Title

Galaxies and their Cosmic Variance in the First Billion Years

Author

Trapp, Adam Charles

Publication Date

2022

UNIVERSITY OF CALIFORNIA
Los Angeles

Galaxies and their Cosmic Variance in the First Billion Years

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Astronomy and Astrophysics

by

Adam Charles Trapp

2022

ABSTRACT OF THE DISSERTATION

Galaxies and their Cosmic Variance in the First Billion Years

by

Adam Charles Trapp
Doctor of Philosophy in Astronomy and Astrophysics
University of California, Los Angeles, 2022
Professor Steven R. Furlanetto, Chair

Cosmic variance is the intrinsic scatter in the number density of galaxies due to fluctuations in the large-scale dark matter density field. We begin by presenting a flexible analytic model of cosmic variance in the high redshift Universe ($z \sim 5$–15). We find that cosmic variance in the luminosity function of galaxies at these times is dominated by the variance in the underlying dark matter halo population, and not by differences in halo accretion nor the specifics of our stellar feedback model. We also find that cosmic variance dominates over Poisson noise except for the brightest sources or at very high redshifts ($z \gtrsim 12$). We provide a linear approximation of cosmic variance via a public Python package `galcv`. We then develop a statistical framework that folds our model of cosmic variance into the measurement of the galaxy luminosity function. Through this framework, we forecast the performance of several major upcoming James Webb Space Telescope (JWST) galaxy surveys. We find that they can constrain field matter densities down to the theoretical limit imposed by Poisson noise and unambiguously identify over-dense (and under-dense) regions on transverse scales of tens of comoving Mpc. We then apply this framework to a real Hubble Space Telescope (HST) data set at $z = 6 - 8$, providing a new measurement of the luminosity function of galaxies, and for the first time, measure the underlying densities of the survey fields, including the most over/under-dense HST fields. We show that the distribution of densities is consistent with current predictions for cosmic variance. Finally, we develop the first quantitative, statistically robust framework to infer the underlying density and ionization environment of regions with elevated densities of Lyman-α emitters (LAEs). We apply this framework to an actual observation of 14 LAEs in a $\sim$50,000 cMpc^3 region at $z = 6.93$, obtaining a measurement of that region's density and ionization state, and a constraint on the average ionization fraction of the Universe.

The dissertation of Adam Charles Trapp is approved.

Tommaso L. Treu

Alice Eve Shapley

Smadar Naoz

Steven R. Furlanetto, Committee Chair

University of California, Los Angeles

2022

To my Family,
born into and found.
Your encouragement lets me fly,
and your love brings me back to ground.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

My parents, Rose and Todd Trapp, have always supported what I want for myself, from when I switched instruments from the easy-to-transport trumpet to the impractically-large tuba, to when I decided to spend six years of my life going to college after I had already graduated college. Thank you for your easy acceptance of who I am.

My advisor, Steven Furlanetto, spent an enormous amount of time teaching me cosmology, responding to my erratic emails, and editing my terrible manuscript drafts. I will be eternally grateful for your mentorship and guidance. Thank you.

My thesis committee, especially Smadar Naoz, helped me navigate the waters of switching projects in the middle of grad school, twice. Thank you for your patience, understanding, and wisdom.

I would like to acknowledge the entire Astronomy graduate student body at UCLA for their support and friendship during my time here. I chose to come to UCLA because of the camaraderie I felt during my visit, and that sense of belonging has persisted through all the difficulties of grad school. I will especially miss Groffee and Grunch with you all...

A huge thanks to Jon K. Zink and Richard H. Mebane for helping me puzzle through my many statistical and cosmological questions, and to John A. Trapp for help with publishing the `galcv` package.

Thank you Frederick B. Davies, Jordan Mirocha, Guochao Sun, Brant E. Robertson, R. Bouwens, and the reviewers of our published papers for helpful comments and suggestions, and C. Williams, S.L. Finkelstein, R.A.A. Bowler for their correspondences and expertise on galaxy surveys.

This work was supported by the National Science Foundation through award AST-1812458. In addition, this work was directly supported by the NASA Solar System Exploration Research Virtual Institute cooperative agreement number 80ARC017M0006. We also acknowledge a NASA contract supporting the “WFIRST Extragalactic Potential Observations (EXPO) Science Investigation Team" (15-WFIRST15-0004), administered by GSFC.

Software used: This work makes use of iPython (Perez & Granger, 2007) and the following Python packages: NumPy (van der Walt et al., 2011), SciPy (Virtanen et al., 2020), Matplotlib (Hunter, 2007), and pandas (McKinney, 2010).

Finally, thank you to my therapist and others who specifically helped support my mental health during my time in graduate school. Mental health is as important as physical health. If you need help, please seek it out; there are people and resources out there for you. In fact, here are some of them:

https://www.cdc.gov/mentalhealth/tools-resources/individuals

VITA

2018-2022	UCLA, Los Angeles, California Graduate student researcher
2018	UCLA, Los Angeles, California M.S. in Astronomy
2016	University of Virginia B.A in Physics and Astronomy

PUBLICATIONS AND PATENTS

A.C. Trapp, Steven R. Furlanetto, Jinghong Yang 2022, *MNRAS*, 2022MNRAS.510.4844T, A Framework for Simultaneously Measuring Field Densities and the High-z Luminosity Function

Holly M. Christenson, George D. Becker, Steven R. Furlanetto, Frederick B. Davies, Matthew A. Malkan, Yongda Zhu, Elisa Boera, **Adam Trapp** 2021, *ApJ*, 2021ApJ...923...87C, Constraints on the End of Reionization from the Density Fields Surrounding Two Highly Opaque Quasar Sightlines

V. Dike, M. R. Morris, R. M. Rich, M. O. Lewis, L. H. Quiroga-Nuñez, M. C. Stroh, **A.C. Trapp**, M. J. Claussen 2021, *AJ*, 2021AJ....161..111D, Ground Vibrational State SiO Emission in the VLA BAaDE Survey

A.C. Trapp & Steven R. Furlanetto 2020, *MNRAS*, 2020MNRAS.499.2401T, A Flexible Analytic Model of Cosmic Variance in the First Billion Years

Pallavi Patil, Kristina Nyland, Mark Whittle, Carol Lonsdale, Mark Lacy, Colin Lons-dale, Dipanjan Mukherjee, **A.C. Trapp**, Amy E. Kimball, Lauranne Lanz, BelindaJ. Wilkes, Andrew Blain, Jeremy J. Harwood, Andreas Efstathiou, Catherine Vlahakis 2020, *ApJ*, 2020ApJ...896...18P, High-resolution VLA Imaging of Obscured Quasars: YoungRadio Jets Caught in a Dense ISM

A.C. Trapp, James Larkin, Ken Magnone, Timothee Greffe, Tim Hardy, Jennifer Dunn, Eric Chisholm, Chris Johnson, Ryuji Suzuki 2018, *SPIE*, 2018SPIE10702E..A1T, The infrared imaging spectrograph (IRIS) for TMT: electronics-cable architecture

A.C. Trapp, R. M. Rich, M. R. Morris, L. O. Sjouwerman, Y. M. Pihlstrom, M. Claussen,M. C. Stroh 2018, *ApJ*, 2018ApJ...861...75T, SiO Masers in the Galactic Bulge and Disk:Kinematics from the BAaDE Survey

CHAPTER 1

Dissertation Introduction and Overview

There was an age, in the Universe we know, that no stars shone. No moons nor galaxies abode in their destined place. There was only soup – very hot, glowing soup – filling all space. Its ingredients: everything in the Universe. It was like a tomato bisque, extremely smooth and homogeneous with only tiny bits of texture. These bits of texture proved fortunate, as about them time and gravity congealed the soup into stars and moons and galaxies, and us.

We know of this age because we have seen it. A stubborn bit of static in a strange-looking radio receiver – at one point thought to be caused by bird droppings – turned out to be a signal from glowing soup at the beginning of time (Penzias & Wilson, 1965; Dicke et al., 1965). This soup and its glow has since been extensively measured by many telescopes and is now one of the best studied and well understood phenomena in all of Astrophysics (Planck Collaboration et al., 2016). Thus, we know – in incredible detail – the nature of the Universe around 13,800,000,000 years ago, a mere 380,000 year after the Big Bang. Paradoxically, the time just after this period – the congealing of the soup and the cosmic dawn of starlight – we know very little about. No telescope in astronomy was able to pierce this particular time period, until this year. The James Webb Space Telescope (JWST) can observe some of the first galaxies in the Universe, helping to illuminate the process by which all things we see in this Universe grew from the soup.

The era after the cosmic dawn is also of great interest to Astronomers: the epoch of reionization (EoR). During this time, the congealing galaxies produced many high-energy photons, so many in fact, that all of the soup in-between galaxies was ionized into a plasma. This process took a few hundred million years, finishing around 1,000,000,000 years after the Big Bang, or 12,800,000,000 years ago. Reionization has been the subject of intense study, but JWST will be able to observe the individual sources that are contributing ionizing photons in much greater detail than ever before.

1.1 The luminosity function and cosmic variance

One of the first things JWST will do is take a census of these galaxies during the EoR and cosmic dawn. With this census Astronomers will build a “luminosity function”, which counts the galaxies categorized by their luminosity. The luminosity function is an extremely important feature of the galaxy population as a whole; its shape and evolution has been studied extensively at the redshifts available to pre-JWST instruments, most importantly, the Hubble Space Telescope (see e.g., Schenker et al., 2013; McLure et al., 2013; Bouwens et al., 2015; Finkelstein et al., 2015; Bowler et al., 2015; Livermore et al., 2017; Atek et al., 2018; Oesch et al., 2018; Behroozi et al., 2019; Bouwens et al., 2021; Finkelstein et al., 2022). The luminosity function typically takes the form of a “Schechter function,” characterized by a power-law slope at faint magnitudes and an exponential cutoff at some characteristic magnitude. This function has the form

$$\Phi_{\mathrm{avg}}(M,z)dM = (0.4\ \ln10)\phi^{*}[10^{0.4(M^{*}-M)}]^{\alpha+1}\mathrm{exp}[-10^{0.4(M^{*}-M)}]dM, \tag{1.1}$$

where $\Phi_{\mathrm{avg}}(M,z)dM$ is the average number density of galaxies with absolute magnitudes between $(M, M+dM)$. The parameters of this function are the normalization ϕ^{*}, the characteristic magnitude M^{*}, and the faint-end power-law slope, α. This “average” luminosity function describes the Universal population of galaxies, and its parameters evolve with redshift.

However, this average luminosity function may not well describe the galaxies in a given volume V. As part of the congealing of the soup, some regions contain much matter and many galaxies, while others are relatively devoid of matter. This effect is called “cosmic variance,” and is the main subject of this dissertation. Thus the “local” luminosity function in the volume V will look different than average. For large volumes, the local luminosity function can be described linearly (Mo & White, 1996; Trapp & Furlanetto, 2020):

$$\Phi_{\mathrm{loc}}(M, z, \delta, V) = \Phi_{\mathrm{avg}}(M, z)\left(1 + \frac{\delta}{\sigma_\delta}\varepsilon_{\mathrm{cv}}(M, z, V)\right), \tag{1.2}$$

where δ is the relative matter density of V, σ_δ is the root mean squared (r.m.s.) fluctuation of δ across many volumes of the same size, and $\varepsilon_{\mathrm{cv}}$ is the "bias function" that indicates the strength by which the luminosity function changes given some matter density, which depends on the magnitude considered and the size and shape of V. The value of σ_δ is relatively well understood (see Chapter 2), so $\varepsilon_{\mathrm{cv}}$ carries most of the information about cosmic variance. Unfortunately, the effects of Poisson noise further complicate this picture. If the local luminosity function predicts there to be N_{exp} galaxies of magnitude M in some volume V with density δ, the actual number of sources found in a given volume N_{obs} will be drawn from the Poisson distribution with probability

$$P(N_{\mathrm{obs}}) = \frac{N_{\mathrm{exp}}^{N_{\mathrm{obs}}} e^{-N_{\mathrm{exp}}}}{N_{\mathrm{obs}}!}. \tag{1.3}$$

The relative width of this distribution goes as approximately $1/\sqrt{N_{\mathrm{exp}}}$, so if a large region or number of regions are sampled, the effects of Poisson noise is reduced. Thus, the local luminosity function is actually an average across all volumes V with density δ.

