

A cross-correlation-based estimate of the galaxy luminosity function

Marcel P. van Daalen^{1,2★} and Martin White¹

¹*Department of Astronomy, Theoretical Astrophysics Center, and Lawrence Berkeley National Laboratory, University of California, Berkeley, CA 94720, USA*

²*Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*

Accepted 2018 February 23. Received 2018 February 23; in original form 2017 March 15

ABSTRACT

We extend existing methods for using cross-correlations to derive redshift distributions for photometric galaxies, without using photometric redshifts. The model presented in this paper simultaneously yields highly accurate and unbiased redshift distributions and, for the first time, redshift-dependent luminosity functions, using only clustering information and the apparent magnitudes of the galaxies as input. In contrast to many existing techniques for recovering unbiased redshift distributions, the output of our method is not degenerate with the galaxy bias $b(z)$, which is achieved by modelling the shape of the luminosity bias. We successfully apply our method to a mock galaxy survey and discuss improvements to be made before applying our model to real data.

Key words: galaxies: distances and redshifts – galaxies: luminosity function, mass function – large-scale structure of Universe – cosmology: theory.

1 INTRODUCTION

Current and planned large galaxy surveys are bringing in enormous amounts of photometric data. Spectroscopic follow-up for even a tenth of these sources is infeasible, and so many techniques have been developed to derive valuable redshift information indirectly for the vast majority of observed galaxies. Classically, estimating redshifts or redshift distributions has been performed using photometry in combination with a library of SEDs and/or spectroscopic sources to train the algorithms used, yielding a redshift (or redshift probability distribution) for each galaxy. However, these methods are not generally designed to yield unbiased redshift distributions, as they rely on the galaxies used in the training set to be representative of and similarly distributed to the overall population. Because of this, the accuracy of photometric redshifts (or photo- z s) can depend strongly on e.g. the magnitude, redshift and type of a galaxy, and the filters used (e.g. Cunha et al. 2009; Bezanson et al. 2016).

While evolved methods exist that counter these problems (e.g. Lima et al. 2008), one can also choose to avoid photometric redshifts altogether. One such way is to obtain redshift distributions for photometric galaxies statistically by examining how strongly they cluster with sources that have a known redshift. Even if these spectroscopic sources are a biased subset with a very different redshift distribution, they should still trace the same large-scale structure as the overall galaxy population. This means that it is statistically likely for galaxies to be at the same redshift as other sources they cluster strongly with, i.e. if two galaxies are close on the sky then they are more likely to be close along the line of sight. Techniques exploiting

clustering to obtain independent redshift information have been applied for a number of years now to improve and/or characterize the errors of a photo- z catalogue (e.g. Padmanabhan et al. 2007; Newman 2008; Erben et al. 2009; Benjamin et al. 2010; Kovač et al. 2010; Quadri & Williams 2010; Choi et al. 2016), to reconstruct the density field (e.g. Jasche & Wandelt 2012; Cucciati et al. 2016; Malavasi et al. 2016), and to derive redshift distributions from clustering directly (e.g. Matthews & Newman 2010; Schulz 2010; McQuinn & White 2013; Ménard et al. 2013; Morrison et al. 2017). However, since this method is necessarily statistical we lose information on the properties of the galaxies in each redshift bin (although recently efforts have been made to introduce a dependence on colour, see Rahman et al. 2016). Additionally, the resulting distribution is often degenerate with the unknown redshift-dependent bias of the photometric sample, $b(z)$, which has to be removed in some way before the outcomes can be used (e.g. Schmidt et al. 2013).

Inspired by Sheth & Rossi (2010), we extend existing methods to find the number density of galaxies in not only bins of redshift, z , but also apparent magnitude, m . By simultaneously fitting for both distributions, luminosity functions in terms of absolute magnitude, M , can be extracted at different redshifts. This has great potential, as the luminosity function is a key observable of the galaxy population that offers powerful constraints on models of galaxy evolution. Extensive cosmological volumes are needed to measure it accurately, particularly at the bright end where galaxies are rare. Large imaging surveys offer this, but their redshift uncertainties lead to uncertainties in the absolute magnitude of the galaxies. Spectroscopic surveys, on the other hand, have small redshift uncertainties but can probe far fewer galaxies. By cross-correlating these two types of survey while taking the observed brightness of the galaxies into

* E-mail: daalen@strw.leidenuniv.nl

account, we can derive luminosity functions for large volumes with smaller redshift uncertainties than would be possible otherwise. This method also allows us to break degeneracies in a new way: by assuming a simple model for just the luminosity dependence of the galaxy bias, the resulting redshift distributions and luminosity functions are independent of $b(z)$, and no bias removal is necessary.

We present our method for simultaneously deriving redshift distributions and luminosity functions from clustering data in Section 2. As a test, we apply our model to a mock galaxy sample in Section 3. Finally, we summarize our results and discuss the possible limitations of our model when applied to real-world data in Section 4.

2 METHODOLOGY

The way in which we link the redshift distribution dN/dz to the clustering signal can be viewed as a combination of the methods employed by Schulz (2010) and Ménard et al. (2013), although we extend previous efforts by also estimating evolving luminosity functions for the photometric galaxies. Our approach is essentially to apply tomography to the luminosity function: the observed distribution of a sample of galaxies over apparent magnitude, $n(m)$, and the distributions of galaxies over redshift in bins of apparent magnitude, $n_m(z)$, can be viewed as projections of the underlying luminosity function as a function of redshift, $\phi(M, z)$, and therefore used to reconstruct it. An added advantage of fitting for the redshift distributions and luminosity functions simultaneously is that it allows one to make optimal use of the information available in the survey – for example, galaxies that appear bright are unlikely to be at high redshift.

In what follows, subscripts ‘p’ denote the photometric sample for which we aim to derive a distribution in magnitude and redshift, while subscripts ‘s’ denote the spectroscopic sample (which has a known redshift distribution).

2.1 The cross-correlation signal

The number of sample galaxies in apparent magnitude bin m_λ and redshift bin z_i is given by

$$N_p(m_\lambda, z_i) = \int_{z_{i,\min}}^{z_{i,\max}} \int_{m_{\lambda,\min}}^{m_{\lambda,\max}} \frac{dN_p}{dm dz}(m, z) dm dz, \quad (1)$$

where ‘i, min’ and ‘i, max’ denote the edges of bin i . The parameter we wish to extract from the data is the fraction of sample galaxies in apparent magnitude bin m_λ that reside in redshift bin z_i , given by

$$f_N(m_\lambda, z_i) = \frac{N_p(m_\lambda, z_i)}{N_p(m_\lambda)}, \quad (2)$$

where $N_p(m_\lambda)$ is the total number of galaxies in bin m_λ , given by

$$N_p(m_\lambda) = \sum_i N_p(m_\lambda, z_i). \quad (3)$$

The $N_p(m_\lambda)$ of the data are known a priori; however, we do not enforce the $N_p(m_\lambda)$ in our model – which we will refer to as $\tilde{N}_p(m_\lambda)$ – to be identical to these. Rather, we interpret those in the data as being drawn from a Poisson distribution with means given by $\tilde{N}_p(m_\lambda)$ (see Section 2.2).

As our signal we choose the integrated angular cross-correlation function of all photometric galaxies in apparent magnitude bin m_λ

with the spectroscopic galaxies in redshift bin z_i , $\bar{w}_{ps}(m_\lambda, z_i)$, given by

$$\bar{w}_{ps}(m_\lambda, z_i) = \int_{\theta_{\min}}^{\theta_{\max}} w_{ps}(m_\lambda, z_i, \theta) W(\theta) d\theta, \quad (4)$$

where $W(\theta)$ is a weight function. We follow Ménard et al. (2013) in choosing $W(\theta) = \theta^{-1}$, and for the purposes of illustration choose $\theta_{\min} = 0.02$ and $\theta_{\max} = 10$ deg.

We will refer to our model for $\bar{w}_{ps}(m_\lambda, z_i)$ as $\tilde{w}_{ps}(m_\lambda, z_i)$. This quantity is related to the integrated angular correlation function between spectroscopic galaxies in redshift bin z_i and those in redshift bin z_j , $\bar{w}_{ss}(z_i, z_j)$, through

$$\tilde{w}_{ps}(m_\lambda, z_i) = \sum_j f_N(m_\lambda, z_j) \frac{\bar{b}_p(m_\lambda, z_j)}{\bar{b}_s(z_j)} \bar{w}_{ss}(z_i, z_j), \quad (5)$$

where \bar{b} is the (linear) galaxy bias averaged over all scales θ between θ_{\min} and θ_{\max} . Here, we have used that the two samples trace the same underlying density field.

Both \bar{w}_{ps} and \bar{w}_{ss} can be directly calculated from the data (e.g. through pair counting), but the galaxy biases are a priori unknown. However, it is not unreasonable to assume that \bar{b}_p and \bar{b}_s evolve similarly with redshift at fixed luminosity, i.e. $\bar{b}_p(m, z) = \bar{b}_{p,0} b_L(m, z) f(z)$ and $\bar{b}_s(z) = \bar{b}_{s,0} f(z)$.¹ Here, b_L is some function of luminosity – assumed to be known, either independently or determined from the spectroscopic sample – with no residual dependence on m or z .

