal down to very small scales, well within the one-halo term. We now apply the same corrections to the eBOSS DR16 LSS catalogues. Since we deal with a single catalogue for each tracer, we increase the number of survey realizations, used to infer the selection probabilities, from 310 used for each mock catalogue to 1860. The implementation of the angular upweighting for the DR16 catalogues is slightly different from that outlined in Section 5 for the mock catalogues. In particular, for each tracer we now use the




Figure 15. Top panel: Measurements of the projected correlation function of the eBOSS DR16 LRG sample using different correction schemes: close-pair weighting (CP, dashed red lines), PIP weighting (PIP, black dash-dotted line), combined PIP and angular upweighting (PIP+ANG, blue continuous lines). The shaded band shows the 1σ errors estimated using 100 mock samples. Bottom panel: Difference between the measurements obtained using CP and PIP weighting with respect to measurements using combined PIP and angular upweighting. The shaded band shows the 1σ statistical error. To reduce the range of variation, each quantity in the bottom panel is multiplied by r_p . As in Fig. 9, the vertical shaded red bands show the transverse scales corresponding to the fibre-collision angle between the minimum and maximum redshift of the sample in the eBOSS fiducial cosmology.

corresponding full input sample as the parent catalogue to compute the quantities labelled with par, and their targeted subsamples to compute the ones denoted with fib in equation (8) regardless of their redshifts. We also compare the novel correction schemes to the standard CP weighting used in eBOSS DR16 cosmology papers.

The shaded bands in Figs 15–20 show the statistical error on a single eBOSS realization derived using a set of 100 EZmocks. These mocks only provide approximate realizations of the galaxy distribution on small scales, and so we do not expect the error bars to fully capture the fluctuations observed in eBOSS data at scales $\lesssim 1 h^{-1}\text{Mpc}$.

6.1 Projected correlation function

Figs 15, 16, and 17 show the measured projected correlation functions from the eBOSS DR16 samples of LRGs, QSOs, and ELGs, respectively. Measurements (top panels) are corrected using the standard CP (red dashed lines), PIP (black dash-dotted lines), and the joint PIP and angular upweighting (PIP+ANG, blue thick lines). In the bottom panels of the same figures, we show the difference of the standard CP and PIP corrections with respect to the joint PIP+ANG upweighting taken as the reference since it is found to be unbiased within data precision over all scales in the tests performed on mock catalogues.

The PIP upweighting matches the PIP+ANG correction for scales larger than the fibre-collision scales (vertical red shaded bands) while it significantly underestimates the clustering at smaller scales. The CP corrections perform very similarly to the PIP ones with deviations consistent with the statistical noise on scales larger than the collision scale. The agreement between the CP and PIP corrections at scales larger than the fibre-collision scale shows that the selection probabilities are highly uncorrelated on these scales and can be well approximated using an empirical prescription such




Figure 16. Equivalent of Fig. 15 for the eBOSS DR16 QSO catalogue.




Figure 17. Equivalent of Fig. 15 for the eBOSS DR16 catalogue of ELGs.

as the NN method. This directly follows from the features of the eBOSS fibre assignment algorithm that uses a random seed to resolve collisions. The strong difference between the CP and PIP corrections with respect to the PIP+ANG weighting below the fibre-collision scale (vertical red shaded bands in Figs 15–17) is due to the effect discussed in Section 5.2 and reflects the trend observed for mocks in Section 5.3. Comparing Figs 15–17 for eBOSS DR16 samples with their equivalent for mocks in Figs 9–11 it is clear that the eBOSS DR16 targets show a higher intrinsic clustering at scales below $1 h^{-1}\text{Mpc}$. Therefore, collisions are expected to occur at a higher rate in eBOSS catalogue with respect to the mocks. This is the source of the increase in the absolute size of the small-scale difference between PIP/CP and the PIP+ANG weighting scheme between the DR16 data and the mocks.