Like a census of people, it's important that JWST's census of galaxies samples many different locations, lest it be biased towards one type of region or subject to large Poisson uncertainties. Prior knowledge of the strength of cosmic variance is valuable for planning a census and interpreting its results. Thus, cosmic variance has been modeled in a variety of ways in the past. For example, analytic models typically start with the linear halo bias function and then connect haloes to galaxies with a halo mass to luminosity relation, or by matching abundances (see e.g., Newman & Davis, 2002; Somerville et al., 2004; Stark et al., 2007; Moster et al., 2011). These models conclude that cosmic variance is a significant source of uncertainty when studying galaxies at high redshifts. However, such models do not allow cosmic variance to affect the halo mass to luminosity connection itself; they assume galaxies are the same in all environments. Cosmic variance can also be estimated using mock observations of galaxy simulations. The early implementations of this method (see e.g., Kitzbichler

& White, 2007; Trenti & Stiavelli, 2008) were very powerful, but also assumed galaxies are the same in all environments. Recently, substantial improvements in computing power have allowed for much higher-volume N-body simulations that also treat star formation in a more complex way (e.g., Bhowmick et al., 2020; Ucci et al., 2021). These studies take into account the difference in environment on individual galaxy growth, and are a major step forward in predicting cosmic variance in the first galaxies. However, these studies (1) are still limited by their volume, as a complete picture of cosmic variance requires extremely large volumes to calculate cosmic variance on all relevant scales and magnitudes; (2) cannot explore how cosmic variance depends on their specific implementation of mass accretion, star formation, feedback, and other parameters, without re-running simulations many times, which would be prohibitively expensive; and (3) can be limited in their redshift or magnitude ranges. Simulations lack large volume and flexibility, while existing analytic models rely on linear theory and lack a fully self-consistent connection between cosmic variance and galaxy growth and star formation.

In Chapter 2 we develop a simple, flexible analytic model of high-z galaxies to study the effects of cosmic variance on the galaxy population. We begin with a close examination of large-scale variations in the dark matter halo population, where galaxies reside. This provides the fundamental basis for the cosmic variance of the galaxy population, and with it we capture some non-linear aspects of these fluctuations. Next, we describe a “minimalist" model of galaxy evolution (Furlanetto et al., 2017) that fits observed luminosity functions reasonably well but is sufficiently flexible to examine how a large range of assumptions about the physics of these sources affects cosmic variance. We use this model to determine which of the many uncertain parameters of galaxy formation have the most impact on cosmic variance, and we account for changes in galaxy growth and star formation in different environments. Finally, we combine our treatments of dark matter haloes and galaxy physics to create a linear approximation of the cosmic variance of galaxies as a function of redshift and absolute magnitude: $\varepsilon_{\rm cv}$ (available via an easy to install public Python package called `galcv`). Unlike other such functions derived from simulations, our results apply across any mass or redshift, and we quantify how uncertainties in galaxy evolution parameters affect the results. This model of $\varepsilon_{\rm cv}$ underpins the rest of the projects in this dissertation.

Our model of $\varepsilon_{\rm cv}$ (as well as others') predicts that the effects of cosmic variance will get stronger as redshift increases through the EoR and into the cosmic dawn. This will increase the difficulty of measuring the luminosity function in these times. So, in Chapter 3, we develop a framework to simultaneously measure the high-z average luminosity function $\Phi_{\rm avg}$ and field densities δ from a galaxy census. Unlike the standard approach to estimating luminosity functions, which acknowledges the existence of cosmic variance but does not attempt to model it, our new framework uses Bayesian statistics to fold in a comprehensive model of cosmic variance (Trapp & Furlanetto, 2020) and its effect on the galaxy population.

Further, we argue cosmic variance is not just a nuisance. It is itself a key driver of both galaxy formation and reionization during the cosmic dawn because it reflects real large-scale structure in the Universe. If these large-scale densities can be measured, they can complement the luminosity function as another bedrock observable. The insights to be gained from such measurements include:

(i) The process of reionization likely began in the densest parts of the Universe and ended in the largest voids. Identifying such over/under-densities is an area of great interest (see e.g. Zitrin et al., 2015; Jung et al., 2020; Tilvi et al., 2020; Hu et al., 2021; Endsley et al., 2021; Becker et al., 2018; Davies et al., 2018; Christenson et al., 2021).

(ii) Large-scale feedback mechanisms, driven by large-scale density, are likely to strongly affect the galaxy population before and during reionization (Thoul & Weinberg, 1996; Iliev et al., 2007; Noh & McQuinn, 2014).

(iii) Measuring large-scale densities at early times facilitates the understanding of the assembly history of rare objects like galaxy clusters (e.g., Chiang et al. 2017), which form from the densest environments.

(iv) Finally, comparing large-scale density measurements from surveys with theoretical predictions of cosmic variance can help test models of the galaxy–halo connection (Trapp & Furlanetto, 2020).

We then use our framework to forecast the performance of a selection of upcoming galaxy surveys, exploring how well different combinations of the surveys can measure the luminosity function and individual field densities.

In Chapter 4, we apply the framework from Chapter 3 to existing HST data, making a joint measurement of galaxy luminosity functions and large-scale structure during the EoR. We obtain a new measurement of the galaxy luminosity function for $z = 6$–8 and, for the first time, measure the underlying large-scale density of many HST survey fields.

1.2 Lyman-α emitters as probes of reionization

Finally, we take a turn towards the process of reionization in Chapter 5. Lyman-α emitters (LAEs) are an exciting probe of reionization (as well as galaxy formation). The young galaxies present in the reionization era are likely to have large intrinsic Lyman-α luminosities (Partridge & Peebles, 1967; Ouchi et al., 2020), but those photons are subject to substantial absorption by the intergalactic medium (IGM) before reionization is complete, thanks to the enormous optical depth of remaining neutral islands (Gunn & Peterson, 1965). We therefore expect to see a decline in the relative abundance of LAEs as we penetrate further into the reionization era, making these galaxies an effective probe of the ionization state (Madau & Rees, 2000; Haiman, 2002; Santos, 2004). A particularly interesting aspect of this modulation is that it exaggerates the existing clustering: galaxies in over-dense regions (which host large ionized bubbles) will remain visible, while even galaxies inside modest over-densities will disappear early enough in reionization. This suggests that focusing on identifying rare ionized regions can be a powerful probe of reionization (e.g., Mesinger & Furlanetto 2008).

Of particular interest for us, several teams have discovered candidate ionized bubbles hosting apparently significant over-densities of LAEs (e.g., Tilvi et al. 2020; Jung et al. 2020; Hu et al. 2021; Endsley et al. 2021). However, to date such inferences are mostly qualitative, with only simple efforts to transform observed features in the galaxy distribution to a quantitative constraint on the underlying ionization and density fields. Finding and analysing these regions, and then connecting the visible sources to the total population of galaxies that are ionizing the Universe is challenging on many fronts. The bulk of ionizing photons are coming from galaxies that we cannot see with HST *nor* JWST (see e.g., Behroozi & Silk, 2015; Furlanetto et al., 2017). In a partially-ionized Universe, even strong Lyman-α lines can be obscured by intervening neutral Hydrogen, leaving only a small number of the brightest LAEs in the most ionized regions visible. The observed number density of these LAEs is then strongly affected by Poisson noise. Poisson noise is uncorrelated between magnitude bins, meaning an over-density of a few bright LAEs does not necessarily mean an over-density in the much-larger underlying population of faint or obscured sources. In other words, determining the large-scale density of a region from a small number of LAEs requires a simultaneous treatment of *(i)* cosmic variance and the corresponding galaxy bias of those LAEs, *(ii)* Poisson noise, and *(iii)* the effects of a partially ionized Universe on the visibility of LAEs.

In Chapter 5, we use our model of cosmic variance to develop the first quantitative, statistically robust framework to infer the underlying density and ionization environment of regions with elevated densities of LAEs. This framework can be applied both to LAEs (which probe both the ionization state and the density) and to other surveys (which probe only the density). We apply this method to an actual observation of 14 LAEs in a $\sim$50,000 cMpc^3 region at $z = 6.93$ (Hu et al., 2021), obtaining a measurement of its density and ionization, and a constraint on the average ionization fraction of the Universe at this time.

We summarize the results of this dissertation and look to future applications in Chapter 6.

CHAPTER 2

A Flexible Analytic Model of Cosmic Variance in the First Billion Years

This chapter can be seen in its published form (Trapp & Furlanetto, 2020) here:
https://ui.adsabs.harvard.edu/abs/2020MNRAS.499.2401T/abstract

2.1 Abstract

Cosmic variance is the intrinsic scatter in the number density of galaxies due to fluctuations in the large-scale dark matter density field. In Chapter 2, we present a simple analytic model of cosmic variance in the high redshift Universe ($z \sim 5$–15). We assume that galaxies grow according to the evolution of the halo mass function, which we allow to vary with large-scale environment. Our model produces a reasonable match to the observed ultraviolet luminosity functions in this era by regulating star formation through stellar feedback and assuming that the UV luminosity function is dominated by recent star formation. We find that cosmic variance in the UVLF is dominated by the variance in the underlying dark matter halo population, and not by differences in halo accretion or the specifics of our stellar feedback model. We also find that cosmic variance dominates over Poisson noise for future high-z surveys except for the brightest sources or at very high redshifts ($z \gtrsim 12$). We provide a linear approximation of cosmic variance for a variety of redshifts, magnitudes, and survey areas through the public Python package `galcv`. Finally, we introduce a new method for incorporating priors on cosmic variance into estimates of the galaxy luminosity function and demonstrate that it significantly improves constraints on that important observable.

2.2 Introduction

Extragalactic astronomy is closing in on arguably the most important era of galaxy evolution: the formation of the first galaxies. These galaxies will allow us to probe the processes that drove the first emergence of complexity in our Universe.

The most fundamental observations for studying this era have been (and will continue to be) deep galaxy surveys. These surveys measure many important features of the galaxy population, most fundamentally the ultra-violet luminosity function (UVLF) of galaxies. The UVLF is a measure of the number of galaxies at each luminosity, and its shape and evolution through cosmic time has important implications for the physics behind galaxy formation and growth, and much more (see e.g., Bouwens et al., 2015; Finkelstein et al., 2015; Livermore et al., 2017; Atek et al., 2018; Oesch et al., 2018; Behroozi et al., 2019). Unfortunately, these deep galaxy surveys will have very small volumes, which will be a key limitation in measuring the UVLF due to the effects of "cosmic variance:" not all regions of the Universe contain the average number of galaxies, and those galaxies did not all grow up in an average environment. We must understand how cosmic variance affects the UVLF in order to inform and correctly interpret future deep galaxy surveys.

Cosmic variance in the UVLF (and other measures) has been modeled in a variety of ways in the past. For example, analytic models typically start with the linear halo bias function and then connect haloes to galaxies with a halo mass to luminosity relation, or by matching abundances (see e.g., Newman & Davis, 2002; Somerville et al., 2004; Stark et al., 2007; Moster et al., 2011). These models conclude that cosmic variance is a significant source of uncertainty when studying galaxies at high redshifts. However, such models do not allow cosmic variance to affect the halo mass to luminosity connection itself; they assume galaxies are the same in all environments. Also, the linear halo bias function does not accurately predict cosmic variance in extreme environments.

Cosmic variance can also be estimated using mock observations of galaxy simulations. The early implementations of this method (see e.g., Kitzbichler & White, 2007; Trenti & Stiavelli, 2008) were very powerful, but also assumed galaxies are the same in all environments.

Recently, substantial improvements in computing power have allowed for much higher-volume N-body simulations that also treat star formation in a more complex way (e.g., Bhowmick et al., 2020; Ucci et al., 2021). These studies take into account the difference in environment on individual galaxy growth, and are a major step forward in predicting cosmic variance in the first galaxies. However, these studies (1) are still limited by their volume, as a complete picture of cosmic variance requires extremely large volumes to calculate cosmic variance on all relevant scales and magnitudes; (2) cannot explore how cosmic variance depends on their specific implementation of mass accretion, star formation, feedback, and other parameters, without re-running simulations many times, which would be prohibitively expensive; and (3) can be limited in their redshift or magnitude ranges.

Simulations lack large volume and flexibility, while existing analytic models rely on linear theory and lack a fully self-consistent connection between cosmic variance and galaxy growth and star formation.

Quantitative interpretations of high-z data require corrections for cosmic variance, especially because most planned surveys subtend relatively small volumes. Such corrections are particularly important when multiple independent surveys are combined, as each such survey contains its own (unknown) intrinsic density. The standard method to account for cosmic variance when fitting a UVLF, originally developed by Sandage et al. (1979) and used by e.g. Efstathiou et al. (1988) and Bouwens et al. (2015), fits a universal *shape* of the UVLF to all fields, ignoring the normalization parameter of the fit in each individual field. After the shape has been optimized, the overall normalization is determined by demanding that it reproduces the correct total number of galaxies across all surveys. This method cannot account for a change in shape of the UVLF between fields, and it does not include a prior for the amount of variance allowed in the normalization parameter of the UVLF.

In this paper, we use a simple, flexible analytic model of high-z galaxies to study the effects of cosmic variance on galaxy surveys. We begin in section 2.3 with a close examination of large-scale variations in the dark matter halo population in the context of excursion set models of halo formation. This provides the fundamental basis for the cosmic variance of the galaxy population, and with it we capture some non-linear aspects of these fluctuations.