Next, we recognize that the redshift evolution $f(z)$ of the biases cancels out when taking the ratio, and absorb all constants in a new term. Then, in the limit of infinitely accurate measurements of $\bar{w}_{ps}(m_\lambda, z_i)$ and $\bar{w}_{ss}(z_i, z_j)$, we can derive f_N simply by solving (for all z_i):

$$\tilde{w}_{ps}(m_\lambda, z_i) = \sum_j f'_N(m_\lambda, z_j) \bar{w}_{ss}(z_i, z_j), \quad (6)$$

where $f'_N(m_\lambda, z_j) = K b_L(m_\lambda, z_j) N_p(m_\lambda, z_j) / N_p(m_\lambda)$ with K an unknown constant and a parameter of the model. This set of equations can be written as $\tilde{w}_{ps}(m_\lambda) = \mathbf{X} f'_N(m_\lambda)$, with $\tilde{w}_{ps}(m_\lambda)$ and $f'_N(m_\lambda)$ vectors of length n_z and \mathbf{X} a matrix of size $n_z \times n_z$, where n_z is the number of redshift bins. Hence, $X_{ij} = \bar{w}_{ss}(z_i, z_j)$. Note that we do not assume the often-used Limber (1953) approximation, but allow for non-zero cross-correlations between redshift bins. Even though such cross-correlations are often serendipitous, they contain additional information on the large-scale density field and therefore can offer additional constraints. In Section 3.2.4, we show how our results are affected if these cross-correlations are assumed to be zero.

For the purposes of illustration, we choose the following simple form for the luminosity bias (motivated by, e.g. Benoist et al. 1996; Norberg et al. 2001; Peacock et al. 2001):

$$b_L(m, z) = 1 + \frac{L(m, z)}{L'}, \quad (7)$$

where $L(m, z)$ is the luminosity of a galaxy of apparent magnitude m at redshift z . We set L' to be the luminosity of a galaxy with absolute magnitude $M' = -23.3$. Note that the normalization of the luminosity bias is indirectly controlled by the model parameter K . Since our model is agnostic about the redshift and therefore

¹ Alternatively, $\bar{b}_s(z)$ could be estimated from the data (propagating the observational uncertainties) and $\bar{b}_p(m, z)$ (or the ratio) could be modelled (e.g. as a polynomial in redshift).




Figure 1. The survey volume and the luminosity function combine to form the redshift distribution dN/dz . Shown here is an example for galaxies in an apparent magnitude bin $m = [22, 22.5]$. *Left:* A Schechter (1976) luminosity function, with arbitrary normalization, as a function of absolute magnitude M . Here, $(\alpha, M_*) = (-1.3, -21.1)$. Coloured regions show the absolute magnitudes corresponding to $m = [22, 22.5]$ at redshifts $z = 0.1, 0.35, 0.6, 0.85,$ and 1.1 as indicated in the figure. In this example, we assume the luminosity function is independent of redshift. *Right:* Shown together here are the comoving volume added by each redshift slice, $\Delta V(z)$, the Schechter function shown on the left integrated over the relevant range in M , $\Delta\Phi(M)$, and the redshift distribution resulting from their product, $dN/dz = \Delta\Phi\Delta V$. The integral over one of the highlighted regions in the left-hand panel corresponds to the highlighted height of $\Delta\Phi(M)$ in the right-hand panel.

luminosity of each individual galaxy, we calculate the luminosity bias only once for each bin (m_λ, z_i) , assuming a naive relation between apparent and absolute magnitude (see Section 2.2).

At this point, we could solve the equations given by $\tilde{w}_{ps}(m_\lambda) = \mathbf{X} f'_N(m_\lambda)$ for every m_λ independently to find the corresponding galaxy redshift distributions. However, this disregards the information inherent in the apparent magnitudes of the galaxies. Since the clustering measurements have uncertainty (and since there may be degenerate solutions), this will likely lead to, for example, at least some galaxies with a very bright apparent magnitude being placed at high redshift – corresponding to an unphysically high luminosity. Luminosity functions fitted to these results will therefore be extremely biased and unrealistic.

By fitting to the redshift distribution and the luminosity functions of the sample galaxies simultaneously, we avoid such biased outcomes. This requires us to explicitly model $N_p(m_\lambda, z_i)$.

For conciseness, we will use a subscript notation for binned quantities, i.e. $N_{p,\lambda i} \equiv N_p(m_\lambda, z_i)$, where Greek subscripts always refer to the apparent magnitude bin and Latin subscripts to the redshift bin. Since we fit our model to all bins simultaneously, it is useful to think in terms of superindices $(\lambda i) = n_z \lambda + i$. We will omit the parentheses where it does not lead to confusion.

2.2 A model for $N_p(m, z)$

$N_{p,\lambda i}$ is shaped by the luminosity function, which determines the number density of galaxies at apparent magnitude m_λ and redshift z_i , and the survey volume at redshift z_i . Fig. 1 illustrates how these two quantities combine to form the redshift distribution of galaxies at fixed apparent magnitude. In this example, we assume that the luminosity function and the total number density of galaxies are constant with redshift. We consider galaxies in a fixed apparent magnitude bin, although the principle applies to any magnitude-limited survey. As the survey volume grows with redshift, the number of

galaxies observed per unit redshift increases. However, galaxies with a fixed apparent magnitude correspond to increasingly more-luminous and more-rare galaxies, and so the number density decreases with redshift, first as a power law and then exponentially. The combined result of these two competing effects is a galaxy redshift distribution $dN/dz \propto \Delta\Phi\Delta V$ that increases as a power law before decreasing exponentially.²

Assuming a cosmology fixes the evolution of the survey volume. The shape of the luminosity function at each z then fixes the redshift distribution. Conversely, knowing both the cosmology and the redshift distribution at several fixed apparent magnitudes gives us information on the shape of the luminosity function through cosmic time.

The comoving distance (for a flat Λ CDM universe) is given by

$$d_c(z) = \int_0^z \frac{c}{H_0 \sqrt{\Omega_{m,0}(1+z')^3 + \Omega_{\Lambda,0}}} dz', \quad (8)$$

and hence the volume in redshift bin i given by

$$\begin{aligned} V_i &= \int_A \int_{d_{i,\min}}^{d_{i,\max}} d_c(z)^2 dd_c(z) dA \\ &= f(A) 4\pi \int_{z_{i,\min}}^{z_{i,\max}} \frac{d_c(z)^2 c}{H_0 \sqrt{\Omega_{m,0}(1+z)^3 + \Omega_{\Lambda,0}}} dz, \end{aligned} \quad (9)$$

where A is the area of the sky the survey covers and $f(A)$ is the fraction of the sky (in units of steradians) covered, and where the limits of integration $z_{i,\min}$ and $z_{i,\max}$ are the minimum and maximum redshift values, respectively, of redshift bin i .

² In the case of an evolving luminosity function, the integral over the sky and the luminosity function do not separate out as neatly as in this example, but the end result is similar. We do not make the assumption of a redshift-independent luminosity function beyond this example.

For the purposes of illustration, we will assume the luminosity function is described well by a single Schechter function. We further assume that its parameters α (the low-luminosity power-law slope) and M_* (the turn-over absolute magnitude) evolve linearly with redshift, that is $\alpha = \alpha_0 + \alpha_e z$ and $M_* = M_{*0} + M_{*e} z$. The normalization of the luminosity function is allowed to evolve with redshift as well; specifically, we model it as the exponential of a fifth-order polynomial as follows:

$$\phi_*(z) = \exp\left(\sum_{j=0}^{j=5} \zeta_j \left[\frac{2z}{z_{\max}} - 1\right]^j\right), \quad (10)$$

with z_{\max} the maximum redshift considered and six free parameters ζ_j .³ Our luminosity function thus has 10 free parameters in total. We note that the luminosity function can be straightforwardly generalized to include, e.g. additional Schechter terms or a more (or less) sophisticated redshift evolution.

To avoid divergence (and because there exists a minimum luminosity to what is considered a galaxy), we define a limiting galaxy absolute magnitude M_{lim} . In this study, we set $M_{\text{lim}} = -16$, but we note that any sufficiently dim value of M_{lim} does not influence the outcome of the model. The (integrated) number density of galaxies in apparent magnitude bin m_λ and redshift bin z_i is then

$$\begin{aligned} \Phi_{\lambda i} &= \frac{2}{5} \ln(10) \int_{z_{i,\min}}^{z_{i,\max}} \phi_*(z) \\ &\times \int_{M_1}^{M_2} \frac{10^{\frac{2}{5}(M_*(z)-M(m,z))(\alpha(z)+1)} e^{-10^{\frac{2}{5}(M_*(z)-M(m,z))}}}{\Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M_{\text{lim}})})} dM dz, \end{aligned} \quad (11)$$

where $M(m, z)$ is the absolute magnitude corresponding to a galaxy with apparent magnitude m at redshift z . Here, we have defined the number density such that $\int_{-\infty}^{M_{\text{lim}}} \frac{d\Phi_{\lambda i}}{dM_{\lambda i}} dM \equiv \int_{-\infty}^{M_{\text{lim}}} \phi_i(M) dM = \int \phi_*(z) dz_i$. The limits of integration for M in equation (11) are determined by the edges of the bins m_λ and z_i , but the former are bounded above by M_{lim} . That is, $M_1 = \min\{M(m_{\lambda, \min}; z), M_{\text{lim}}\}$ and $M_2 = \min\{M(m_{\lambda, \max}; z), M_{\text{lim}}\}$.

To convert between absolute and apparent magnitudes, we make the simplified assumption of a flat galaxy spectrum, in which case the K -correction is zero and $M(m, z) = m + 5[1 - \log_{10}(d_L(z))]$, with $d_L(z)$ the luminosity distance. As we show in Appendix C, the K -correction for the galaxies in our sample is typically very small. However, for faint high-redshift galaxies the difference between the real and naively derived magnitude can be up to ~ 1.5 dex, and so K -corrections should be included in the model before applying it to real data.