6.2 Multipoles

In Figs 18, 19, and 20, we show the measurements of the redshift-space multipole correlation functions for the eBOSS DR16 samples of LRGs, QSOs, and ELGs, respectively. The top panels show only the measurements performed using our reference PIP+ANG upweighting. In the bottom panels, we plot the deviations of the CP and PIP corrections with respect to the measurements that use the joint PIP and angular upweighting.




Figure 18. Top panel: Measurements of the redshift-space multipole correlation functions of eBOSS DR16 LRG sample corrected using the combined PIP and angular upweighting. The shaded band shows 1σ errors estimated using 100 mock samples. Bottom panels: Difference of the measurements obtained using CP and PIP weighting with respect to the one using combined PIP and angular upweighting. The shaded bands show the 1σ statistical error from the top panels.




Figure 19. Same as in Fig. 18, here for the eBOSS DR16 QSO catalogue.

As for the projected correlation functions shown in Figs 15–17, CP and PIP corrections provide similar results with discrepancies consistent with statistical uncertainties. The difference of the CP and PIP upweighting with respect to the joint PIP and angular corrections appears to be significant at scales smaller than $\sim 10 h^{-1}\text{Mpc}$. Such a difference is due to the presence of zero-probability pairs at transverse separations below the fibre-collision scale as discussed in Section 5.2. These results are consistent with the measurements from the mock catalogues discussed in Section 5.4. At scales below $\sim 10 h^{-1}\text{Mpc}$ CP and PIP upweighting perform better for the QSO




Figure 20. Same as in Fig. 18, here for the eBOSS DR16 catalogue of ELGs.

catalogue and ELGs with respect to the sample of LRGs. This is expected since the LRGs show a higher small-scale clustering compared to the QSOs and ELGs (see Figs 15–17).

7 SUMMARY AND CONCLUSIONS

We have applied the PIP and the angular upweighting to the 3D clustering measurements from eBOSS DR16 catalogues in order to correct for the systematic bias arising from fibre collisions in the spectroscopic observations. We inferred selection probabilities by means of multiple survey realizations obtained rerunning several times the SDSS tiling code on the input target catalogues. We focused on the measurements of the projected correlation function and the multipoles of the redshift-space 2PCFs.

We used the approximate ‘effective Zel’dovich’ EZmocks to test the performance of the correction method. To this end, we edited raw mock catalogues for the LRGs, ELGs, and QSOs, to obtain synthetic samples as close as possible to real catalogues in terms of number density of targets. These mock catalogues were processed using the same SDSS tiling code used for the actual eBOSS target selection for the spectroscopic follow-up, in order to reproduce fibre collisions. Clustering measurements were then corrected for missing observations using PIP and angular upweighting schemes. PIP corrections applied to the projected correlation function $w_p(r_p)$ provided unbiased measurements on scales larger than $r_p \sim 1 h^{-1}\text{Mpc}$ but were strongly biased at smaller scales. This is consistent with the presence of the so-called zero-probability pairs due to the fibre collisions in single-pass regions. Including the angular upweighting provided unbiased estimates on scales as small as $0.1 h^{-1}\text{Mpc}$ within data precision. When averaged over 100 EZmocks, a systematic bias was found in the $w_p(r_p)$ measurements close to the fibre-collision scales due to the eBOSS strategy of maximizing the targeting efficiency in collision groups. These results were reproduced also in the tests on the multipoles of the redshift-space correlation function. However, due to the projection on Legendre polynomials, the systematic bias for the PIP correction alone gets spread over angle-averaged scales as large as $\sim 10 h^{-1}\text{Mpc}$. Combining the PIP corrections with angular upweighting, we successfully removed the

residual systematic offset in the measured multipoles at scales below $\sim 10 h^{-1}\text{Mpc}$.