Next, in section 2.4, we describe a “minimalist” model of galaxy evolution (Furlanetto et al., 2017) that fits observed luminosity functions reasonably well but is sufficiently flexible to examine how a large range of assumptions about the physics of these sources affects cosmic variance. We use this model to determine which of the many uncertain parameters of galaxy formation have the most impact on cosmic variance, and we account for changes in galaxy growth and star formation in different environments. We combine our treatments of dark matter haloes and galaxy physics in section 2.5, where we also provide a linear approximation to the cosmic variance of galaxies as a function of redshift and absolute magnitude. Unlike other such functions derived from simulations, our results apply across any mass or redshift, and we quantify how uncertainties in galaxy evolution parameters affect the results.

In section 2.6, we then describe the importance of our cosmic variance results for future surveys of high-z galaxies with the James Webb Space Telescope (JWST) and the Nancy Grace Roman Space Telescope (hereafter, the Roman Space Telescope; Spergel et al., 2015; Akeson et al., 2019; Dore et al., 2019). We show how cosmic variance limits inferences about the average UVLF of the Universe. Additionally, we introduce a method that fully incorporates cosmic variance into UVLF estimates, essentially treating our estimates for cosmic variance as a prior on the measurements. Most commonly, UVLF estimates allow for an arbitrary amount of cosmic variance between fields, by ignoring the normalization of the UVLF in each field (e.g., Bouwens et al. 2015; Finkelstein et al. 2015, though see Livermore et al. 2017 for a contrasting case). We show that our method provides tighter constraints in mock surveys. Finally, in section 2.7, we summarize our results.

In htis chapeter, we take the following cosmological parameters: $\Omega_m = 0.308$, $\Omega_\Lambda = 0.692$, $\Omega_b = 0.0484$, $h = 0.678$, $\sigma_8 = 0.815$, and $n_s = 0.968$, consistent with recent Planck Collaboration XIII results (Planck Collaboration et al., 2016). We give all distances in comoving units.

2.3 Dark Matter Haloes

We follow the methods described in Furlanetto et al. (2017) to model dark matter haloes. In this section we give a brief summary of those methods and also describe some additions.

2.3.1 Conditional halo mass function

We define the dark matter halo mass function as $n_{\rm h}(m, z)dm$: the comoving number density of dark matter haloes between masses $(m, m + dm)$ at redshift z. By convention,

$$n_{\rm h}(m, z) = f(\sigma)\frac{\bar{\rho}}{m}{\rm dln}\frac{(1/\sigma)}{{\rm d}m}, \tag{2.1}$$

where $\bar{\rho}$ is the comoving average matter density, $\sigma(m, z)$ is the linear rms fluctuation of the matter density field at redshift z smoothed over a spherical region of mass m (see section 2.3.3 for the calculation of $\sigma(m, z)$), and $f(\sigma)$ is a dimensionless function that modifies the shape of the mass function. Following Furlanetto et al. (2017), we use $f_{\rm Trac}(\sigma)$ from a fit to the average mass function of a high-z cosmological simulation (Trac et al., 2015):

$$f_{\rm Trac}(\sigma) = 0.150\left[1 + \left(\frac{\sigma}{2.54}\right)^a\right]e^{b/\sigma^2}, \tag{2.2}$$

with $a = -1.36$ and $b = -1.14$.

The key aspect of this model is allowing the halo mass function to depend on its environment. This environmental dependence is introduced via the conditional mass function (CMF). The CMF $n_{\rm cond}(m, z, \delta_b, R)$ describes the number density of haloes in a spherical region of mass M, with a corresponding Lagrangian radius[1] $R^3 = 3M/(4\pi\bar{\rho})$, and relative density $\delta_b = (\rho - \bar{\rho})/\bar{\rho}$, where ρ is the linearly extrapolated matter density in that region at redshift z. The CMF is what adds cosmic variance into the model.

We determine $n_{\rm cond}(m, z, \delta_b, R)$ using a coordinate transfer method described in Tramonte

---

[1] Note that the radius R is really a mass scale, as it does not correspond to the real radius of a region except for regions that happen to be at cosmological average density.

et al. (2017) that we will call "ν-scaling" applied to equation (2.2)[2]

$$\sigma^2(m,z) \rightarrow [\sigma^2(m,z) - \sigma^2(M,z)] \left[\frac{\delta_{crit}}{\delta_{crit} - \delta_b}\right]^2, \tag{2.3}$$

where M is the mass corresponding to the region R, and $\delta_{crit} \approx 1.69$ is the linear halo collapse threshold (see Loeb & Furlanetto, 2013, eq.(3.13)). The resulting CMF is Lagrangian in that it assumes all regions of fixed mass have the same volume. To convert into a real-space (Eulerian) CMF, we calculate the real-space radius $R_{\rm e}$ of a region of mass M and density δ_b assuming spherical collapse: $R_{\rm e} = R/(1+\delta_r)^{1/3}$, where δ_r is the real-space (non-linear) relative density[3](Mo & White, 1996, see Appendix 2.8 for more details). Applying this adjustment to the radius of each region results in an Eulerian CMF, $n_{\rm cond}(m,z,\delta_b,R_{\rm e}) = n_{\rm cond}(m,z,\delta_b,R) \times (1+\delta_r)$.

Tramonte et al. (2017) justify ν-scaling by noting δ_{crit} enters into $f(\sigma)$ only through the variable $\nu = \delta_{crit}/\sigma$ in previous parameterizations. They then apply the "standard" coordinate transfer

$$\begin{gathered} \delta_{crit} \rightarrow \delta_{crit} - \delta_b \\ \sigma^2(m,z) \rightarrow \sigma^2(m,z) - \sigma^2(M,z) \end{gathered} \tag{2.4}$$

to the variable ν, giving equation (2.3). Tramonte et al. (2017) validate this method using an N-body simulation by Tinker et al. (2008), finding that this scaling technique accurately describes the CMF except for the most underdense regions ($\delta_b \lesssim -1.5$), where it overestimates halo abundance.

We note that while we consider halo masses down to $m \sim 10^8 \mathrm{M}_\odot$, Tramonte et al. (2017) test their prescription only down to a mass of $m \sim 3 \times 10^{10} \mathrm{M}_\odot$, leaving it untested for the lowest masses we consider. We also note that we use a different mass function (Trac et al., 2015), for which this method has not been explicitly tested. Tramonte et al. (2017) also test "local scaling", a more rigorous method for constructing a CMF developed in Patiri et al. (2006) and expanded in Rubiño-Martín et al. (2008), and find that method produces

---

[2] We construct the CMF this way because simulations of dark matter haloes (e.g. McBride et al., 2009; Goerdt et al., 2015; Trac et al., 2015) do not provide a full CMF.

[3] In practice, δ_r and δ_b are very similar, especially at the redshifts considered in this paper.

a slightly better CMF. We do not test our model using local scaling as it cannot be easily applied to the Trac et al. (2015) mass function.

Mo & White (1996) developed a method to linearly approximate their CMF by use of a bias function b_{PS}. We follow the same steps to calculate a linear bias factor b_{Trac} for the Trac et al. (2015) mass function. We first substitute equation (2.3) into equation (2.2), and then Taylor expand to linear order about $\delta_b = 0$, and set $\sigma(M, z) = 0$. We then make a linear volume change correction of $(1 + \delta_b)$, giving us

$$b_{\mathrm{Trac}} = 1 + \frac{a}{\delta_{crit}} \frac{(\sigma/2.54)^a}{1 + (\sigma/2.54)^a} - \frac{2b}{\sigma^2 \delta_{crit}}. \tag{2.5}$$

The CMF can then be approximated as

$$n_{\mathrm{cond,\ lin}}(m, z, \delta_b) = n_{\mathrm{h}}(m, z)(1 + b_{\mathrm{Trac}} \delta_b). \tag{2.6}$$

We will use this linear approximation of the CMF to compare to our results when using the full CMF.

For another comparison, we calculate the CMF by scaling the Trac et al. (2015) mass function by the ratio of the conditional to non-conditional Press-Schechter mass functions.

$$\frac{n_{\mathrm{cond}}}{n_{\mathrm{h}}} = \frac{n_{\mathrm{PS,cond}}}{n_{\mathrm{PS}}} \tag{2.7}$$

where n_{PS} is defined in Press & Schechter (1974), and represented here as $f_{\mathrm{PS}}(\sigma)$ (which is plugged into eq. 2.1).

$$f_{\mathrm{PS}}(\sigma) = \sqrt{\frac{2}{\pi}} \nu \ e^{-\nu^2/2} \tag{2.8}$$

The Press-Schechter CMF $n_{\mathrm{PS,cond}}(m, z, \delta_b, R)$ is obtained with the “standard” coordinate transfer in equation (2.4). We then multiply by the same $(1+\delta_r)$ factor to obtain an Eulerian CMF. We note that this scaling, which was introduced in the high-z context by Barkana & Loeb (2004), is commonly used in analytic, semi-analytic, and semi-numeric calculations of galaxy populations at this time (e.g., in Mesinger et al. 2011).

At these high redshifts, we ignore the effects of assembly bias (e.g., Gao & White 2007) on the CMF, as it is a small effect compared to the other uncertainties in our model.

2.3.2 Dark matter density fluctuations

As stated above, $\sigma(M, z)$ is the linear rms fluctuation of the matter density field at redshift z smoothed over a spherical region of mass M (and corresponding R). Thus, the probability distribution of dark matter density for a given scale R and redshift z, $p(\delta_b|R, z)$, is by definition equal to a zero-mean Gaussian with variance $\sigma^2(M, z)$. However, galaxy surveys measure Eulerian volumes, so to make predictions for them we must convert this distribution to that system. A fixed Eulerian volume will correspond to a range of masses, because each has a different density. In Appendix 2.8, we convert the probability distribution of densities at fixed region mass $p(\delta_b|R, z)$ to fixed real-space volume $p(\delta_b|R_\mathrm{e}, z)$. While $p(\delta_b|R, z)$ is a Gaussian, $p(\delta_b|R_\mathrm{e}, z)$ is closer to an inverse Gaussian. Fortunately, these two distributions are very similar to one another at the region sizes and redshifts we consider in this paper.

However, we do find that $p(\delta_b|R_\mathrm{e}, z)$ predicts that underdense regions occupy a larger volume-fraction of the Universe than overdense regions at all scales, by as much as ~12 per cent when considering very small scales. This result indicates that surveys will be slightly more likely to probe underdense regions (see Appendix 2.8 for more details). Using different methods, Muñoz et al. (2010) also found that surveys are more likely to probe an underdense region because of those regions' more rapid cosmic expansion.

With the distribution of densities $p(\delta_b|R_\mathrm{e}, z)$ and the CMF $n_\mathrm{cond}(m, z, \delta_b, R_\mathrm{e})$, we can compute the scatter in halo number density on various scales R_e and redshifts. Figure 2.1 shows some example results. Cosmic variance in the mass function is substantial for the haloes in which high-z galaxies form. For example, at redshift $z = 9$ on a 50 Mpc (radius) scale, massive haloes ($\sim 10^{12}$ $\mathrm{M}_\odot$) have a typical relative standard deviation of ~65 per cent, while haloes at the atomic cooling limit (see Section 2.4.1) have a relative standard deviation of ~10 per cent. At fixed halo mass, these relative standard deviations increase at higher redshifts and decrease when considering larger volumes.

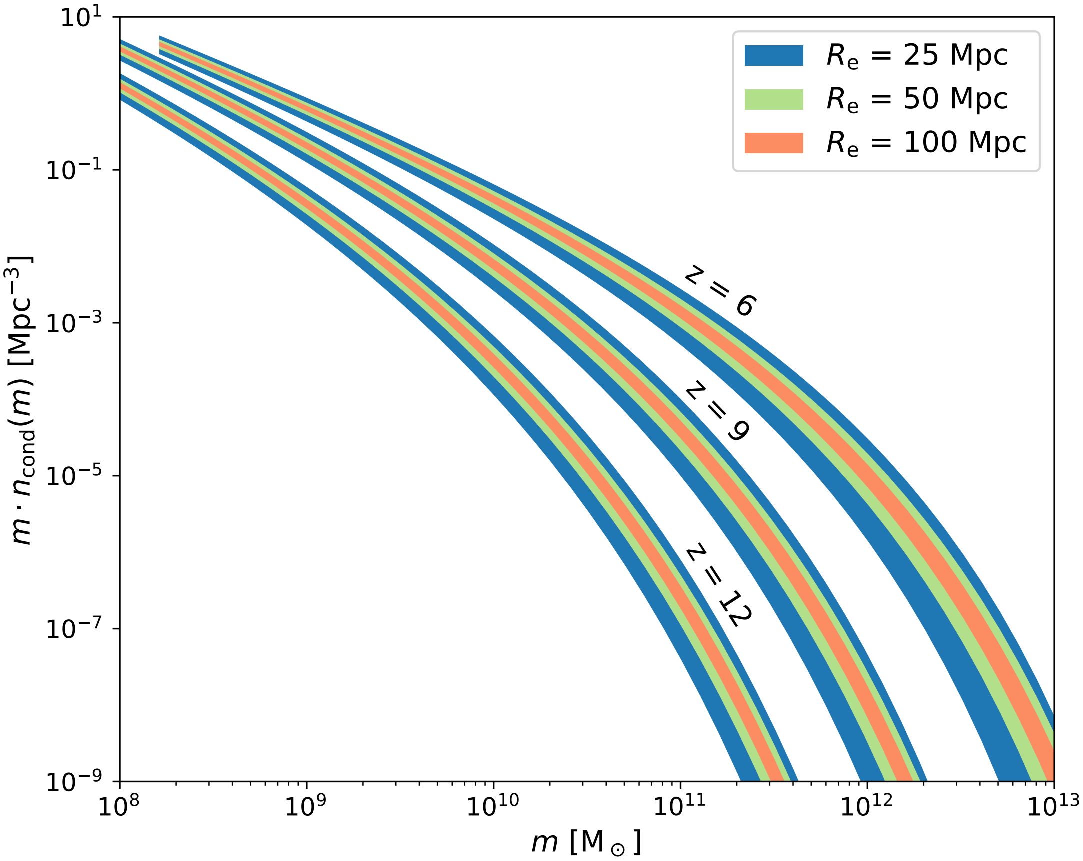

Figure 2.1: The CMF n_{cond} and its 2σ scatter due to cosmic variance on various scales (identified with their radius R_{e}) at three redshifts. The scatter increases at high mass and for smaller scales of the Universe (widest shaded area corresponds to the smallest scale).