To account for evolution of the different functions within each redshift bin, we simultaneously integrate the volume and the luminosity function. The expected (Poisson mean) number of galaxies in apparent magnitude bin m_λ and redshift bin z_i is then given by⁴

$$\begin{aligned} \tilde{N}_{p, \lambda i} &= \int_{z_{i,\min}}^{z_{i,\max}} \int_{m_1}^{m_2} \frac{d\phi_i(M)}{dz_i} \frac{dV}{dz} dm dz \\ &= \frac{2}{5} \ln(10) B \int_{z_{i,\min}}^{z_{i,\max}} \frac{d_c(z)^2 \phi_*(z)}{\sqrt{\Omega_{m,0}(1+z)^3 + \Omega_{\Lambda,0}}} \\ &\times \int_{m_1}^{m_2} \frac{10^{\frac{2}{5}(M_*(z)-M(m,z))(\alpha(z)+1)} e^{-10^{\frac{2}{5}(M_*(z)-M(m,z))}}}{\Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M_{\text{lim}})})} dm dz, \end{aligned} \quad (12)$$

where some of the constants have been absorbed into the constant B ; specifically, $B = 4\pi f(A) c/H_0$. We have switched the integral over M to an integral over m , but similar to before, $m_1 = \min\{m_{\lambda, \min}; m(M_{\text{lim}}, z)\}$ and $m_2 = \min\{m_{\lambda, \max}; m(M_{\text{lim}}, z)\}$. The integral over apparent magnitude has an analytical solution, and so we can reduce the above expression for the Poisson mean to an integral over only the redshift bin z_i :

$$\begin{aligned} \tilde{N}_{p, \lambda i} &= B \int_{z_{i,\min}}^{z_{i,\max}} \frac{d_c(z)^2 \phi_*(z)}{\sqrt{\Omega_{m,0}(1+z)^3 + \Omega_{\Lambda,0}}} \\ &\times \left[\Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M_{\text{lim}})}) \right]^{-1} \\ &\times \left[\Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M(m_2, z))}) \right. \\ &\left. - \Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M(m_1, z))}) \right] dz. \end{aligned} \quad (13)$$

The total number of model galaxies in apparent magnitude bin m_λ at any redshift is then

$$\begin{aligned} \tilde{N}_{p, \lambda} &= \sum_i \tilde{N}_{p, \lambda i} \\ &= B \int_{z_{\min}}^{z_{\max}} \frac{d_c(z)^2 \phi_*(z)}{\sqrt{\Omega_{m,0}(1+z)^3 + \Omega_{\Lambda,0}}} \\ &\times \left[\Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M_{\text{lim}})}) \right]^{-1} \\ &\times \left[\Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M(m_2, z))}) \right. \\ &\left. - \Gamma(\alpha(z)+1, 10^{\frac{2}{5}(M_*(z)-M(m_1, z))}) \right] dz. \end{aligned} \quad (14)$$

Here, z_{\min} and z_{\max} are the limits of the redshift range probed by the spectroscopic sample. These model estimates of the mean can be directly compared to the $N_{p, \lambda}$ of the data as a measure of our model's accuracy for a given set of parameters.

2.3 Fitting the model

Using 11 free parameters in total (1 parameter for the bias ratio, 6 for the normalization of the luminosity function, and 4 for its shape parameters), our model predicts a distribution of galaxies in both absolute magnitude and redshift, and – using the observed integrated autocorrelation of the spectroscopic sample $\bar{w}_{ss, ij}$ – the corresponding cross-correlation signal $\tilde{w}_{ps, \lambda i}$. The best-fitting set of parameters is determined by comparing the model outcomes $\tilde{w}_{ps, \lambda i}$ and $\tilde{N}_{p, \lambda}$ to their observed counterparts. We fit for these two quantities simultaneously by minimizing

$$\chi^2 = (\bar{w}_{ps} - \tilde{w}_{ps})^T \mathbf{C}^{-1} (\bar{w}_{ps} - \tilde{w}_{ps}) + R \sum_{\lambda} \frac{(N_{p, \lambda} - \tilde{N}_{p, \lambda})^2}{\tilde{\sigma}_{\lambda}^2}, \quad (15)$$

where \mathbf{C} is a joint covariance matrix combining different sources of uncertainty in both the data and the model (see Appendix A), R is a

³ The number of parameters used to fit $\phi_*(z)$ should be high enough to allow enough versatility, but much smaller than the number of redshift bins to ensure that it varies smoothly and that no (additional) degeneracies are introduced. We found that using a fifth-order polynomial strikes a nice balance. This particular form for $\phi_*(z)$ was chosen for numerical reasons (e.g. an easily calculable derivative).

⁴ We note here that we ignore the modulation of observed galaxy number densities due to lensing magnification, which causes a magnification bias.

constant determining the relative weight of the two observables, and $\tilde{\sigma}_\lambda^2$ is the variance of $\tilde{N}_{p,\lambda}$. Since $\tilde{N}_{p,\lambda}$ is a Poisson mean, $\tilde{\sigma}_\lambda^2 = \tilde{N}_{p,\lambda}$. The ideal value of R is unknown, but it should be set such that $\tilde{N}_{p,\lambda}$ is not fit at the expense of $\tilde{w}_{ps,\lambda i}$, but rather used to break degeneracies in the clustering. In what follows, we set $R = n_z$, so as to give the cross-correlation signal and galaxy number counts equal weight (after all, the former provides $n_z \times n_m$ data points while the latter only provides n_m). When the Limber approximation is taken, we set $R = 1$, since in this case the effective number of data points obtained from the clustering signal goes down by a factor of n_z . Very similar results are obtained if we vary R within a factor of 10.

3 TESTING THE MODEL

3.1 Mock catalogues

To test our model, we extract a mock galaxy survey from one of the publicly available Planck Millennium all-sky lightcones released with Henriques et al. (2015).⁵ The semi-analytical model that forms the base for this lightcone is detailed in Henriques et al. (2015), while information on how the lightcone was constructed and magnitudes were assigned can be found in Henriques et al. (2012). In order to measure our model’s performance, we have to know the luminosity function of the data. A potential mismatch in our final results may be due to either inaccuracies in the model or to the fact that a single Schechter function is not a perfect fit to the intrinsic luminosity function of the mock galaxies. In order to separate these effects, we reassign the absolute magnitude of each galaxy (in the i band) so that it is consistent with an input luminosity function. This is done in redshift bins 0.01 wide, and in such a way that the rank ordering of galaxies in brightness in each redshift bin is preserved (i.e. the N brightest galaxies at each redshift before reassignment are still the N brightest galaxies after reassignment, for every N). We choose the shape parameters of our Schechter function to be $\{\alpha_0, \alpha_e, M_{*0}, M_{*e}\} = \{-1.01, -0.15, -21.5, -0.8\}$ (see Section 2.2). These parameters were chosen in fair approximation of the intrinsic luminosity function of the galaxies in the lightcone. We do not change the redshifts, locations or number densities as a function of redshift of the galaxies; hence, the normalization of the luminosity function and the clustering bias of the galaxies is still determined by the processes that formed them. The apparent magnitude of each galaxy is recalculated to match its new absolute magnitude, assuming again a naive relation between these and redshift (i.e. without K -corrections). We then make cuts in apparent magnitude and redshift, only keeping galaxies for which $m \leq 21$ and $z \leq 0.8$. Next, we arbitrarily select the region with right ascension within $[100, 200]$ deg and declination within $[10, 50]$ deg, equivalent to 3394 deg^2 or about 8 per cent of the sky.

The galaxies that are left comprise our photometric galaxy sample. From it, we select a spectroscopic sample by selecting the N brightest galaxies in each redshift bin with stellar masses $M_* \geq 10^{10} h^{-1} M_\odot$ and star formation rates $\dot{M}_* \geq 1 h^{-1} M_\odot \text{ yr}^{-1}$, where N is chosen such that the number density of spectroscopic galaxies is at most $10^{-4} (\text{Mpc}/h)^{-3}$ at every redshift.⁶ Note that for

the purpose of demonstrating the effectiveness of our model, all that matters is that the spectroscopic sample is a small and highly biased subset of the total population, not that it is realistically selected. For the photometric sample, we retain only the position on the sky and apparent magnitude. Finally, we take $n_z = 16$ redshift bins, $\Delta z = 0.05$ wide, in the range $z = [0, 0.8]$, and $n_m = 16$ apparent magnitude bins, $\Delta m = 0.5$ wide, in the range $m = [13, 21]$, and calculate the relevant (cross-)correlation functions and covariance matrices. In total, our photometric sample contains 14 280 584 galaxies that fall in these ranges, and our spectroscopic sample contains 250 372 sources.

We note that the spectroscopic sample is (realistically) a biased subset of the galaxy population. As we show in Fig. 2, the spectroscopic galaxies have a radically different redshift distribution and only probe the most luminous end of the total luminosity function. However, since both samples still trace the same large-scale distribution, and since the bias ratio of the two samples is a free parameter in the model, this is not an issue in our approach. Indeed, Scottez et al. (2016) recently showed that for the similar methodology of Ménard et al. (2013), accurate redshift distributions can be obtained for galaxies fainter than those of the spectroscopic sample. Our own results in the following section confirm this.

3.2 Results

3.2.1 Fiducial model

Using only the clustering amplitude of the spectroscopic sources and the total distribution of photometric galaxies over apparent magnitude, our model is able to reproduce the input luminosity function of the mock catalogue to very high accuracy. The results for our fiducial model are shown in Fig. 3. Since error bars on the data shown or the model are not straightforwardly calculated, due to the many interrelated sources of uncertainty, we instead show just the 1σ variation due to cosmic variance on the data as lightly shaded bands. This was calculated from 1000 randomly placed surveys of the lightcone catalogue, each with the same sky area as our fiducial survey area.