The EZmocks used in this work are based on the Zel'dovich approximation. As such they do not faithfully reproduce the small-scale clustering observed in real data. Mock targets exhibit a lower level of clustering compared to the eBOSS DR16 samples (see Figs 9–11 and 15–17). This makes the fibre-collision issue less severe in the mocks with respect to real data. However, we tested the method using three different tracers. We recovered the input clustering within 1σ errors in all three cases despite the fact that different tracers exhibit significantly different intrinsic clustering strength and are affected by fibre collisions in different ways. This assures us that the technique adopted in this paper is universal and as such does not depend on intrinsic features of a particular sample.

We finally applied the PIP+ANG weighting to the eBOSS DR16 catalogues to correct for fibre collision when estimating the projected and multipole correlation functions. We compared the joint PIP and angular upweighting (PIP+ANG) with the PIP-only and with the standard ‘CP’ weighting, a modified version of the NN method, that is adopted in eBOSS cosmological analyses and previous SDSS spectroscopic samples such as those from the BOSS. As expected from tests performed on mock catalogues, the three correction methods perform similarly at scales above $\sim 1 h^{-1}\text{Mpc}$ for the projected correlation function and $\sim 10 h^{-1}\text{Mpc}$ for the multipole correlation functions, with differences consistent with a statistical noise. On smaller scales, both PIP and CP weighting provide a very low clustering compared to the PIP+ANG corrections.

Our analysis provides a robustness test for the standard technique adopted in BOSS and eBOSS cosmological analyses to correct for fibre collisions on scales larger than $10 h^{-1}\text{Mpc}$. Additionally, the PIP and angular upweighting scheme tested in this paper is more robust and provides the most accurate 3D clustering measurements down to $\sim h^{-1}\text{kpc}$ scales. This gives us access to scales where one-halo term dominates in the clustering signal. As such these scales are crucial to constrain the HOD models that study how different tracers populate dark matter haloes of different masses. Moreover, analyses that rely on numerical simulation or HOD formalism to model redshift-space distortions can now be pushed to smaller scales where the clustering signal is measured with high significance. This will allow putting even tighter constraints on the growth rate of structure, a key parameter to constrain gravity models at cosmological scales. These analyses will be presented in future work.

ACKNOWLEDGEMENTS

This research was supported by the Centre for the Universe at the Perimeter Institute. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ontario Ministry of Economic Development, Job Creation and Trade. We acknowledge support provided by Compute Ontario (www.computeontario.ca) and Compute Canada (www.computecanada.ca).

HJS is supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics under Award Number DE-SC0014329.

Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS web site is www.sdss.org.

SDSS-IV is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, the Chilean Participation Group, the French Participation Group, Harvard-Smithsonian Center for Astrophysics, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU) / University of Tokyo, the Korean Participation Group, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatório Nacional / MCTI, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University.

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 693024).

DATA AVAILABILITY

The bitwise weights, used to compute the PIP and angular weights, will be made available after acceptance of the manuscript through the SDSS Science Archive Server (<https://sas.sdss.org/>).