2.3.3 The CMF in realistic survey volumes

The CMF presented in section 2.3.1 assumes a spherical region of radius $R_{\rm e}$. However, real surveys subtend elongated regions pointing away from Earth, commonly referred to as pencil beams. Here, we describe a method for building a CMF for a pencil-beam region.

We start with the variance in the dark matter density field σ^2 in a pencil-beam region (following e.g., Newman & Davis, 2002; Stark et al., 2007; Muñoz et al., 2010; Robertson, 2010). For an arbitrarily-shaped volume $\mathbf{V}$,

$$\sigma^2(\mathbf{V}) = \frac{F_g(z)}{(2\pi)^3}\int P(\mathbf{k})|\hat{W}_{\mathbf{V}}(\mathbf{k})|^2 d\mathbf{k}, \tag{2.9}$$

where $F_g(z)$ is the growth function (nearly equal to $1/(1+z)$), $\mathbf{k}$ is wave vector, $P(\mathbf{k})$ is the power spectrum of dark matter (we use the transfer function from Eisenstein & Hu, 1998), and $\hat{W}_{\mathbf{V}}(\mathbf{k})$ is the Fourier transform of a real-space top hat in the shape of the region $\mathbf{V}$, normalized such that its integral in real-space is equal to unity. In the case of a rectangular pencil-beam volume with side lengths $a_x,\ a_y,\ a_z$, $\hat{W}_{\mathbf{V}}(\mathbf{k}) = \hat{W}(k_x)\hat{W}(k_y)\hat{W}(k_z)$ with $\hat{W}(k_i) = \sin(a_i k_i/2)/(a_i k_i/2)$. When constructing a survey volume, we define a_z as the radial distance corresponding to some Δz centered at z. We define a_x and a_y such that the physical area $a_x \cdot a_y$ at z gives the survey area A as seen from Earth.

We then make the simple approximation that a pencil-beam region has the same CMF as a (larger) spherical region of radius $R_{\rm eff}$, such that $\sigma_{\rm sphere}(R_{eff}) = \sigma_{\rm PB}(a_x, a_y, a_z)$. This prescription is analogous to how pencil-beam volumes are treated in other analytic studies of cosmic variance. In such studies, $\sigma_{\rm PB}$ is multiplied by a halo bias function to find cosmic variance (e.g., Stark et al., 2007; Muñoz et al., 2010; Robertson, 2010; Moster et al., 2011). The halo bias functions used in these studies come from CMFs determined assuming spherical regions, so that pencil-beam volumes are treated as spherical volumes with equivalent σ.

2.3.4 Accretion rates

We now consider how dark matter haloes accrete matter. This accretion will be used in the next section to determine the rate of star formation.

Many simulations provide similar predictions of halo mass accretion rates (e.g., McBride et al., 2009; Fakhouri et al., 2010; van den Bosch et al., 2014; Goerdt et al., 2015; Trac et al., 2015). However, these rates have not been tested at the very high redshifts and very low masses relevant to our model. For this paper, we calculate the accretion rates using the method described in Furlanetto et al. (2017), which is analogous to abundance matching (Vale & Ostriker, 2004): haloes maintain a constant number density as they evolve according to the mass function of the region they are in. That is, we require that at any two nearby redshifts z_1 and z_2, a halo has masses $m_1(z_1)$ and $m_2(z_2)$ such that:

$$\begin{aligned} \int_{m_1}^{\infty} dm \; n_{\rm cond}(m, z_1, \delta_b, R) \; &= \\ \int_{m_2}^{\infty} dm \; n_{\rm cond}(m, z_2, \delta_b, R), \end{aligned} \tag{2.10}$$

where $n_{\rm cond}(m, z, \delta_b, R)$ is the Lagrangian CMF from equation (2.7). We define the accretion rate of haloes $\dot{m}_{\rm h}$ in a region such that they satisfy equation (2.10) at all masses over a small redshift interval ($\Delta z \sim 0.1$).

In practice this accretion method means that in a given region, the most massive halo at one time step is also the most massive halo at the next time step, and the same goes for the second and third most massive haloes, etc. In this treatment, accretion is continuous and smooth, increases monotonically with halo mass, and has zero scatter at a fixed mass. This treatment is obviously not entirely correct, but it is in line with our goal of simplicity and has the added benefit maintaining the CMF across cosmic time in a way that conserves mass. Furlanetto et al. (2017) show that accretion rates obtained using this method are similar to the simulation accretion rates mentioned above in the redshift and mass ranges they probe. We do neglect mergers in this model, which will provide an additional source of scatter (see the discussion in Furlanetto et al. 2017).

This method allows haloes in over and underdense regions to accrete at different rates. However, we find that this is actually a small effect. The shaded areas in Figure 2.2 show how the accretion rates depend on large scale environment. For most masses, redshifts, and scales, the variation in accretion is less than 5 per cent. A 5 per cent difference in accretion can be significant over a Hubble time, but we will show it has a small effect on the UVLF, which is most sensitive to instantaneous star formation.

2.4 Feedback-Regulated Star Formation

In this section we transform the mass accretion rates of haloes into ultraviolet (UV) luminosities. Furlanetto et al. (2017) provides a detailed explanation of our star formation model; we briefly summarize it here. We intentionally choose this simple, "minimalist" model so as to make our assumptions about the mass-luminosity relation transparent. Estimates of cosmic variance must necessarily account for the many uncertainties about high-z galaxies, and a simple, flexible model allows us to estimate how important the specifics of galaxy formation are for the variance.

2.4.1 Models of feedback

We assume that haloes only form stars when they exceed a threshold mass $m_{\rm min}$. This mass corresponds to a halo virial temperature $T_{vir} = 10^4 K$, when atomic line cooling becomes efficient enough for gas clouds to collapse and fragment for star formation (Loeb & Furlanetto, 2013). This mass is typically $m_{\rm min} \sim 10^8 M_\odot$. At the redshifts considered in this paper, haloes at the threshold are always far below the detection limit of next generation telescopes.

Gas accreting onto a galaxy can be turned into stars. When stars form, they expel baryons from their host galaxy through radiation pressure, supernovae, or some other process like grain heating (e.g., Faucher-Giguère et al., 2013; Hayward & Hopkins, 2017; Krumholz et al., 2018). Balancing this stellar feedback with accretion provides a simple estimate of the star formation rate $\dot{m}_*$ of a galaxy via

$$\dot{m}_* = \dot{m}_{\rm b} - \dot{m}_{\rm w}, \tag{2.11}$$

where $\dot{m}_{\rm b}$ is the mass accretion rate of the halo times the baryon fraction $\dot{m}_{\rm b} = [\Omega_{\rm b}/\Omega_{\rm m}]\dot{m}_{\rm h}$, and $\dot{m}_{\rm w}$ is the rate of baryon loss through feedback. The fraction of accreting baryons that are converted into stars is defined as $f_* = \dot{m}_*/\dot{m}_{\rm b}$. Finally, we write the mass ejection rate

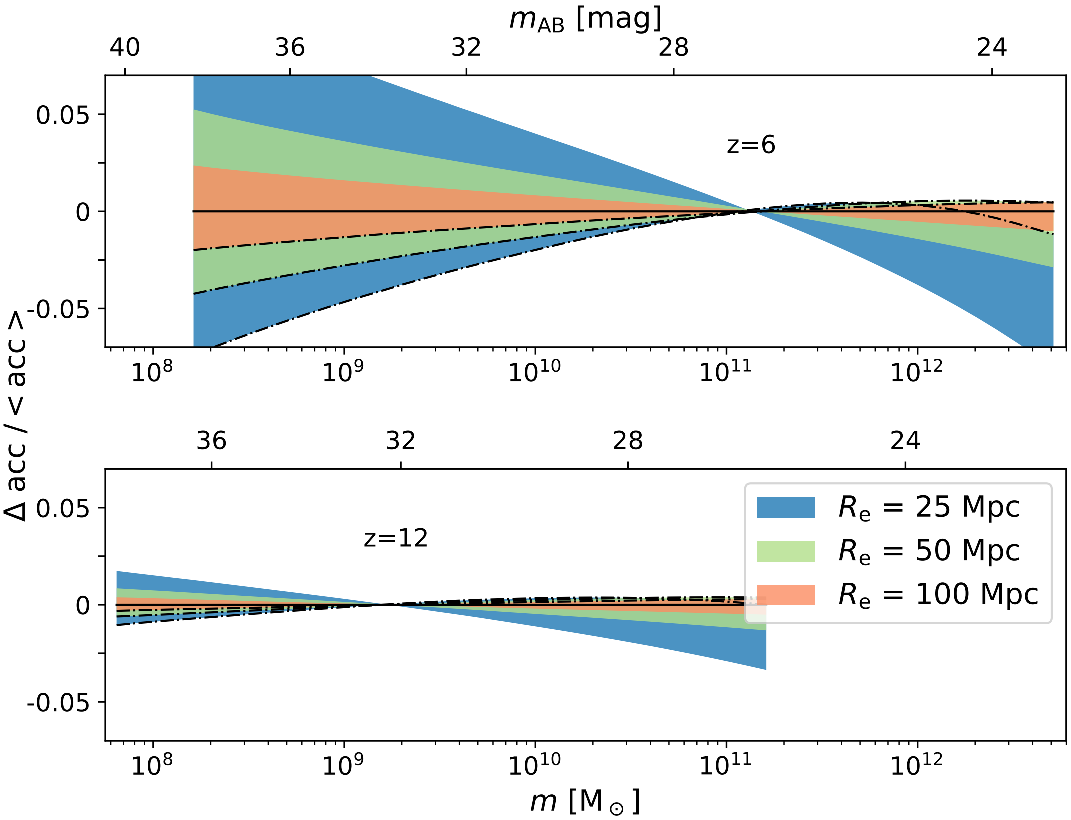

Figure 2.2: Accretion rate scatter (shaded areas) due to cosmic variance compared to the average accretion rate of haloes (solid black line) for various scales (identified with their radius R_e). Adding cosmic variance to our models changes accretion rates by < 5 per cent for most haloes; smaller regions have larger variance in accretion. The dot dashed lines show the accretion of a 1σ overdense region, and the opposite end of the shaded regions show the 1σ underdense region. Higher mass haloes overaccrete in overdense environments and underaccrete in underdense environments. For lower mass haloes, the opposite is true. The threshold mass where haloes in all environments accrete nearly equally evolves to lower mass with redshift. The top axes show the approximate apparent magnitude of the haloes “m_{AB}” (we assign haloes their magnitudes in section 2.4).

as a multiple of the star formation rate $\dot{m}_{\rm w} = \eta(m, z)\ \dot{m}_*$, yielding

$$f_* = \frac{1}{1 + \eta(m, z)}. \tag{2.12}$$

Many models suggest that massive haloes accrete gas more slowly than our simple argument suggests, because of the heating at the virial shock. Furlanetto et al. (2017) show that virial shock heating only has a modest effect on the results of this model, but we include it because it helps match the observed densities at large luminosities. Faucher-Giguère et al. (2011) show the fraction of gas that can cool onto a galaxy in the presence of a virial shock is [4]

$$f_{shock} = 0.47 \left(\frac{1+z}{4}\right)^{0.38} \left(\frac{m}{10^{12} M_\odot}\right)^{-0.25}. \tag{2.13}$$

Because we only include stellar feedback, which limits star formation at small masses, we also impose a maximum efficiency $f_{*,max}$ that limits star formation when $\eta(m, z) \to 0$ at large halo masses. We impose it in a way that keeps f_* smoothly differentiable. Combining f_{shock} and $f_{*,max}$ with equation 2.12 gives

$$f_* = \frac{f_{shock}}{f_{*,max}^{-1} + \eta(m, z)}. \tag{2.14}$$

Finally, we parameterize the strength of stellar feedback, $\eta(m, z)$, as

$$\eta = C \left(\frac{10^{11.5} M_\odot}{m}\right)^{\xi} \left(\frac{9}{1+z}\right)^{\sigma}. \tag{2.15}$$

For energy driven supernova feedback: $C = 1$, $\xi = 2/3$, $\sigma = 1$, and $f_{*,max} = 0.1$. We will also consider a redshift-independent version ($C = 2$, $\xi = 2/3$, $\sigma = 0$, and $f_{*,max} = 0.1$) and a momentum-driven version ($C = 5$, $\xi = 1/3$, $\sigma = 1/2$, and $f_{*,max} = 0.2$) for comparison (for more details on η and its parameterization, see Sun & Furlanetto, 2016; Mirocha et al., 2017; Furlanetto et al., 2017). These alternate parameterizations of η will allow us to test how cosmic variance depends on our galaxy formation model.