In the top-left panel, we show the luminosity functions as a function of absolute magnitude in for each redshift bin. Solid lines show the luminosity function as measured directly from the mock catalogue with full redshift information, thereby including realization noise (which plays a significant role in the first two redshift bins). At low redshift – specifically in the first two redshift bins – the model tends to overestimate the number of dim galaxies, although we note that the difference is in large part due to cosmic variance, as we will show. For the highest redshift bin, too, the model slightly overestimates the number of galaxies observed. Even with these caveats, in most regimes the luminosity function of the mock galaxies are very accurately reproduced by the best-fitting galaxy distribution, including the bright and dim end drop-offs. The latter is due to the cut-off apparent magnitude shifting to brighter galaxies within each redshift bin, and therefore only captured when the model luminosity function and volume are integrated together (see equation 12).

We show the distribution over apparent magnitude for each redshift bin in the top right-hand panel. The total distribution is shown in black and is used as a constraint in the model to break the

⁵ Specifically, we use the catalogues ‘cones.AllSky_M05_001’ and ‘MRsc-Planck1’ from the ‘Henriques2015a’ part of the Millennium public data base.

⁶ To clarify, after our pre-selection by stellar mass and star formation rate, we choose the N galaxies at each redshift that are brightest in absolute magnitude. One might argue that a more natural choice is apparent magnitude;

however, since we make our selection in separate redshift bins, the difference is minimal.

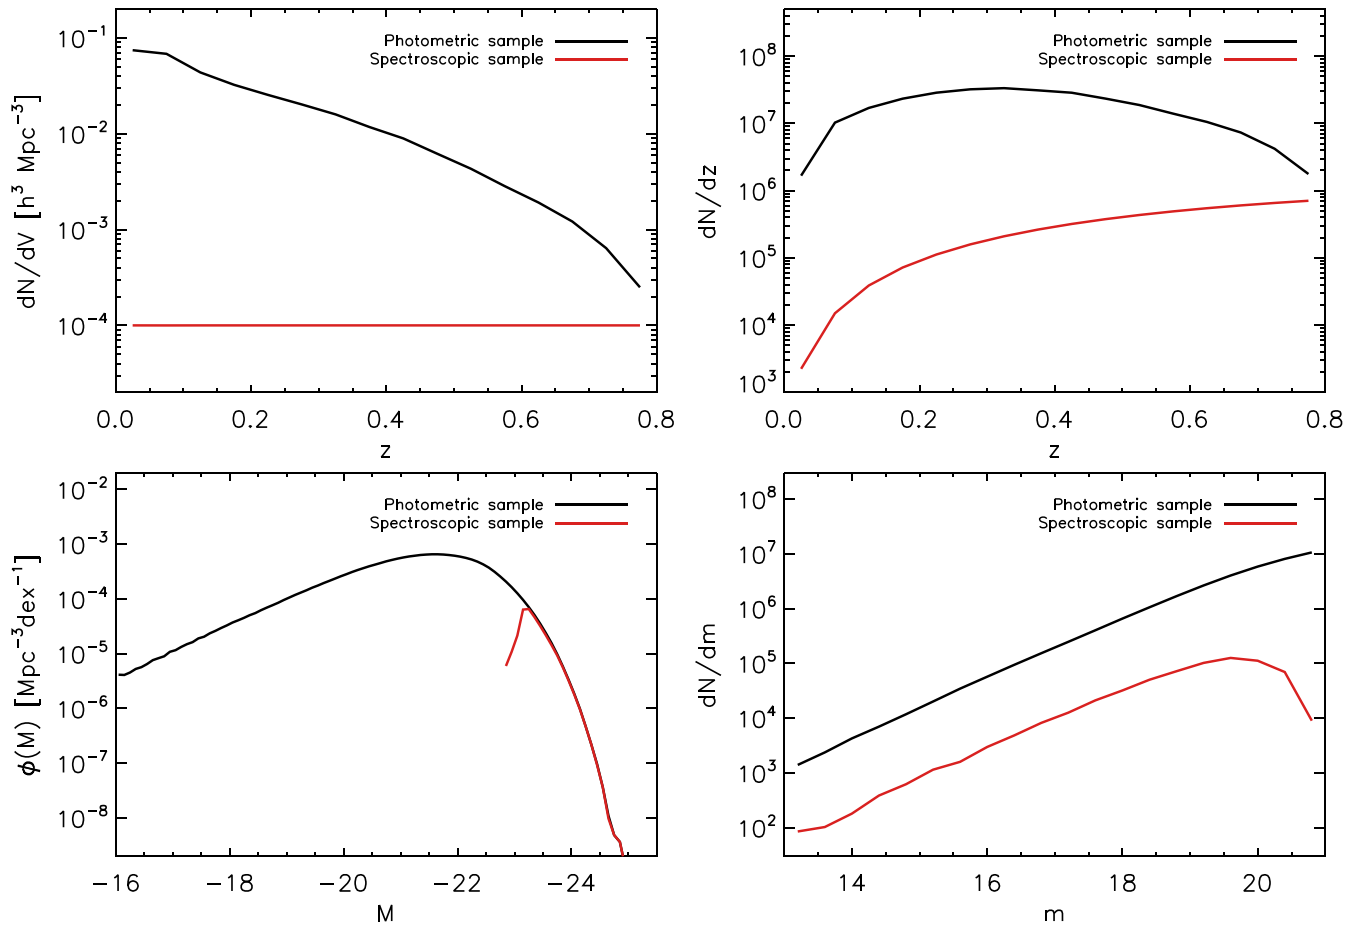


Figure 2. A comparison of the distributions of our photometric (in black) and spectroscopic (in red) mock galaxy samples over cosmic (comoving) volume (top left), redshift (top right), absolute magnitude (bottom left), and apparent magnitude (bottom right). By construction, the spectroscopic sample has a spatial density of $10^{-4} (\text{Mpc}/h)^{-3}$ over the entire redshift range and contains only the most luminous (star-forming) galaxies. Even though the spectroscopic sample is, realistically, a biased subset of the total galaxy population, it can still be used to derive accurate redshift distributions and luminosity functions for the photometric galaxies, as they trace the same large-scale structure and the clustering bias of the samples does not need to be known in our model.

clustering degeneracies (see equation 15). The model again tends to overestimate the number of dim galaxies in the lowest redshift bin, where the cosmic variance is largest and the clustering signal has a relatively large uncertainty. Overall, though, the model does very well in reproducing the true distribution of galaxies in apparent magnitude, at any redshift.

The bottom left-hand panel of Fig. 3 shows the redshift distributions in each apparent magnitude bin, as well as the total. Note that we are showing the absolute number of galaxies assigned to each redshift bin. The clustering model does an excellent job at reproducing these, even for the bright galaxies with relatively low number densities. As before, the fit is particularly accurate at intermediate redshifts (for all apparent magnitudes), where most of the galaxies in our sample reside and therefore where the uncertainty on the constraints is smallest.

Finally, in the bottom right-hand panel, we show the normalization of the luminosity function as a function of redshift. Black crosses show the effective normalization of the mock galaxies in the survey area in each redshift bin. The dashed red line shows the fit (see equation 10) that best reproduced the clustering data, with red crosses showing its volume-averaged values in each redshift bin to allow for a more direct comparison to the input data. The fit captures the shape of the input data, even if it tends to

overestimate the normalization. However, due to the degeneracies between different Schechter parameters,⁷ a mismatch in the value of the normalization parameter does not necessarily mean that the luminosity function itself is not accurately reproduced, as the other panels show.

If we compare the normalization measured for our catalogue to the shaded band showing the 1σ range of cosmic variance, we see that our survey area contains significantly less galaxies than average in the first redshift bin, and significantly more than average in the second redshift bin. This uncommon feature is the main reason why our model has trouble matching the measured number densities in these redshift bins. Also note the sharp downturn to lower number densities observed for the very highest redshift bin, which is not fully captured by our fit, causing the model to overestimate the number of galaxies in that bin.

⁷ One easily seen example of such a degeneracy is between the high-redshift normalization and the slope parameters of the Schechter function. At high redshifts, galaxies above the knee ($M > M_*$) are not or barely probed, as they are too dim to observe, and so in this regime, the slope parameters only serve to normalize the profile.