REFERENCES

- Alam S. et al., 2017, *MNRAS*, 470, 2617
- Alam S., Peacock J. A., Kraljic K., Ross A. J., Comparat J., 2020a, *MNRAS*, 497, 581
- Alam S. et al., 2020b, preprint ([arXiv:2007.09004](https://arxiv.org/abs/2007.09004))
- Avila S. et al., 2020, preprint ([arXiv:2007.09012](https://arxiv.org/abs/2007.09012))
- Bautista J. E. et al., 2018, *ApJ*, 863, 110
- Bautista J. E. et al., 2020, preprint ([arXiv:2007.08993](https://arxiv.org/abs/2007.08993))
- Bianchi D., Percival W. J., 2017, *MNRAS*, 472, 1106
- Bianchi D., Verde L., 2020, *MNRAS*, 495, 1511
- Bianchi D. et al., 2018, *MNRAS*, 481, 2338
- Blake C. et al., 2012, *MNRAS*, 425, 405
- Blanton M. R., Lin H., Lupton R. H., Maley F. M., Young N., Zehavi I., Loveday J., 2003, *AJ*, 125, 2276
- Blanton M. R. et al., 2017, *AJ*, 154, 28
- Chen Y.-M. et al., 2013, *MNRAS*, 429, 2643
- Cole S. et al., 2005, *MNRAS*, 362, 505
- Dawson K. S. et al., 2016, *AJ*, 151, 44
- de la Torre S. et al., 2013, *A&A*, 557, A54
- de Mattia A. et al., 2020, preprint ([arXiv:2007.09008](https://arxiv.org/abs/2007.09008))
- Drinkwater M. J. et al., 2010, *MNRAS*, 401, 1429
- du Mas des Bourboux H. et al., 2020, preprint ([arXiv:2007.08995](https://arxiv.org/abs/2007.08995))
- eBOSS Collaboration, 2020, preprint ([arXiv:2007.08991](https://arxiv.org/abs/2007.08991))
- Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
- Eisenstein D. J. et al., 2011, *AJ*, 142, 72
- Feldman H. A., Kaiser N., Peacock J. A., 1994, *ApJ*, 426, 23
- Fisher K. B., Davis M., Strauss M. A., Yahil A., Huchra J. P., 1994, *MNRAS*, 267, 927
- Gil-Marín H. et al., 2020, preprint ([arXiv:2007.08994](https://arxiv.org/abs/2007.08994))
- Gunn J. E. et al., 2006, *AJ*, 131, 2332
- Guo H., Zehavi I., Zheng Z., 2012, *ApJ*, 756, 127