This simple model undoubtedly ignores many important elements of galaxy formation, but it suffices to consider a wide range of possible halo mass-luminosity relations. For

---

[4] We require $f_{shock} \leq 1$, and we smooth the function near where $f_{shock} \to 1$ in order to ensure that it is smoothly differentiable.

example, we do not take into account that gas should cycle through the interstellar medium (ISM) before forming stars. However, in the 'bath tub' model of galaxy formation, galaxies evolve towards a quasi-equilibrium state between mass accretion and star formation such that the ISM maintains roughly constant mass (Dekel & Mandelker, 2014). Once this equilibrium is reached, our model more accurately describes star formation.

2.4.2 From star formation to luminosity

We now convert star formation rate to UV luminosity. UV luminosity is a good tracer of star formation because it is produced only by massive, short-lived stars. We take the standard conversion

$$\dot{m}_* = K_{UV} \times L_{UV}, \tag{2.16}$$

where ${L_{UV}}^5$ is the rest-frame continuum ($1500 - 2800$ Å)[6] intrinsic luminosity (without extinction). K_{UV} is a conversion from luminosity to star formation rate, and it is dependent on the initial mass function, metallicity, star formation history, binaries, etc. We take $K_{UV} = 1.15 \times 10^{-28} M_\odot yr^{-1}/(erg\ s^{-1} Hz^{-1})$ from Madau & Dickinson (2014). We will show that K_{UV} will not substantially affect our predictions for the relative cosmic variance of the UVLF, even though it can have significant effects on the UVLF itself.

We do ignore dust in our fiducial model, because the extinction in these sources is only poorly constrained. Models suggest that it is modest, and most importantly is not a strong function of halo mass (Mirocha et al., 2020). We do however test the effects of dust on our results at $z < 8$ using an empirical dust correction (Vogelsberger et al., 2020, "Model A"). In order to match the data when applying this dust correction to our energy-driven model, we set $C = 2$ and $f_{*,max} = 0.3$, making the galaxies intrinsically brighter at fixed halo mass. For simplicity, we ignore f_{shock} in this case, because it also affects the bright end of the luminosity function.

---

[5] For the remainder of the paper, we will display luminosity as absolute and apparent AB magnitudes ($M_{\rm AB}$ and $m_{\rm AB}$).

[6] This wavelength range corresponds to H-band in the redshift range of $z \approx 5 - 9$, and K-band for $z \approx 8 - 12$.

We also ignore scatter in the halo mass to luminosity relation, which would have the effect of flattening out the exponential drop off of the UVLF, as upward scatter in the luminosity has a larger relative effect on the luminosity function in that regime. While we expect this effect to be small (at least on population-level statistics such as the UVLF), we plan to explore it in the future by introducing a scatter in the accretion rates and/or star formation rates.

2.5 Cosmic Variance in the UVLF

In this section we present the conditional UVLF generated by our model. We provide a fit to the conditional UVLF with a simple Gaussian approximation. We then test the robustness of our results against model choices. Finally, we compare our results to recent works.

We show the conditional UVLF $\phi_{\mathrm{cond}}(M_{\mathrm{AB}}, z, \delta_b, R_{\mathrm{e}})$ and its 2σ scatter due to cosmic variance in Figure 2.3. As in the CMF, cosmic variance increases with increasing galaxy luminosity and also with increasing redshift. The data points shown in Figure 2.3 are from Bouwens et al. (2015) and Bouwens et al. (2016); for a more in-depth analysis of this model's agreement with current data, see Furlanetto et al. (2017).

The mapping from halo mass to luminosity in our model is nearly independent of environment because accretion is also nearly independent of environment (see Fig. 2.2). Thus, nearly all of the cosmic variance of the UVLF comes directly from the variance in the CMF (see Fig. 2.1). Similarly, simulations by Lovell et al. (2021) find the star formation rate of a galaxy is independent of the dark matter environment, although, they compare star formation rates at fixed stellar mass, not halo mass.

2.5.1 Calculating cosmic variance

As shown in Figure 2.3, the amount of cosmic variance in a given galaxy formation model will depend on luminosity, redshift, and the survey characteristics. In this section we provide a simple descriptor of cosmic variance across all these parameters. We quantify cosmic variance

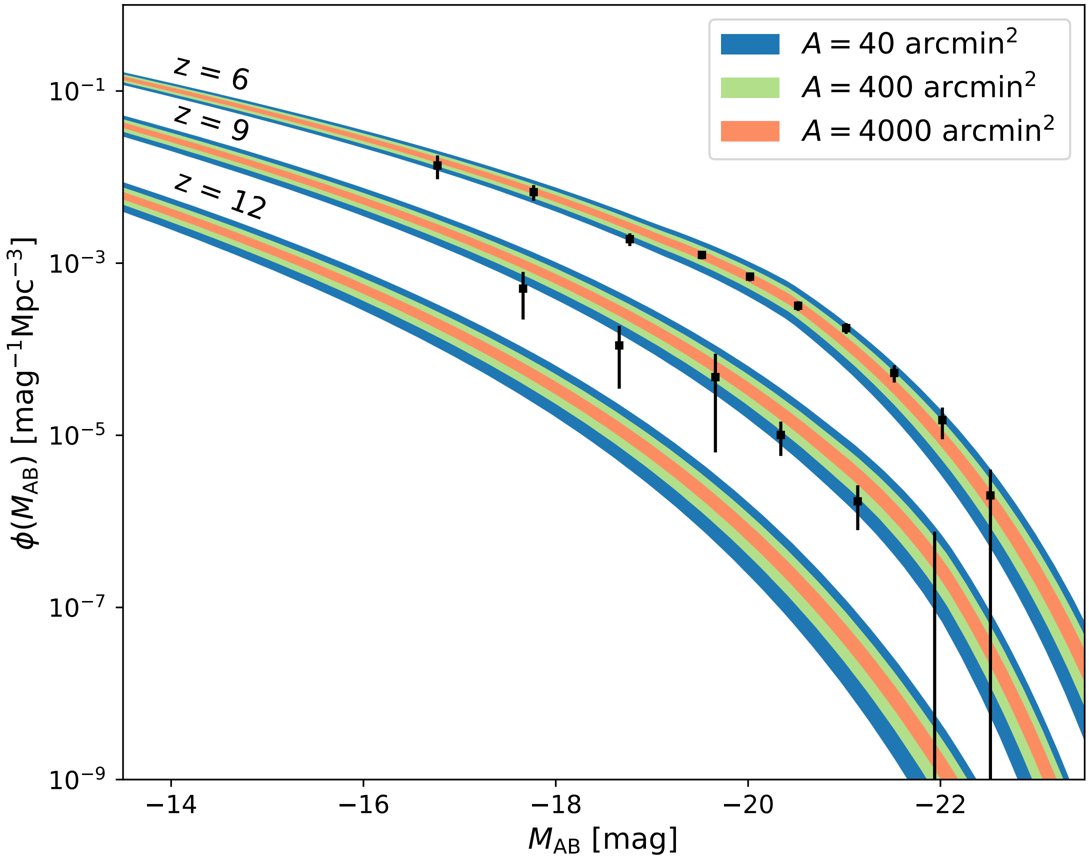

Figure 2.3: The UVLF and its 2σ scatter from cosmic variance for three selections of survey area at three redshifts ($z = 6$, 9, and 12; $z = 9$ and 12 are off-set in log space by -0.5 dex and -1 dex, respectively). The scatter in the UVLF increases at the bright end and for smaller survey areas, similar to the CMF in Figure 2.1. The survey areas $A = 40$, 400, 4000 arcmin^2 have volumes equivalent to spheres with radii $R_\mathrm{e} \approx 29$, 63, 135 Mpc at $z = 6$ (redshift bin $\Delta z = 1$). The data points are from Bouwens et al. (2015) and Bouwens et al. (2016).

ε_{cv} as the relative standard deviation of the conditional UVLF at fixed redshift, apparent magnitude, survey area, and redshift bin width[7]:

$$\varepsilon_{cv}^2 = \frac{\langle \phi_{\rm cond}^2 \rangle - \langle \phi_{\rm cond} \rangle^2}{\langle \phi_{\rm cond} \rangle^2}, \tag{2.17}$$

with $\langle \phi_{\rm cond}^n \rangle$ defined as

$$\langle \phi_{\rm cond}^n \rangle = \int \phi_{\rm cond}^n (m_{\rm AB}, z, \delta_b, R_{\rm e}) \times p(\delta_b | R_{\rm e}, z) d\delta_b, \tag{2.18}$$

where $R_{\rm e}$ is determined from the survey area and redshift bin width as described in Section 2.3.3. Figure 2.4 shows ε_{cv} as a function of survey area for various redshifts and apparent magnitudes (all with $\Delta z = 1$). This definition of ε_{cv} uses $\phi_{\rm cond}$ at fixed z, but applies it to the entire volume defined by A and Δz. This approximation breaks down if cosmic variance evolves significantly over the range defined by Δz. Thus, the choice of Δz should be made with care, especially at lower z where ε_{cv} evolves most rapidly in a relative sense (ε_{cv} evolves more rapidly in an absolute sense at high z). Over $\Delta z = 1$, ε_{cv} evolves 10% – 30% (at $z =$ 14 & 5, respectively). Choosing $\Delta z < 0.5$ keeps the change in ε_{cv} below 10% for most cases[8].

Figure 2.4 shows that the relative importance of cosmic variance varies widely across the galaxy population, with a strong dependence on survey parameters. Using a redshift bin width $\Delta z = 1$, ε_{cv} is low at the faint end of the UVLF ($m_{\rm AB} = 32$), ranging from $\sim 5\%$ at large survey area (1000 arcmin^2) to $\sim 15\%$ at small survey area (1 arcmin^2) at $z = 6$. As redshift increases, so does ε_{cv}; at $z = 12$ and $m_{\rm AB} = 32$, ε_{cv} ranges from $\sim 12\%$ at large survey area (1000 arcmin^2) to $\sim 35\%$ at small survey area (1 arcmin^2).

Cosmic variance also increases significantly at the bright end of the UVLF. At $m_{\rm AB} = 26$, ε_{cv} ranges from $\sim 15\%$ at large survey area (1000 arcmin^2) to $\sim 40\%$ at small survey area (1 arcmin^2) at $z = 6$. At $z = 12$ and $m_{\rm AB} = 26$, ε_{cv} ranges from $\sim 30\%$ at large survey area (1000 arcmin^2) to $\sim 90\%$ at small survey area (1 arcmin^2).

Cosmic variance flattens out at small survey areas, which is largely due to the effects of the pencil-beam shape of surveys. Even at small survey areas, such a geometry still contains

---

[7] Our definition of $m_{\rm AB}$ assumes the galaxy is at the specified 'fixed redshift', regardless of redshift bin width.

[8] Our public Python package `galcv` can be used to explore the evolution of ε_{cv} over any desired parameter.