Fiducial model

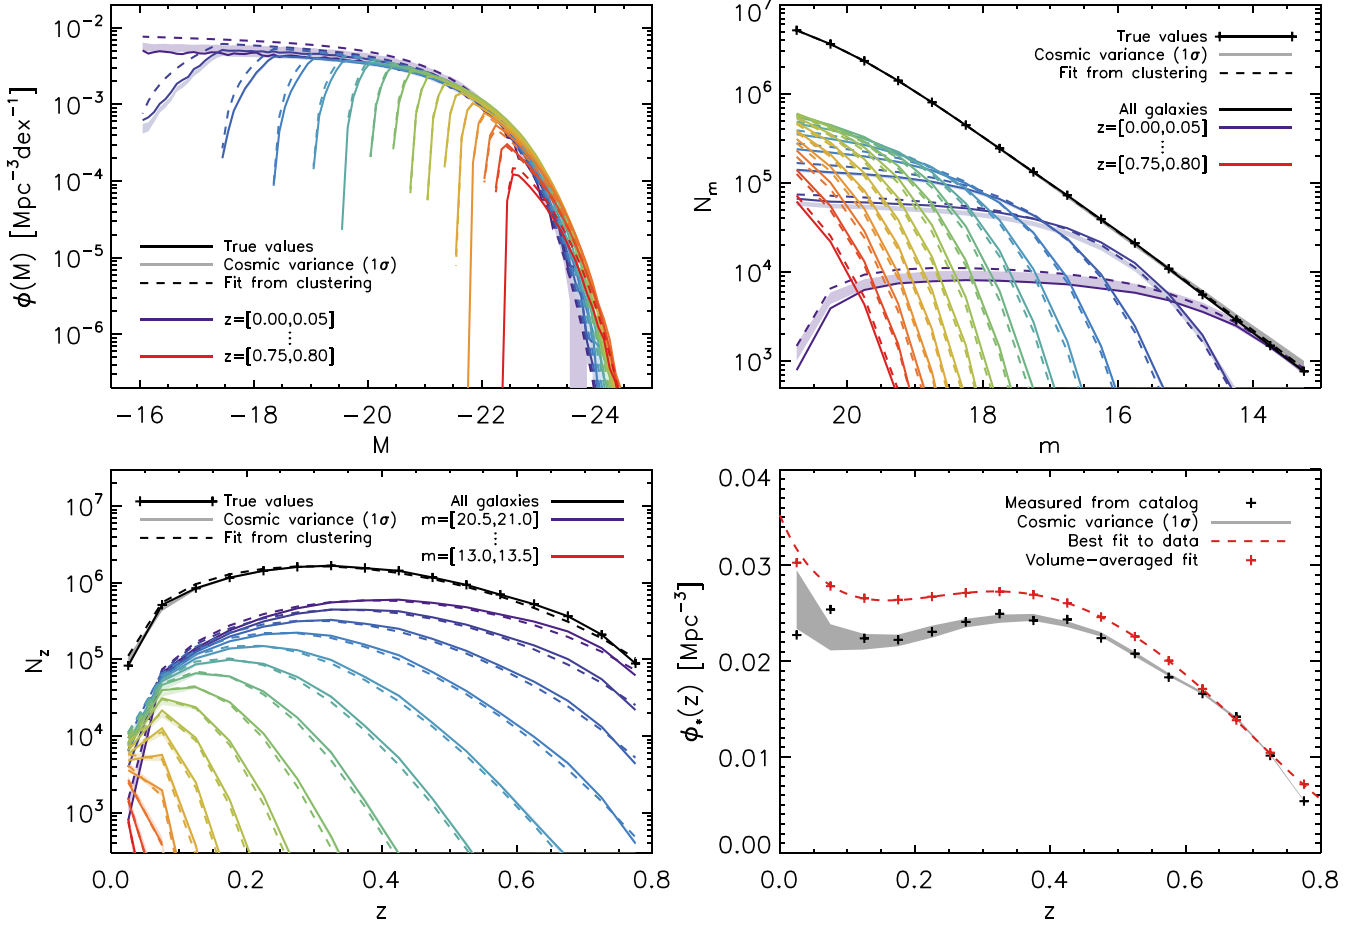


Figure 3. The results for our fiducial model, minimized using equation (15). The model is constrained by two sets of data, one being the cross-correlation signal between photometric and spectroscopic galaxies (in bins of the apparent magnitude of the former and the redshift of the latter), the other being the total number of photometric galaxies in each bin of apparent magnitude. Lighter shaded bands show the effect of cosmic variance (see the main text). *Top left:* the number of galaxies in different bins of redshift as a function of absolute magnitude. Solid lines show the data, and dashed lines show the outcome of the model. Note that the power-law part of the Schechter function is only probed by low-redshift galaxies. Overall, the luminosity function of the data is reproduced very well. For the first redshift bin, where the deviation between the derived and true galaxy densities is largest, the vertical offset is in large part due to cosmic variance for the sky area we are using here (see the main text). *Top right:* the number of galaxies in different bins of redshift as a function of apparent magnitude. The total overall redshifts, shown by the black line, is one of the constraints of the model. *Bottom left:* the number of galaxies in different bins of apparent magnitude as a function of redshift. Black lines show the total overall apparent magnitudes. *Bottom right:* the normalization of the Schechter luminosity function as a function of redshift. The normalizations as inferred from the mock catalogue are shown as black crosses, while what the best-fitting model prefers is shown as a red dashed line. Red crosses show the result of volume averaging the best fit over each redshift bin. Note that degeneracies exist between $\phi_*(z)$ and the other parameters of the Schechter function, which is why the number densities of the galaxies can be reproduced quite well for different sets of parameters.

3.2.2 Direct maximum-likelihood fit

To show that the mismatch at low redshift is indeed not due to the clustering signal or shortcomings of our clustering model, we show in Fig. 4 the results of performing a maximum-likelihood fit directly to the absolute magnitudes and redshifts of the galaxies in the survey catalogue, both of which our fiducial model is agnostic about. We do not bin the data here, instead using the individual M and z of each galaxy as input to the maximum-likelihood function (see Appendix B). Even in this case, the number of galaxies at low redshift is overpredicted due to realization noise (which includes cosmic variance). Comparing Figs 3 and 4, we see that the result of our fiducial model is extremely close to the maximum-likelihood luminosity function, showing the power of using the cross-correlation signal even without any prior redshift information. Additionally,

this shows that the cumulative impact of binning, uncertainties in the clustering data, and perhaps most significantly our assumptions regarding the clustering bias is small.

Adding more parameters to the luminosity function, for example including a second Schechter function or higher-order terms in the normalization, would allow us to compensate for the realization noise and possibly yield a better match to the data. However, doing so would also introduce additional degeneracies.

3.2.3 Fixed slope at low redshift

Our fiducial model has no prior information on the parameters of the luminosity function. However, it is not unreasonable to assume that the power-law slope of the luminosity function at redshift zero, α_0 , is well constrained. To see how much the model outcome is

Direct maximum-likelihood fit




Figure 4. As Fig. 3, but now the model has been replaced by a maximum-likelihood fit to the absolute magnitudes and redshifts of the mock galaxies, which are (realistically) inaccessible to the fiducial model. The luminosity functions derived here are extremely close to those of the fiducial model, showing that most of the already slight mismatch in Fig. 3 is not due to the clustering signal or our clustering model, but to realization noise and the limitations of the parametrization of the luminosity function.

influenced by the uncertainty at low redshift, we therefore also ran our model with α_0 fixed to the input value. The results of this test are shown in Fig. 5. As expected, the panels show a marginal improvement at low redshift in comparison to the results for our fiducial model, but our results at high redshift are slightly worse than before. This is again because of the unusually large realization noise at low redshift: as one parameter is held fixed, the model loses some freedom to compensate for this, which in this case leads to a mismatch at high redshift.

3.2.4 Limber approximation

Finally, we have also tested the consequences of assuming the often-used Limber approximation, by setting the clustering signal (and its covariance) to zero for the cross-correlations of spectroscopic sources in different redshift, the results of which are shown in Fig. 6. In this case, the model performs less well in regimes where the cross-correlations between different redshift bins contribute significantly – that is, at both the low and high redshift ends, and for the brightest galaxies, which have relatively low number densities. At the lowest redshift, depending on the choice of θ_{\max} (see Section 2.1) the typical distance between galaxies may be larger than the distances probed by the clustering signal, and so no or barely

any clustering is observed. Without the information contained in the cross-correlation signal between these and higher-redshift bins, the model therefore prefers to place as few galaxies as possible at low redshift. At high redshift, depending on the choice of θ_{\min} the scales probed may be larger than the scales on which those galaxies cluster strongly, and so a weak signal with a relatively large uncertainty is observed. Increasing θ_{\max} /decreasing θ_{\min} gives better results at low/high redshifts but increases the uncertainty at higher/lower redshifts. It is, therefore, best to not take the Limber approximation but make use of all available information. If the Limber approximation has to be taken, it is better to calculate the clustering at a fixed physical scale instead of a fixed angular scale (e.g. Schulz 2010).

For completeness, we present the best-fitting Schechter parameters corresponding to all figures in this section in Table 1. Note that the reproduced luminosity functions can be quite accurate even when the parameters are not, because of the degeneracies of some of these parameters with the normalization.

4 DISCUSSION

The methods presented in this paper extend previous work by not only deriving the redshift distribution of photometric sources through clustering but also deriving their luminosity function