- Guo H. et al., 2019, *ApJ*, 871, 147
- Guzzo L. et al., 2014, *A&A*, 566, A108
- Hahn C., Scoccimarro R., Blanton M. R., Tinker J. L., Rodríguez-Torres S. A., 2017, *MNRAS*, 467, 1940
- Hou J. et al., 2020, preprint ([arXiv:2007.08998](https://arxiv.org/abs/2007.08998))
- Jing Y. P., Mo H. J., Börner G., 1998, *ApJ*, 494, 1
- Krywult J. et al., 2017, *A&A*, 598, A120
- Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
- Laurent P. et al., 2017, *J. Cosmol. Astropart. Phys.*, 2017, 017
- Le Fèvre O. et al., 2003, in Iye M., Moorwood A. F. M., eds, Proc. SPIE Conf. Ser. Vol. 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes. SPIE, Bellingham, p. 1670
- Lewis I. et al., 2002, *MNRAS*, 334, 673
- Lin S. et al., 2020, preprint ([arXiv:2007.08996](https://arxiv.org/abs/2007.08996))
- Lyke B. W. et al., 2020, preprint ([arXiv:2007.09001](https://arxiv.org/abs/2007.09001))
- Mohammad F. G. et al., 2018, *A&A*, 619, A17
- Myers A. D. et al., 2015, *ApJS*, 221, 27
- Neveux R. et al., 2020, preprint ([arXiv:2007.08999](https://arxiv.org/abs/2007.08999))
- Percival W. J., Bianchi D., 2017, *MNRAS*, 472, L40
- Percival W. J. et al., 2001, *MNRAS*, 327, 1297
- Pezzotta A. et al., 2017, *A&A*, 604, A33
- Prakash A. et al., 2016, *ApJS*, 224, 34
- Raichoor A. et al., 2017, *MNRAS*, 471, 3955
- Raichoor A. et al., 2020, preprint ([arXiv:2007.09007](https://arxiv.org/abs/2007.09007))
- Reid B. A., Seo H.-J., Leauthaud A., Tinker J. L., White M., 2014, *MNRAS*, 444, 476
- Reid B. et al., 2016, *MNRAS*, 455, 1553
- Ross A. J. et al., 2012, *MNRAS*, 424, 564
- Ross A. J. et al., 2020, preprint ([arXiv:2007.09000](https://arxiv.org/abs/2007.09000))
- Rossi G. et al., 2020, preprint ([arXiv:2007.09002](https://arxiv.org/abs/2007.09002))
- Scodeggio M. et al., 2018, *A&A*, 609, A84
- Smee S. A. et al., 2013, *AJ*, 146, 32
- Smith A. et al., 2019, *MNRAS*, 484, 1285
- Smith A. et al., 2020, preprint ([arXiv:2007.09003](https://arxiv.org/abs/2007.09003))
- Sunayama T. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 057
- Tamone A. et al., 2020, preprint ([arXiv:2007.09009](https://arxiv.org/abs/2007.09009))
- Yang L., Jing Y., Yang X., Han J., 2019, *ApJ*, 872, 26
- York D. G. et al., 2000, *AJ*, 120, 1579
- Zarrouk P. et al., 2018, *MNRAS*, 477, 1639
- Zehavi I. et al., 2002, *ApJ*, 571, 172
- Zehavi I. et al., 2005, *ApJ*, 630, 1
- Zhao C. et al., 2020, preprint ([arXiv:2007.08997](https://arxiv.org/abs/2007.08997))
- ³Perimeter Institute for Theoretical Physics, 31 Caroline St North, Waterloo, ON N2L 2Y5, Canada
- ⁴Department of Physics and Astronomy, Ohio University, 251B Clippinger Labs, Athens, OH 45701, USA
- ⁵Institut de Ciències del Cosmos, Universitat de Barcelona, ICCUB, Martí i Franquès 1, E-08028 Barcelona, Spain
- ⁶Center for Cosmology and AstroParticle Physics, Ohio State University, Columbus, OH 43210, USA
- ⁷Institute of Physics, Laboratory of Astrophysics, École Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, CH-1290 Versoix, Switzerland
- ⁸Institute of Cosmology & Gravitation, Dennis Sciama Building, University of Portsmouth, Portsmouth PO1 3FX, UK
- ⁹Apache Point Observatory and New Mexico State University, P.O. Box 59, Sunspot, NM 88349, USA
- ¹⁰Department Physics and Astronomy, University of Utah, 115 S 1400 E, Salt Lake City, UT 84112, USA
- ¹¹IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France
- ¹²Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, 452 Lomita Mall, Stanford, CA 94305, USA
- ¹³Aix Marseille Univ, CNRS, CNES, LAM, F-13388 Marseille, France
- ¹⁴Instituto de Ciencias Físicas, Universidad Nacional Autónoma de México, Av. Universidad s/n, 62210 Cuernavaca, Mor., México
- ¹⁵Institut d'Estudis Espacials de Catalunya (IEEC), E-08034 Barcelona, Spain
- ¹⁶Max-Planck-Institut für Extraterrestrische Physik, Postfach 1312, Giessenbachstr, D-85748 Garching bei München, Germany
- ¹⁷Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK
- ¹⁸Department of Physics and Astronomy, Sejong University, Seoul 143-747, Korea
- ¹⁹Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA
- ²⁰Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA
- ²¹Center for Cosmology and Particle Physics, Department of Physics, New York University, New York, NY 10003, USA
- ²²School of Physics and Astronomy, University of St Andrews, St Andrews KY16 9SS, UK
- ²³Instituto de Física, Universidad Nacional Autónoma de México, Apdo. Postal 20-364, 04510 Ciudad de México, México
- ²⁴National Astronomy Observatories, Chinese Academy of Science, Beijing 100101, P.R. China
- ²⁵School of Astronomy and Space Science, University of Chinese Academy of Sciences, Beijing 100049, P.R. China

¹Waterloo Centre for Astrophysics, Department of Physics and Astronomy, University of Waterloo, Waterloo, ON N2L 3G1, Canada

²Department of Physics and Astronomy, University of Waterloo, Waterloo, ON N2L 3G1, Canada

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