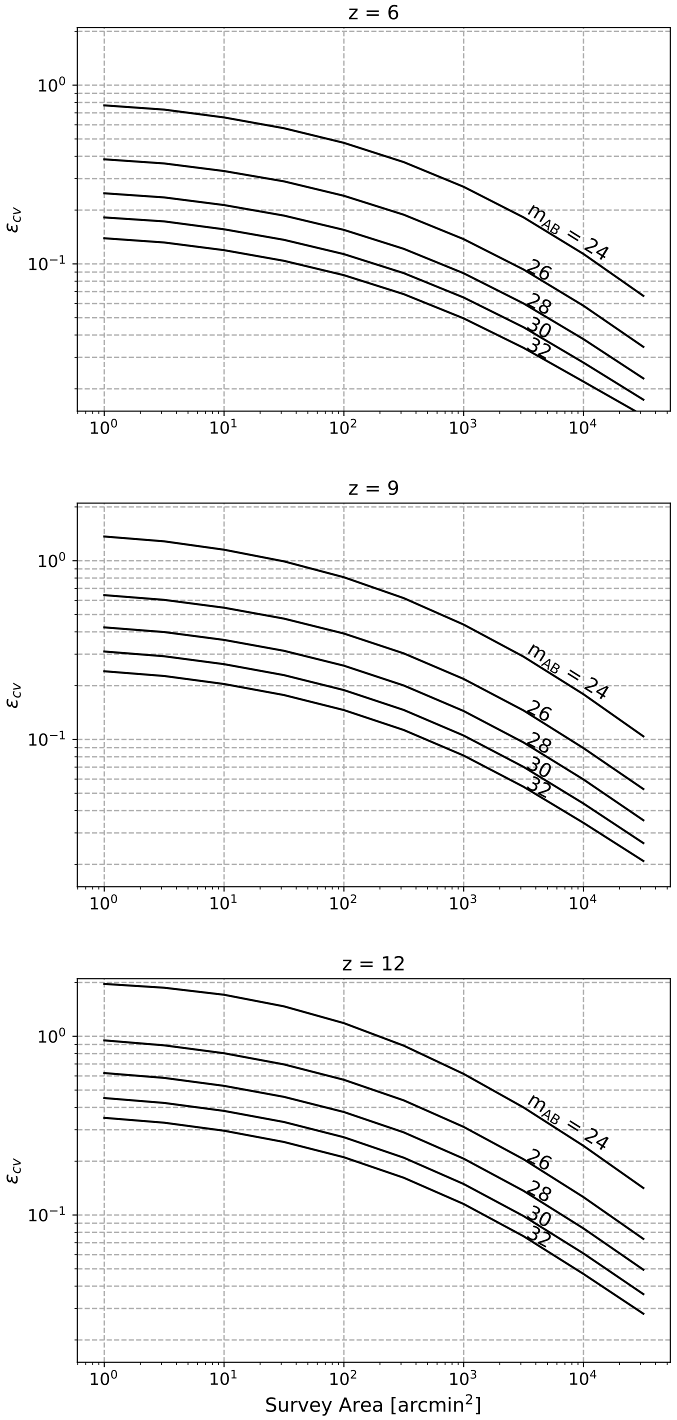

Figure 2.4: The relative cosmic variance of the UVLF ε_{cv} as a function of survey area for various apparent magnitudes with a redshift bin width of $\Delta z = 1$ (*black lines*; shown from apparent magnitude $m_{\rm AB} = 32$ on the bottom, decreasing by $\Delta m_{\rm AB} = 2$ towards the top).

Table 2.1: Parameters for fits to ε_{cv} (eq. 2.19). We provide ε_{cv} on a much wider range of parameters via our public Python package `galcv` (see Data Availability section).

Redshift	App. UV mag.	Fit parameters		
z	$\mathrm{m_{AB}}$	Ψ	γ	b
6	32	-0.223	0.167	-0.608
	30	-0.189	0.184	-0.529
	28	-0.174	0.192	-0.410
	26	-0.161	0.200	-0.235
	24	-0.165	0.199	0.070
9	32	-0.198	0.184	-0.399
	30	-0.190	0.188	-0.298
	28	-0.178	0.194	-0.175
	26	-0.173	0.197	-0.002
	24	-0.197	0.188	0.353
12	32	-0.195	0.188	-0.240
	30	-0.188	0.191	-0.137
	28	-0.184	0.193	-0.001
	26	-0.185	0.193	0.182
	24	-0.202	0.189	0.533

a relatively large range of environments due to its elongated shape. This effect keeps cosmic variance much lower than what one would obtain with a spherical region of the same volume.

We approximate ε_{cv} with a simple functional form; a polynomial in $\log_{10}(\varepsilon_{cv})$ fits well:

$$\log_{10}(\varepsilon_{cv}) \approx \Psi A^{\gamma} + b, \tag{2.19}$$

where A is in arcmin^2, and $\Psi,\ \gamma,$ and b are fit parameters[9]. Table 2.1 displays the parameter fits at a selection of redshifts and magnitudes. These fits have a typical/maximum fractional error of 3/5%. We provide ε_{cv} for a wider range of parameters via a public python package `galcv` (see Data Availability section for more details).

With ε_{cv}, we define a linear approximation of the conditional UVLF in a region with density δ_b, angular extent A, and redshift bin width Δz:

$$\begin{aligned} &\phi_{\text{cond}}(m_{\text{AB}}, z, \delta_b, A, \Delta z) = \\ &\langle \phi(m_{\text{AB}}, z) \rangle \left[1 + \varepsilon_{cv}(A, m_{\text{AB}}, z, \Delta z) \frac{\delta_b}{\sigma_{\text{PB}}} \right], \end{aligned} \tag{2.20}$$

where $\langle \phi(m_{\text{AB}}, z) \rangle$ is the average UVLF and $\delta_b/\sigma_{\text{PB}}$ is the density of the region relative to a 1σ fluctuation. This conditional UVLF is similar in construction to that in Livermore et al. (2017), used to fit to lensed high-z galaxy survey data.

2.5.2 Parameter dependence of ε_{cv}

Our calculations so far have assumed our fiducial choices for the galaxy model (assuming the minimalist energy-regulated prescription) and mass function parameters. Here, we explore how sensitive our results are to variations in these assumptions.

First, we consider how cosmic variance depends on the star formation model. Figure 2.5 shows shows the difference in ε_{cv} when using our fiducial energy-regulated feedback (solid lines) vs a redshift-independent version of feedback (dotted lines, see Section 2.4.1). There is little difference in the predictions for ε_{cv}. We find a similarly small difference when using momentum-regulated feedback and when using the dust correction from Vogelsberger et al.

---

[9] Note that our fit assumes the survey subtends a square area on the sky.

(2020) for $z < 8$ (the effects of these two are similar to the dotted lines; so are not plotted to reduce clutter). Also, our choice of K_{UV} (see eq. 2.16) will not significantly affect our results, as ε_{cv} is not a particularly strong function of magnitude. These results suggest that cosmic variance is not strongly dependent on the details of star formation or dust.

Second, we explore if cosmic variance is strongly affected by large-scale galaxy environment, namely through differences in accretion. With the linear halo bias function from equation (2.5), we approximate

$$\varepsilon_{cv} \approx b_{\rm Trac}\sigma_{\rm PB}, \tag{2.21}$$

and show it in Figure 2.5 (faded solid line)[10]. While our full model allows for galaxies to have an environment-dependent accretion and thus luminosity, this linear method does not. However, it provides very similar results to the full method, though it slightly underpredicts cosmic variance at the bright end and overpredicts at the faint end due to the variance in accretion for those haloes (see Fig. 2.2). This result suggests that approximating the CMF via a simple bias factor is sufficient to capture the effects of cosmic variance (at least to linear order; large density excursions are discussed later in this section).

We conclude that the level of cosmic variance is not sensitive to the particulars of the galaxy formation model. Rather, cosmic variance is dominated by the underlying conditional halo mass function. In section 2.3.1, we described an alternate method of creating a CMF: scaling the Trac et al. (2015) mass function by the conditional Press & Schechter (1974) mass function (eq. 2.7). In Figure 2.5 we show ε_{cv} when using that CMF (dashed lines). This change results in $\sim 25\%$ more cosmic variance across the board, the largest effect of any model choice. Thus, in our model, the biggest uncertainty in ε_{cv} is in our understanding of the CMF.

While equation (2.20) provides a good approximation to the conditional UVLF for $\delta_b/\sigma_{\rm PB} \lesssim 2$, the assumption of a Gaussian bias distribution breaks down at larger density excursions. Figure 2.6 shows (at $z = 9$) the difference between using equation (2.20) (*dashed lines)* and our full treatment (*solid lines*) for the conditional UVLF (ϕ) for a 3σ

---

[10] We connect the bias function $b_{\rm Trac}$ to galaxies using our model's average halo mass–UV luminosity relation

density excursion. Equation (2.20) underestimates the number of galaxies in very underdense regions (even giving unphysical negative densities at the smallest survey areas), and it also underestimates the number of galaxies in very overdense regions. Equation (2.20) underpredicts the number of galaxies in both wings, even though it is more reliable near $\delta_b = 0$, because the true bias distribution (at fixed magnitude) is closer to a log-normal. However, where the deviation from the Gaussian approximation is most pronounced, Poisson shot noise usually dominates the error. Thus, for most applications, equation (2.20) (ε_{cv}) adequately captures the behaviour of the conditional UVLF.

2.5.3 Comparison to other works

Here we compare our predictions of cosmic variance to those from two recent models in the literature: Bhowmick et al. (2020) and Ucci et al. (2021).

Bhowmick et al. (2020) provide public estimates for cosmic variance in a redshift range $z = 7$–14 and for apparent H-band magnitudes between $m_{\rm AB} = 25$ and 30. They determine cosmic variance first by calculating the two-point correlation function of galaxies in their simulation box. They then fit the correlation function to a power law and integrate it in a pencil-beam volume (see their eq. 2) to estimate the relative cosmic variance . They provide estimates of cosmic variance for all sources *brighter* than the listed magnitude, rather than for sources *at* the listed magnitude. This choice means their estimates of cosmic variance are higher than they would be at fixed magnitude, as cosmic variance increases for brighter sources. However, when we mimic this cumulative method, we find the effect is relatively small (cosmic variance $\lesssim 10\%$ larger than fixed magnitude method).

Ucci et al. (2021) provide public estimates for cosmic variance in a redshift range $z = 6$–12 and for apparent magnitudes $m_{\rm AB} = 24$–38 (at $z = 9$). Ucci et al. (2021) calculate cosmic variance as the relative standard deviation of galaxy number counts in many pencil-beam sub-volumes of their simulation box. They have a slightly more limited survey area coverage, providing estimates between $A = 1$ and 1000 arcmin^2. We compare to the predictions from their "photoionization" model. These predictions include Poisson variance, making them an

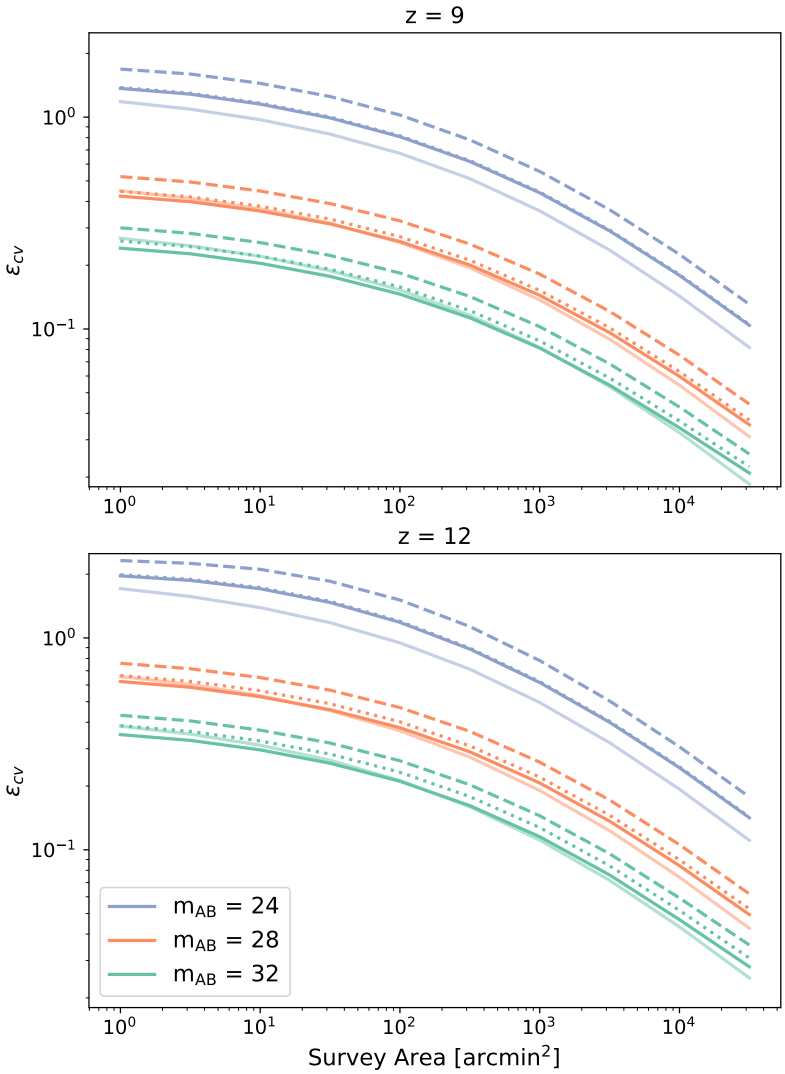

Figure 2.5: The relative cosmic variance of the UVLF ε_{cv} as a function of survey area for various apparent magnitudes (*solid lines*, same as Fig. 2.4). The *dotted lines* show the effects of switching to a z-independent version of star formation. The *faded solid lines* show the linear bias method for estimating cosmic variance (eq. 2.21). The *dashed lines* show the effects of using a different method for creating the CMF, specifically the "Press-Schechter scaling" approach applied to the Trac et al. (2015) mass function (see eq. 2.7). The three sets of lines correspond to magnitudes 32, 28, and 24 (bottom to top).