Known local slope

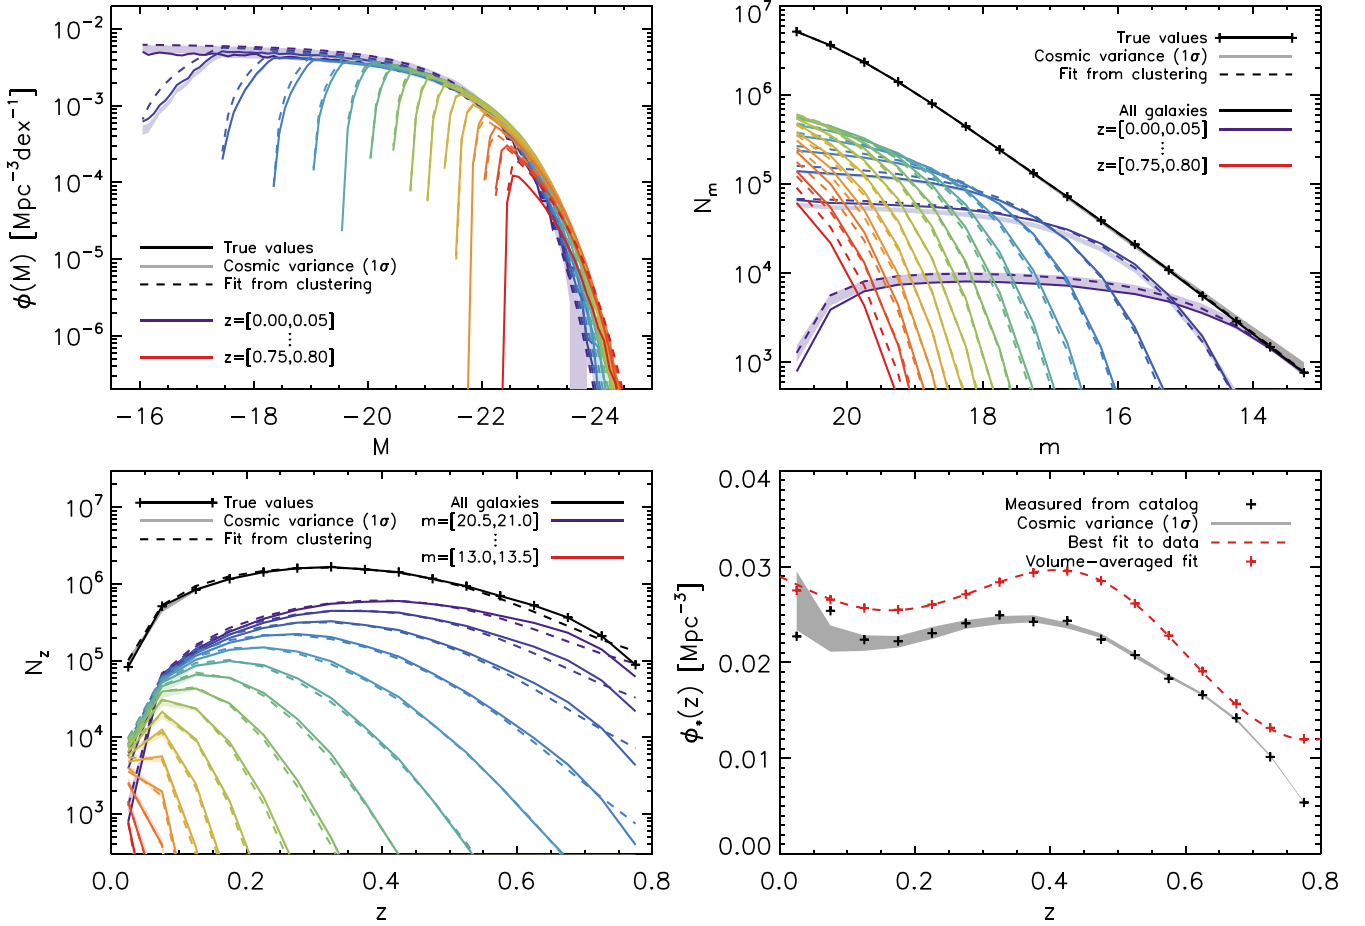


Figure 5. As Fig. 3, but now α_0 , the power-law slope of the luminosity function at $z = 0$, was assumed known and held fixed to the input value in the fitting. As expected, this marginally improves the fit at low redshift, but at the cost of some model freedom that is felt mainly at high redshift, where the fit worsens with respect to the fiducial model.

through cosmic time. By testing this method on a mock galaxy survey, we have demonstrated that an input galaxy distribution over redshift and luminosity can be very accurately recovered in this way for large surveys, even when these are relatively shallow. The redshift distributions derived in this way are not biased by having the spectroscopic sources be selected differently from the photometric sources. As we have shown, the method returns accurate distributions and luminosity functions even if the only galaxies with spectra are the brightest members of the sample and their number densities have a vastly different redshift evolution, so long as they are in the same area of sky. Additionally, our results are not degenerate with the unknown redshift-dependent galaxy bias, $b(z)$.

Our goal has been to introduce a technique for measuring the luminosity function from the co-spatial combination of a deep imaging survey and a sparse spectroscopic survey and to illustrate its potential. The performance of our simple algorithm on mock data is sufficiently encouraging that further development appears warranted. In particular, application to real data would need to consider the possible effects of lensing magnification and incorporate K -corrections in the conversion between apparent and absolute magnitudes (see Appendix C for more on this). Additionally, in this paper we have taken the following assumptions, which should be kept in mind and modified where necessary:

(i) First of all, we have assumed that the form of the luminosity function is known (in our case, a single Schechter function), which in real surveys may not be the case. However, one generally finds that a sum of Schechter functions is a good fit to real data (e.g. Peng et al. 2010). Additionally, the form of the luminosity function that one assumes in this formalism can be very versatile, and is allowed to contain many parameters to be constrained at once. We therefore do not anticipate this to be an issue in the application of the model.

(ii) Secondly, we have assumed a simple luminosity bias relation (equation 7) with a known parameter L' . We have also assumed that the redshift evolution of the remaining bias terms cancel out. However, we have imposed neither bias relation on the mock data, and our results imply these assumptions were sufficiently valid. There is no reason to assume, therefore, that the same would not apply to real data – except perhaps if the clustering bias in the real data had some residual dependence on redshift and/or magnitude that the mock data do not. Any potential scale dependence of the clustering bias (insofar not already implicitly included in the mock data) is not expected to be important, as the bias in our model is an effective one, averaged over a large range in scales. Finally, while the value of L' was fit to a subset of the data prior to running the model, it could in principle be a free parameter constrained by the model.

Limber approximation

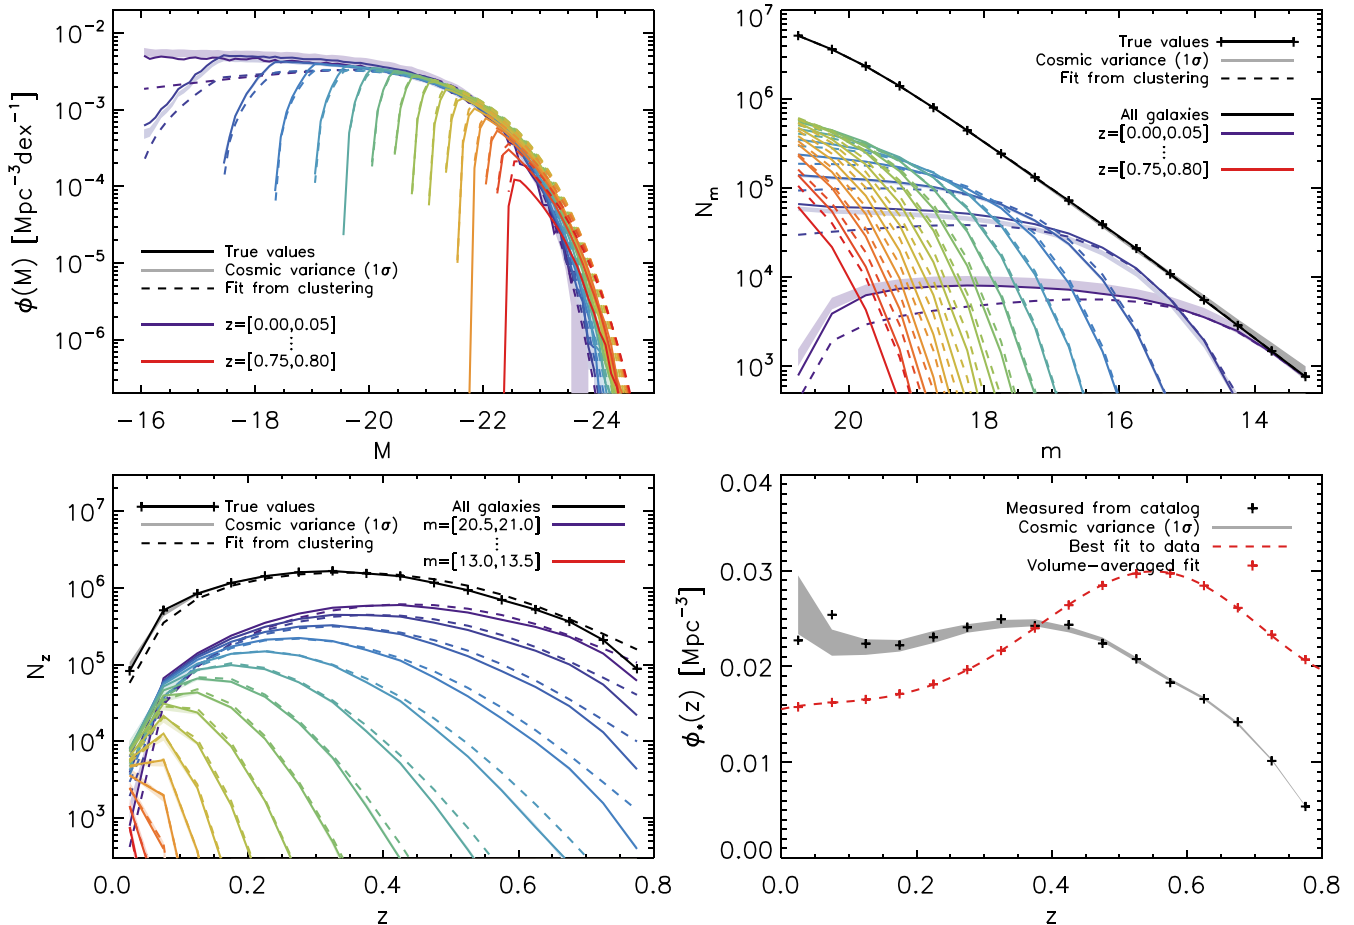


Figure 6. As Fig. 3, but now with the Limber approximation taken. While the fit is still good at intermediate redshifts, both the luminosity function and dN/dz are considerably less well reproduced at both low and high redshifts. This is due to the small autocorrelation clustering amplitude and relatively high uncertainty at these redshifts, meaning contributions from cross-correlations between different redshift bins – which are ignored in the Limber approximation – are relatively more important.

Table 1. The best-fitting luminosity function parameters derived from the clustering data for each of our model runs. Parentheses indicate that the parameter was held fixed to this value. In the run labelled ‘Direct’ no clustering information was used, and instead a maximum-likelihood fit to the galaxies absolute magnitudes and redshifts was performed. Note that the luminosity function may be highly accurately reproduced even for parameters other than the input parameters, due to degeneracy with the normalization and realization noise (including cosmic variance).