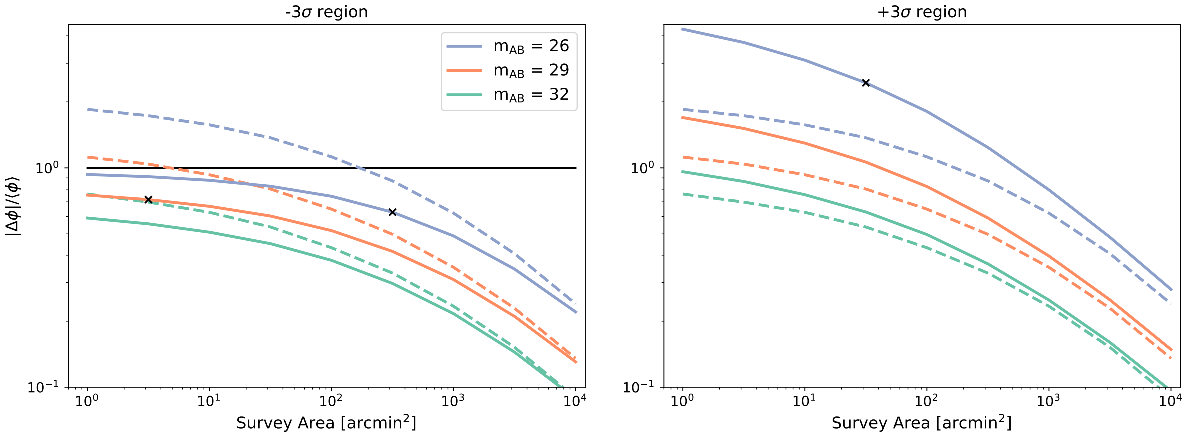

Figure 2.6: The relative excursion from the average UVLF $\langle\phi\rangle$ for a 3σ underdense region (*left*) and a 3σ overdense region (*right*) at redshift 9. We compare the full treatment (*solid lines*) and the linear approximation in equation (2.20) (*dashed lines*). In the left panel, the linear approximation predicts there will be fewer galaxies than the full approach, and can even predict unphysical negative galaxy number densities (when above the horizontal black line). In the right panel, again the linear approximation underpredicts the expected number of galaxies. The black 'x' marks the survey area where our model predicts there to be $\sim$1 source in the corresponding magnitude bin. The solid and dashed lines show apparent magnitudes 32, 29, and 26 (bottom to top).

estimate of the total variance rather than just cosmic variance.

Figure 2.7 shows our predictions compared to those of Bhowmick et al. (2020) (blue dotted lines) and Ucci et al. (2021) (yellow dashed lines) at $z = 9$, with a redshift window of $\Delta z = 1$, at apparent magnitudes of $m_{\rm AB} = 30$ (lower, thick curves) and $m_{\rm AB} = 27$ (upper, thin curves).

As Ucci et al. (2021) report total variance, their predictions should be compared with the red dashed lines (our prediction plus Poisson noise). Our predictions agree closely with those of Ucci et al. (2021) at $m_{\rm AB} = 30$, and agree within $\sim$50% at $m_{\rm AB} = 27$ (though worsening towards low survey area). Our predictions diverge more significantly at a survey area of 100 arcmin^2, where Ucci et al. (2021) have the fewest independent volumes in their simulation. Also, note that differences in the underlying UVLF can strongly affect the strength of Poisson noise; when Poisson noise begins to dominate, our results should not be too closely compared with those of Ucci et al. (2021).

Our results (red solid lines) are systematically lower than those from Bhowmick et al. (2020) (dotted blue lines). However, our predictions remain within $\sim$25% of each other except at smaller survey areas. At $m_{\rm AB} = 30$, our predictions diverge at an area of ~ 10 arcmin^2. At $m_{\rm AB} = 27$, our predictions diverge for survey areas where Poisson noise begins to dominate.

Numerical simulations have the benefit of being able to capture the non-linear bias of haloes. This effect, along with differences in Poisson noise from differing mass functions, could help explain the discrepancy between our predictions and those of the simulations at small survey areas.

In comparison to estimates with numerical simulations, the principal benefit of our model is its flexibility. We can test our model with any mass function or star formation and feedback prescription. Simulations must also subtract their intrinsic Poisson noise (which is not known perfectly) to estimate the cosmic variance, while analytic models can easily separate the two effects. Finally, we cover a wider range of redshifts ($z = 5$–15) and magnitudes ($m_{\rm AB} = 22$–38), and we can study larger volumes than simulations.

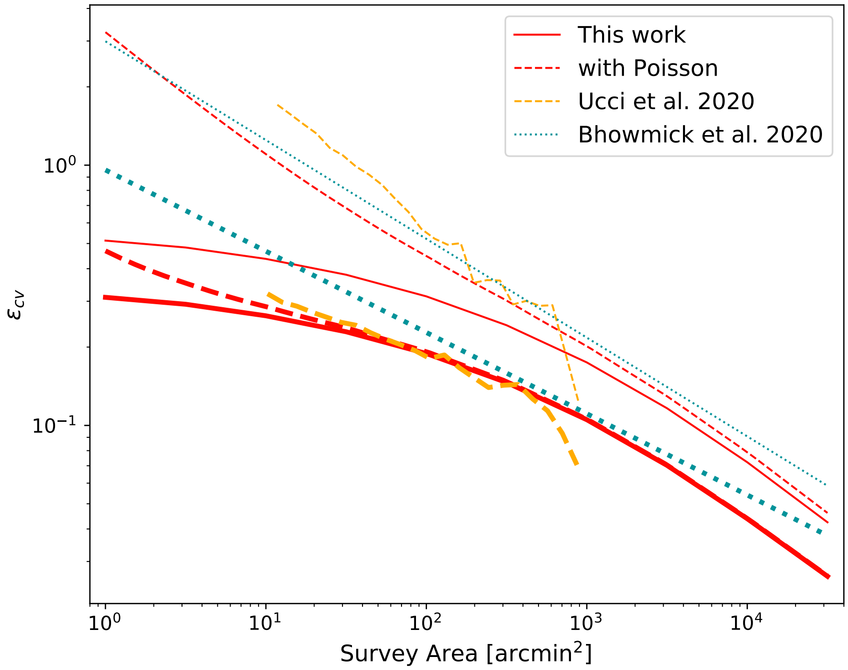

Figure 2.7: Comparison of cosmic variance predictions at $z = 9$ (with a redshift window of $\Delta z = 1$). The lower set of lines (thick) is at an apparent magnitude of $m_{\mathrm{AB}} = 30$, while the upper set of lines (thin) is at $m_{\mathrm{AB}} = 27$. The dashed red lines show our cosmic variance predictions with Poisson noise added, for comparison with Ucci et al. (2021).

Our results agree quite well with those of Ucci et al. (2021), over the range to which we can compare, especially at faint luminosities. For bright sources, our estimates are slightly below theirs, but the discrepancy is comparable to the apparent uncertainty in the CMF (~25%; see Section 2.5.2). We agree reasonably well with Bhowmick et al. (2020) on large scales as well.

2.6 Impact on Future Surveys

Cosmic variance will provide an unavoidable source of error for next generation telescopes, especially at the highest redshifts. It will dominate over Poisson noise for all but the brightest sources, and it is not easily avoided with deeper observations. Instead, it can only be minimized by probing larger volumes (at the cost of missing the more numerous faint sources)

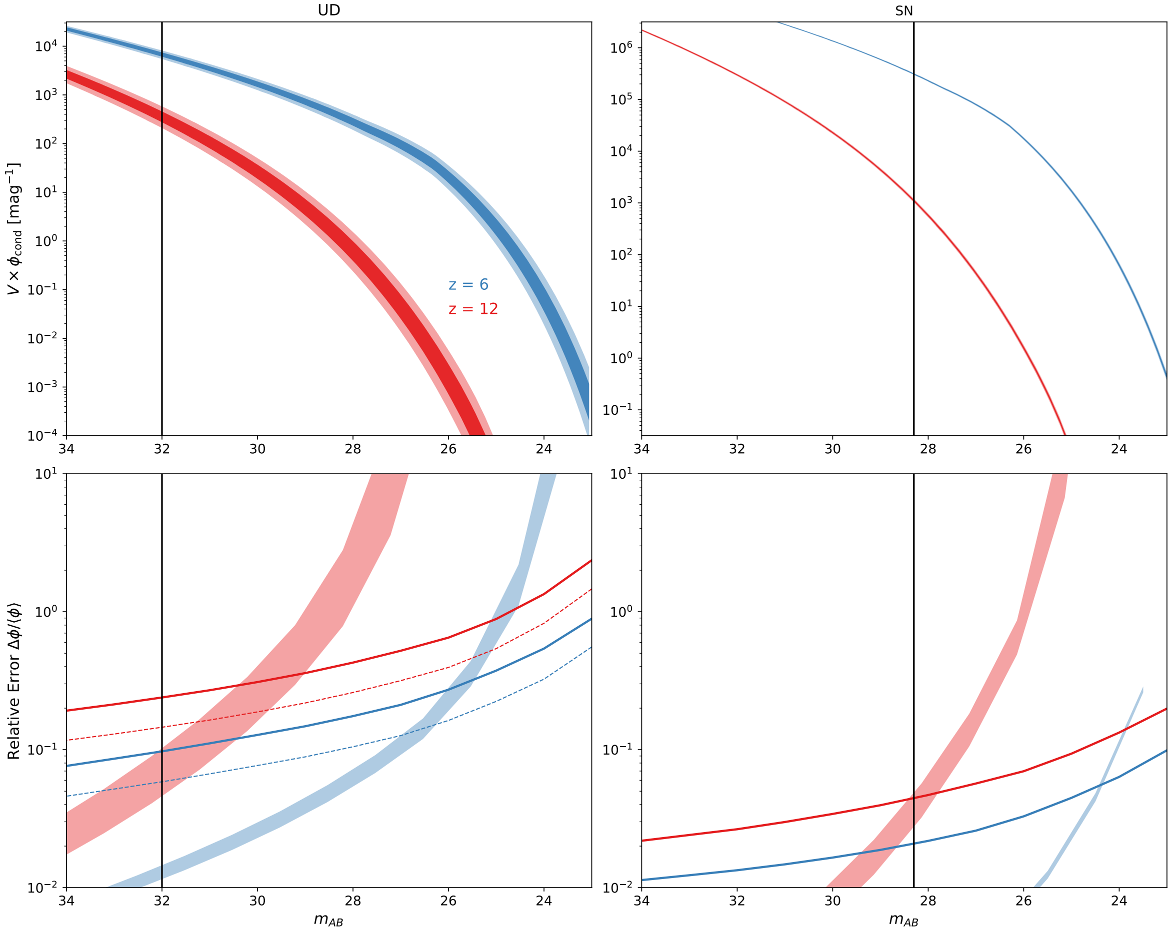

Figure 2.8: The effects of cosmic variance on the UVLFs of two high redshift surveys (UD and SN). In the upper panel, the width of the curves represents the 1σ and 2σ (inner and outer shading) ranges of intrinsic UVLFs that could be found in that survey's volume. In the lower panels, the lines show ε_{cv} (upper set of lines are at $z = 12$, lower set at $z = 6$). If the UD survey is broken up into four independent pointings of JWST, each smaller sub-pointing has a higher variance than a large mosaic, but they may be combined for an overall reduction in measuring the average UVLF. This improvement is represented with the thin dashed lines. The vertical black lines are the magnitude limits of the surveys. The shaded band provides an estimate of Poisson shot noise. Cosmic variance acts as a noise floor for measuring the average UVLF, bounded on the faint end by the magnitude limit, and on the bright end by Poisson noise, except for the SN survey, which is entirely dominated by Poisson noise at high redshift.

or by splitting up surveys into multiple independent pointings (at the cost of missing large-scale structure and making clustering measurements more difficult).

In this section, we perform a case study of the effects of cosmic variance on two upcoming instruments, JWST and the Roman Space Telescope. We consider two potential high redshift surveys: a JWST ultradeep (UD) survey (following Mason et al., 2015) along with a much wider-field Roman Space Telescope survey (similar to their planned supernova survey, and which we refer to as our SN survey). The UD survey has a detection limit of $m_{\rm AB} \approx 32.0$ and survey area of $A = 40$ arcmin^2, while the SN survey has a detection limit of $m_{\rm AB} \approx 28.3$ and a survey area of $A = 9$ deg^2. To begin, we assume that both are performed over a single contiguous area, requiring at minimum four and 30 separate pointings (neglecting overlap between the pointings).