Run	α_0	M_{*0}	α_e	M_{*e}
Input	-1.01	-21.5	-0.15	-0.8
Fiducial	-1.050	-21.429	-0.155	-0.927
Direct	-1.019	-21.520	-0.178	-0.783
Fixed α_0	(-1.01)	-21.415	-0.337	-1.061
Limber	-0.747	-21.349	-0.854	-1.344

(iii) Thirdly, as we mentioned in Section 2.3, it is difficult to define an objective value for the relative weight R of the two terms in our model’s χ^2 in equation (15). Fortunately, the outcome of the model turns out not to be very sensitive to its value.

We plan to test our method on a large catalogue of observed galaxies in a follow-up publication.

ACKNOWLEDGEMENTS

The authors thank Joanne Cohn for helpful discussions and Yu Feng for the use of his clustering code. We also thank the anonymous referee for their useful comments, which have led to the improvement of this manuscript. This work was supported in part by the Theoretical Astrophysics Center at UCB. The Millennium Simulation data bases used in this paper and the web application providing online access to them were constructed as part of the activities of the German Astrophysical Virtual Observatory.

REFERENCES

- Benjamin J., van Waerbeke L., Ménard B., Kilbinger M., 2010, MNRAS, 408, 1168
 Benoist C., Maurogordato S., da Costa L. N., Cappi A., Schaeffer R., 1996, ApJ, 472, 452
 Bezanson R. et al., 2016, ApJ, 822, 30
 Choi A. et al., 2016, MNRAS, 463, 3737

- Cucciati O., Marulli F., Cimatti A., Merson A. I., Norberg P., Pozzetti L., Baugh C. M., Branchini E., 2016, MNRAS, 462, 1786
- Cunha C. E., Lima M., Oyaizu H., Frieman J., Lin H., 2009, MNRAS, 396, 2379
- Erben T. et al., 2009, A&A, 493, 1197
- Henriques B. M. B., White S. D. M., Lemson G., Thomas P. A., Guo Q., Marleau G.-D., Overzier R. A., 2012, MNRAS, 421, 2904
- 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
- Jasche J., Wandelt B. D., 2012, MNRAS, 425, 1042
- Kovač K. et al., 2010, ApJ, 708, 505
- Lima M., Cunha C. E., Oyaizu H., Frieman J., Lin H., Sheldon E. S., 2008, MNRAS, 390, 118
- Limber D. N., 1953, ApJ, 117, 134
- Malvasi N., Pozzetti L., Cucciati O., Bardelli S., Cimatti A., 2016, A&A, 585, A116
- Matthews D. J., Newman J. A., 2010, ApJ, 721, 456
- McQuinn M., White M., 2013, MNRAS, 433, 2857
- Ménard B., Scranton R., Schmidt S., Morrison C., Jeong D., Budavari T., Rahman M., 2013, preprint (arXiv:1303.4722)
- Morrison C. B., Hildebrandt H., Schmidt S. J., Baldry I. K., Bilicki M., Choi A., Erben T., Schneider P., 2017, MNRAS, 467, 3576
- Newman J. A., 2008, ApJ, 684, 88
- Norberg P. et al., 2001, MNRAS, 328, 64
- Padmanabhan N. et al., 2007, MNRAS, 378, 852
- Peacock J. A. et al., 2001, Nature, 410, 169
- Peng Y.-j. et al., 2010, ApJ, 721, 193
- Quadri R. F., Williams R. J., 2010, ApJ, 725, 794
- Rahman M., Mendez A. J., Ménard B., Scranton R., Schmidt S. J., Morrison C. B., Budavári T., 2016, MNRAS, 460, 163
- Schechter P., 1976, ApJ, 203, 297
- Schmidt S. J., Ménard B., Scranton R., Morrison C., McBride C. K., 2013, MNRAS, 431, 3307
- Schulz A. E., 2010, ApJ, 724, 1305
- Scottz V. et al., 2016, MNRAS, 462, 1683
- Sheth R. K., Rossi G., 2010, MNRAS, 403, 2137

APPENDIX A: JOINT COVARIANCE MATRIX

There are three sources of uncertainty when fitting our model to the data: uncertainties in the integrated cross-correlation function of photometric and spectroscopic galaxies, $\bar{\mathbf{w}}_{\text{ps},\lambda i} \equiv \bar{\mathbf{w}}_{\text{ps}}(m_\lambda, z_i)$, in the integrated cross-correlation function of spectroscopic galaxies in different redshift bins, $\bar{w}_{\text{ss},ij} \equiv \bar{w}_{\text{ss}}(z_i, z_j)$, and finally in the number of galaxies in the volume at some apparent magnitude and redshift, $\tilde{N}_{\text{p},\lambda i} \equiv \tilde{N}_{\text{p}}(m_\lambda, z_i)$. For the first two, we use 20 000 bootstrap resamplings to calculate full covariance matrices, while the latter is modelled as a Poisson variable with a mean given by the volume-weighted integral over the luminosity function over bins m_λ and z_i . Here, we derive the total covariance matrix, which incorporates the uncertainties from all three sources.

To find the best-fitting model, we aim to minimize χ^2 as given by equation (15), where \mathbf{C} is the joint covariance matrix. As such, \mathbf{C} is a $(n_m n_z) \times (n_m n_z)$ matrix with element $(\lambda i), (\mu j)$ given by

$$\begin{aligned} \mathbf{C}_{(\lambda i)(\mu j)} &= \sigma(\bar{\mathbf{w}}_{\text{ps},\lambda i} - \tilde{\bar{\mathbf{w}}}_{\text{ps},\lambda i}; \bar{\mathbf{w}}_{\text{ps},\mu j} - \tilde{\bar{\mathbf{w}}}_{\text{ps},\mu j}) \\ &= \sigma\left(\bar{\mathbf{w}}_{\text{ps},\lambda i} - \sum_k X_{ik} \mathbf{f}'_{N,\lambda k}; \bar{\mathbf{w}}_{\text{ps},\mu j} - \sum_l X_{jl} \mathbf{f}'_{N,\mu l}\right), \end{aligned} \quad (\text{A1})$$

where $\sigma(A; B)$ denotes the covariance between A and B . Note that \mathbf{C} is symmetric. As before, $X_{ij} = \bar{w}_{\text{ss},ij}$ and $\mathbf{f}'_{N,\lambda i} = K b_{L,\lambda i} \tilde{N}_{\text{p},\lambda i} / \tilde{N}_{\text{p},\lambda}$, where K is a constant, $b_{L,\lambda i}$ is the part of the bias that scales with the luminosity of a galaxy of apparent magnitude

m_λ at redshift z_i (see equation 7) and $\tilde{N}_{\text{p},\lambda} = \sum_i \tilde{N}_{\text{p},\lambda i}$ is the total number of galaxies observed in apparent magnitude bin m_λ . While $\tilde{N}_{\text{p},\lambda}$ and $b_{L,\lambda i}$ are known a priori, K is a parameter of the model. Expanding equation (A1), we find

$$\begin{aligned} \mathbf{C}_{(\lambda i)(\mu j)} &= \sigma(\bar{\mathbf{w}}_{\text{ps},\lambda i}; \bar{\mathbf{w}}_{\text{ps},\mu j}) \\ &\quad - \sum_l b_{L,\mu l} \sigma\left(\bar{\mathbf{w}}_{\text{ps},\lambda i}; \frac{\tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\mu}} \bar{w}_{\text{ss},jl}\right) \\ &\quad - K \sum_k b_{L,\lambda k} \sigma\left(\bar{\mathbf{w}}_{\text{ps},\mu j}; \frac{\tilde{N}_{\text{p},\lambda k}}{\tilde{N}_{\text{p},\lambda}} \bar{w}_{\text{ss},ik}\right) \\ &\quad + K^2 \sum_{k,l} b_{L,\lambda k} b_{L,\mu l} \sigma\left(\frac{\tilde{N}_{\text{p},\lambda k}}{\tilde{N}_{\text{p},\lambda}} \bar{w}_{\text{ss},ik}; \frac{\tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\mu}} \bar{w}_{\text{ss},jl}\right). \end{aligned} \quad (\text{A2})$$

It is clear that $\bar{w}_{\text{ss},ij}$ and $\tilde{N}_{\text{p},\lambda k}$ should be uncorrelated, and we assume the same for $\tilde{N}_{\text{p},\lambda k}$ and $\bar{\mathbf{w}}_{\text{ps},\mu i}$. With this in mind, we can write

$$\sigma\left(\bar{\mathbf{w}}_{\text{ps},\lambda i}; \frac{\tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\mu}} \bar{w}_{\text{ss},jl}\right) = \frac{\tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\mu}} \sigma(\bar{\mathbf{w}}_{\text{ps},\lambda i}; \bar{w}_{\text{ss},jl}), \quad (\text{A3})$$

and

$$\begin{aligned} &\sigma\left(\frac{\tilde{N}_{\text{p},\lambda k}}{\tilde{N}_{\text{p},\lambda}} \bar{w}_{\text{ss},ik}; \frac{\tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\lambda}} \bar{w}_{\text{ss},jl}\right) \\ &= \frac{\tilde{N}_{\text{p},\lambda k} \tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\lambda} \tilde{N}_{\text{p},\mu}} \sigma(\bar{w}_{\text{ss},ik}; \bar{w}_{\text{ss},jl}) \\ &\quad + [\bar{w}_{\text{ss},ik} \bar{w}_{\text{ss},jl} + \sigma(\bar{w}_{\text{ss},ik}; \bar{w}_{\text{ss},jl})] \sigma\left(\frac{\tilde{N}_{\text{p},\lambda k}}{\tilde{N}_{\text{p},\lambda}}; \frac{\tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\mu}}\right). \end{aligned} \quad (\text{A4})$$

All remaining covariances involving the clustering terms are calculated directly through bootstrapping. This leaves only the last covariance in equation (A4). The $\tilde{N}_{\text{p},\lambda i}$ are mutually independent Poisson variables, but are not independent of $\tilde{N}_{\text{p},\mu}$ when $\mu = \lambda$. So

$$\begin{aligned} \mathbf{C}_{(\lambda i)(\mu j)} &= \sigma(\bar{\mathbf{w}}_{\text{ps},\lambda i}; \bar{\mathbf{w}}_{\text{ps},\mu j}) \\ &\quad - \sum_k \mathbf{f}'_{N,\lambda k} \sigma(\bar{\mathbf{w}}_{\text{ps},\mu j}; \bar{w}_{\text{ss},ik}) \\ &\quad - \sum_l \mathbf{f}'_{N,\mu l} \sigma(\bar{\mathbf{w}}_{\text{ps},\lambda i}; \bar{w}_{\text{ss},jl}) \\ &\quad + \sum_{k,l} \mathbf{f}'_{N,\lambda k} \mathbf{f}'_{N,\mu l} \sigma(\bar{w}_{\text{ss},ik}; \bar{w}_{\text{ss},jl}) \\ &\quad + \delta_{\lambda\mu} \sum_{k,l} K^2 b_{L,\lambda k} b_{L,\mu l} [\bar{w}_{\text{ss},ik} \bar{w}_{\text{ss},jk} \\ &\quad + \sigma(\bar{w}_{\text{ss},ik}; \bar{w}_{\text{ss},jk})] \sigma\left(\frac{\tilde{N}_{\text{p},\lambda k}}{\tilde{N}_{\text{p},\lambda}}; \frac{\tilde{N}_{\text{p},\mu l}}{\tilde{N}_{\text{p},\mu}}\right). \end{aligned} \quad (\text{A5})$$

Since $\tilde{N}_{\text{p},\lambda}$ is a sum of independent Poisson variables, and therefore a Poisson distributed variable itself, we need to know the covariance between ratios of dependent Poisson variables in the domain $[0,1]$.

Analytical expressions for this (co)variance and its derivatives can be derived, and the former are given below for completeness. Here, γ is Euler's constant and Ei is the exponential integral function.

If $k = l$:

$$\sigma^2 \left(\frac{\tilde{N}_{p,\lambda k}}{\tilde{N}_{p,\lambda}} \right) = \frac{\tilde{N}_{p,\lambda k}}{\tilde{N}_{p,\lambda}^2} e^{-\tilde{N}_{p,\lambda}} \left\{ \tilde{N}_{p,\lambda k} \left(1 - e^{-\tilde{N}_{p,\lambda}} \right) + \left(\tilde{N}_{p,\lambda} - \tilde{N}_{p,\lambda k} \right) \left(\text{Ei} \left[\tilde{N}_{p,\lambda} \right] - \gamma - \ln \left[\tilde{N}_{p,\lambda} \right] \right) \right\}. \quad (\text{A6})$$

In all other cases,

$$\sigma \left(\frac{\tilde{N}_{p,\lambda k}}{\tilde{N}_{p,\lambda}} ; \frac{\tilde{N}_{p,\lambda l}}{\tilde{N}_{p,\lambda}} \right) = \frac{\tilde{N}_{p,\lambda k} \tilde{N}_{p,\lambda l}}{\tilde{N}_{p,\lambda}^2} e^{-\tilde{N}_{p,\lambda}} \times \left(1 - e^{-\tilde{N}_{p,\lambda}} - \text{Ei} \left[\tilde{N}_{p,\lambda} \right] + \gamma + \ln \left[\tilde{N}_{p,\lambda} \right] \right). \quad (\text{A7})$$

APPENDIX B: DIRECT LIKELIHOOD FUNCTION

In order to test our clustering-based approach to finding the galaxy redshift distribution and luminosity function, as well as test the cumulative impact of binning, the uncertainties of the clustering data and our model choices regarding the clustering data (e.g. the bias model), in Section 3.2.2 we considered the luminosity function one would obtain when doing a direct maximum-likelihood fit to the individual absolute magnitudes and redshifts of the galaxies. The likelihood function we maximized is constructed as follows.

We assume that the set of observed galaxies is a Poisson realization, with Poisson means determined by an underlying luminosity function and cosmology (see equation 13). Let us now consider these Poisson means in bins in z and m that are sufficiently small such that each contains at most one galaxy, and index these bins with j (previously λi). If μ_j is the Poisson mean for the apparent magnitude and redshift corresponding to j , N_j is the number of galaxies in this bin, and \mathbf{p} is a vector of all parameters, then the likelihood is given by

$$\mathcal{L}(\mathbf{p}) = \prod_j \frac{\mu_j(\mathbf{p})^{N_j} e^{-\mu_j(\mathbf{p})}}{N_j!}. \quad (\text{B1})$$

Using that the number of galaxies in bin j is by construction equal to either 0 or 1, we can write the log-likelihood as

$$\begin{aligned} \ln \mathcal{L}(\mathbf{p}) &= \sum_j \{ N_j \ln[\mu_j(\mathbf{p})] - \mu_j(\mathbf{p}) - \ln[N_j!] \} \\ &= \sum_j \ln[\mu_j(\mathbf{p})] - \int \int \mu(\mathbf{p}) dm dz. \end{aligned} \quad (\text{B2})$$

The second term is a sum over all bins, regardless of whether there is a galaxy in that bin, and so can be replaced by an integral over all the probed redshifts and apparent magnitudes. The first term, on the other hand, is only non-zero for bins that contain a galaxy, and so can be viewed as a sum over all galaxies in the sample, rather than a sum over bins.

Similar to equation (13), we can write

$$\begin{aligned} \int \int \mu(\mathbf{p}) dm dz &= B \int_0^{z_{\max}} \frac{d_c(z)^2 \phi_*(\mathbf{p}, z)}{\sqrt{\Omega_{m,0}(1+z)^3 + \Omega_{\Lambda,0}}} \\ &\times \left[\Gamma \left(\alpha(\mathbf{p}, z) + 1, 10^{\frac{2}{5}(M_*(\mathbf{p}, z) - M_{\text{lim}})} \right) \right]^{-1} \\ &\times \left[\Gamma \left(\alpha(\mathbf{p}, z) + 1, 10^{\frac{2}{5}(M_*(\mathbf{p}, z) - M_{\max}(z))} \right) \right. \\ &\left. - \Gamma \left(\alpha(\mathbf{p}, z) + 1, 10^{\frac{2}{5}(M_*(\mathbf{p}, z) - M_{\min}(z))} \right) \right] dz, \end{aligned} \quad (\text{B3})$$

where $M_{\max}(z) = \min[M(m_{\max}, z), M_{\text{lim}}]$ and $M_{\min}(z) = M(m_{\min}, z)$. In our application, $M_{\text{lim}} = -16$, $z_{\max} = 0.8$, $m_{\min} = 13$, and $m_{\max} = 21$.

In applying this maximum-likelihood method (Section 3.2.2), we assume full information on each galaxy, meaning we can use its absolute magnitude and redshift directly. For galaxy j , we can therefore write (see equation 12):

$$\begin{aligned} \ln[\mu_j(\mathbf{p})] &= \ln \left[\frac{2}{5} \ln(10) B \right] + 2 \ln [d_c(z_j)] + \ln [\phi_*(\mathbf{p}, z_j)] \\ &- \frac{1}{2} \ln [\Omega_{m,0}(1+z_j)^3 + \Omega_{\Lambda,0}] + \frac{2}{5} \ln(10) (\alpha(\mathbf{p}, z_j) + 1) \\ &\times (M_*(\mathbf{p}, z_j) - M_j) - 10^{\frac{2}{5}(M_*(\mathbf{p}, z_j) - M_j)} \\ &- \ln \left[\Gamma \left(\alpha(\mathbf{p}, z_j) + 1, 10^{\frac{2}{5}(M_*(\mathbf{p}, z_j) - M_{\text{lim}})} \right) \right]. \end{aligned} \quad (\text{B4})$$

APPENDIX C: IGNORED EFFECTS

The most important effects that we ignore in our model are (i) the fact that galaxy spectra are not flat and that therefore the relation between apparent and absolute magnitude is not straightforward, and (ii) biasing due to lensing (de)magnification. Incorporating either effect into our model or even quantify how ignoring them impacts our results is far from trivial and outside the scope of this paper. However, in order to gauge the importance of the former effect – that is, ignoring K -corrections – we show in Fig. C1 a comparison between the true i -band apparent magnitude of each mock galaxy as calculated by Henriques et al. (2015) and the one naively derived from that galaxy's i -band absolute magnitude through $m_i = M_i - 5[1 - \log_{10}(d_L(z))]$, with $d_L(z)$ the luminosity distance. All galaxies in the lightcone that satisfy the redshift and (true) magnitude cuts of our catalogue are included.

In the left-hand panel of Fig. C1, brightness indicates the logarithmic density of galaxies at each point in the space. This shows that the effect of ignoring K -corrections is largest for the faintest galaxies and that the naive relation tends to overestimate the apparent brightness of these galaxies. In the right-hand panel, we show now colour-code by redshift. As the redshift increases, the mean true apparent magnitude and the mean difference between it and the naive apparent magnitude increase as well. This shows that, as one might expect, the effect of ignoring K -corrections is strongest at high redshifts.

In our case (with $\Delta m = 0.1$), taken overall redshifts the difference between the true and naive apparent magnitude is at most one bin for the majority of galaxies, but even this small effect may be enough to significantly impact the results when applying our current model to real data. Other photometric bands may be affected differently. More work is needed to explore and account for this.




Figure C1. A comparison of the apparent i -band magnitude naively derived from a galaxy’s absolute magnitude, and its true value. The left-hand panel shows the log-density of galaxies in the plane, while the right-hand panel shows the distribution with redshift. Fainter galaxies and in particular high-redshift galaxies need larger K -corrections and are therefore more sensitive to these being ignored.

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