Additionally, we present a method of fitting an average UVLF to data from the UD and SN surveys simultaneously. Our method makes use of our model effectively as a “prior” on the cosmic variance in each survey field. The fitting process also accounts for the difference in shape between the local UVLF in each field and the average UVLF that we wish to fit.

2.6.1 Effects of cosmic variance on UD and SN surveys

We show the effects of cosmic variance on the UVLF of the UD and SN surveys in Figure 2.8. The upper panels show the 1 and 2σ fluctuations of the UVLF at $z = 6$ and $z = 12$. The lower panels show ε_{cv} for these surveys (lines) and Poisson shot noise (shaded bands[11]). The vertical lines denote the magnitude limit of the surveys.

A given survey has access to the UVLF over a limited magnitude range, bound on the faint side by the magnitude limit and on the bright side by Poisson noise. In between, the noise floor of cosmic variance determines the maximum accuracy one can achieve in measuring the average UVLF over the accessible magnitude range if using just the one survey.

Splitting up a survey into independent pointings can improve the measurement of the

---

[11] Poisson shot noise is model-dependent, so we represent it as a band that encompasses the predictions from the three different feedback prescriptions described in section 2.4 as well as the variety of number counts predicted from cosmic variance itself.

average UVLF. While each individual pointing has higher cosmic variance than a large mosaic, they may be combined, which results in a reduction by the square root of the number of fields. The effect of splitting the UD survey into 4 pointings is represented by the dashed lines in Figure 2.8. Robertson (2010) found a similar improvement in cosmic variance when splitting surveys into independent volumes. See Section 2.6.4 for more details on the benefits/drawbacks of splitting up surveys

When interpreting survey results, it is crucial to note that cosmic variance is correlated across all magnitudes. If a survey probes a 1σ underdense region, the expected number counts in *each* magnitude bin will be below the average by 1σ. In contrast, Poisson noise is uncorrelated between each magnitude bin, depending only on the expected number of sources in that bin.

2.6.2 Measuring the average UVLF

Here we introduce a method to account for cosmic variance in measuring the *average* UVLF of the Universe given data from multiple independent survey volumes. As an example of this method, we simulate mock UD and SN surveys of the UVLF and fit a model that extracts the average UVLF parameters; we then repeat this many times and compare those fits to the "true" parameters predicted by our model.

In this section, we model the average UVLF as a modified Schechter function:

$$\phi(L)dL = \frac{\phi^*}{L^*}\left(\frac{L}{L^*}\right)^{\alpha} e^{-(L/L^*)^{\Gamma}} dL, \tag{2.22}$$

where $\phi(L)dL$ is the number density of galaxies with luminosities in the range $(L, L + dL)$, ϕ^* is a normalization constant, L^* is the location of the exponential cutoff, α is the faint end slope, and Γ is a parameter that governs the strength of the exponential cutoff. $\Gamma = 1$ corresponds to a normal Schechter function. Our models are fit best with $\Gamma = 0.5$, so we will use that value for this paper. We note that our use of $\Gamma = 0.5$ predicts a higher number of bright galaxies than a normal Schechter function. This effect is reminiscent of recent studies of very high-z surveys, which have found that the UVLF can be better fit by a double power-law due to an excess of bright galaxies (Bowler et al., 2014, 2020). We do not use a double power law as our models are better fit by the modified Schechter function.

We explore four possible methods to measure the average UVLF of the Universe.

1. *NoCV*: We assume cosmic variance does not exist. Every region of the Universe has the exact same underlying UVLF, so Poisson noise is the only source of error.
2. *Naive*: Cosmic variance exists, but we fit the average UVLF without attempting to account for it.
3. *Standard*: We fit for the average UVLF using a common method to account for cosmic variance
4. *Full*: Our fiducial method. We fit for the average UVLF parameters using the conditional UVLF developed in this paper[12] (eq. 2.20).

The *NoCV* method assumes (unrealistically!) that cosmic variance does not exist. We simulate galaxy counts for the UD and SN surveys by drawing from the average UVLF that our model predicts, adding Poisson noise, and then fitting equation (2.22) to the combined mock data[13]. The solid curves in Figures 2.9 and 2.11 show the probability density functions (pdfs) of the best fits of the UVLF to 2000 sets of simulated data with *no* cosmic variance (for $z = 9$ and 12, *solid lines*). Unsurprisingly, this method recovers the "true" values (black crosses) of the average UVLF parameters well, as the SN probes the bright end and the UD the faint end, with some overlap between. Of course, cosmic variance *does* exist; this method is only to be used as a comparison to our more realistic scenarios.

For the other three methods we use our model to simulate data for each survey, including cosmic variance. We first draw from the distribution of possible density environments for the UD survey $p(\delta_b|R_{\rm e}, z)$ (see Appendix 2.8) and then use equation (2.20) to generate the UVLF

---

[12] A similar method is implemented in Livermore et al. (2017); they consider cosmic variance in lensed surveys, and construct a conditional luminosity function using cosmic variance estimates from Robertson et al. (2014).

[13] We assume in this paper that the UD and SN surveys are perfect, in that they detect every galaxy and are able to accurately place each source in a magnitude bin of width $\Delta m_{\rm AB} = 0.5$ and a redshift bin of $\Delta z = 1$. These are clearly not all accurate assumptions, especially the first one, but this treatment may be taken as a best possible scenario.

for that survey[14]. We then calculate the expected number of galaxies in each magnitude bin and apply Poisson shot noise. We repeat these steps for the SN survey. We then repeat this process 2000 times to generate many possible pairs of surveys.

In the *Naive* method, we simply joint fit equation (2.22) to the 2000 UD+SN mock data pairs with no attempt to correct for cosmic variance. The red dotted lines in Figures 2.9–2.12 show the resulting best fit pdf for the average UVLF. The recovered parameter range is far wider than the *NoCV* method because the *Naive* method completely ignores the effects of cosmic variance; the measured luminosity functions in the two surveys are *not* the same so cannot easily be reconciled by a single fit.

The *Standard* method, originally developed by Sandage et al. (1979) and used by e.g. Efstathiou et al. (1988) and Bouwens et al. (2015), fits a universal *shape* to the UVLF, ignoring the field-to-field normalization. Then, the normalization is fixed at the end to reproduce the correct total number of galaxies across all surveys. Using this method, we fit to the mock data with cosmic variance. The blue dashed lines in Figures 2.10 and 2.12 show the pdfs of the best fit parameters for the average UVLF. This method recovers the average UVLF parameters much more accurately than the *Naive* method.

While the *Standard* method is relatively robust to cosmic variance, it does not take into account any changes in the shape of the UVLF due to environment. Additionally, it does not incorporate any information about expected levels of cosmic variance, and it can produce biased results, as seen in this example by its systematic underprediction of the values of ϕ^* and α (most noticeably in Fig. 2.10, upper-left panel).

Finally, in the *Full* method, we fit equation (2.20) (with the modified Schecter function in eq. 2.22 as $\langle\phi(m_{\rm AB})\rangle$) simultaneously to each of the pairs of mock surveys, allowing for different values of δ_b for each survey. The green solid lines in Figures 2.10 and 2.12 show the pdfs of best fit parameters for the average UVLF. Unsurprisingly (because we are fitting with the same function used to generate the mock data), the "true" parameters are recovered better than with the *Standard* method.

---

[14] For $\langle\phi(m_{\rm AB})\rangle$ in equation (2.20), we use the average UVLF predicted by our model, fit by equation (2.22) to obtain the "true" parameters.

The upper right panels of Figures 2.9–2.12 show the total emissivity of the Universe as inferred from the parameters of the best fit (integrating down to $m_{\rm min}$), compared to the "true" average emmisivity that our model predicts (vertical line). The *Full* method does a slightly better job at recovering the average emissivity of the Universe compared to the *Standard* method, and both do much better than the *Naive* method.

While it is certainly to be expected that the *Full* method performs better than the *Standard* method in our calculations (given that we use our model to generate the mock data *and* to fit to the data), the *Full* method still has benefits. First, it provides estimates of the dark matter overdensity δ_b for each survey field, while the *Standard* method by design throws out field-to-field variance information. Thus, the *Full* method can be used to test our understanding of cosmic variance, because it effectively has a prior on the level of cosmic variance allowed. It penalizes very high field-to-field variance, unlike the *Standard* method that effectively uses a flat prior on the amount of cosmic variance that is allowed during fitting. If real data were fit with the *Full* and *Standard* methods, and the *Standard* method provided a better fit, that would indicate that our understanding of cosmic variance is flawed. We could use our model to investigate *where* and *why* our understanding of cosmic variance breaks down in terms of our physically motivated inputs.

One could investigate the time evolution of UVLF parameters after determining the best fit values at a variety of redshifts. However, this experiment would need to be done with care, as Figures 2.9–2.12 show that the UVLF parameters are highly correlated. Thus, their time evolution must be fit jointly and with a good estimation of their covariance. That covariance can depend strongly on the treatment of cosmic variance.

2.6.3 The benefits of multiple surveys

Next we consider the importance of measuring the UVLF with multiple complementary surveys. Figure 2.13 shows the range of best fit parameters for the $z = 9$ UVLF when fitting to SN survey data alone (*dotted contours*), the UD survey alone (*dashed contours*), and with both simultaneously fit (*solid contours*, identical to those in Figure 2.10). The SN survey

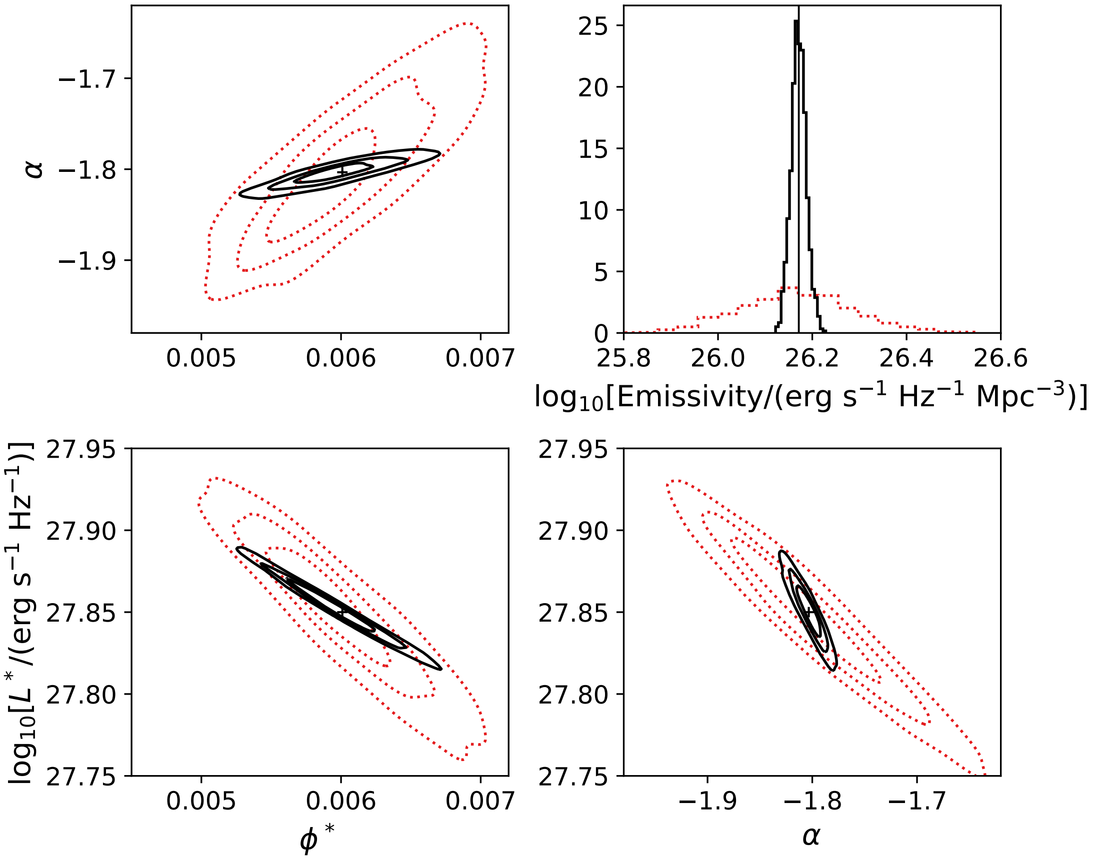

Figure 2.9: Recovered luminosity function parameters and uncertainties in the *NoCV* (*solid contours*) and *Naive* (*red dotted contours*) methods. The contours represent the distribution of best fit average UVLF parameters (see eq. 2.22) for 2000 simulated pairs of UD and SN surveys at $z = 9$. Cosmic variance adds a large amount of uncertainty to the determination of the "true" parameters (*black crosses*) if not treated properly in the fitting technique. The top-right panel shows the emissivity of each fit (integrated down to $m_{\rm min}$), with the "true" average emissivity shown as the vertical line. The contours in this and all other figures are equally spaced between zero and the peak value of each normalized distribution.

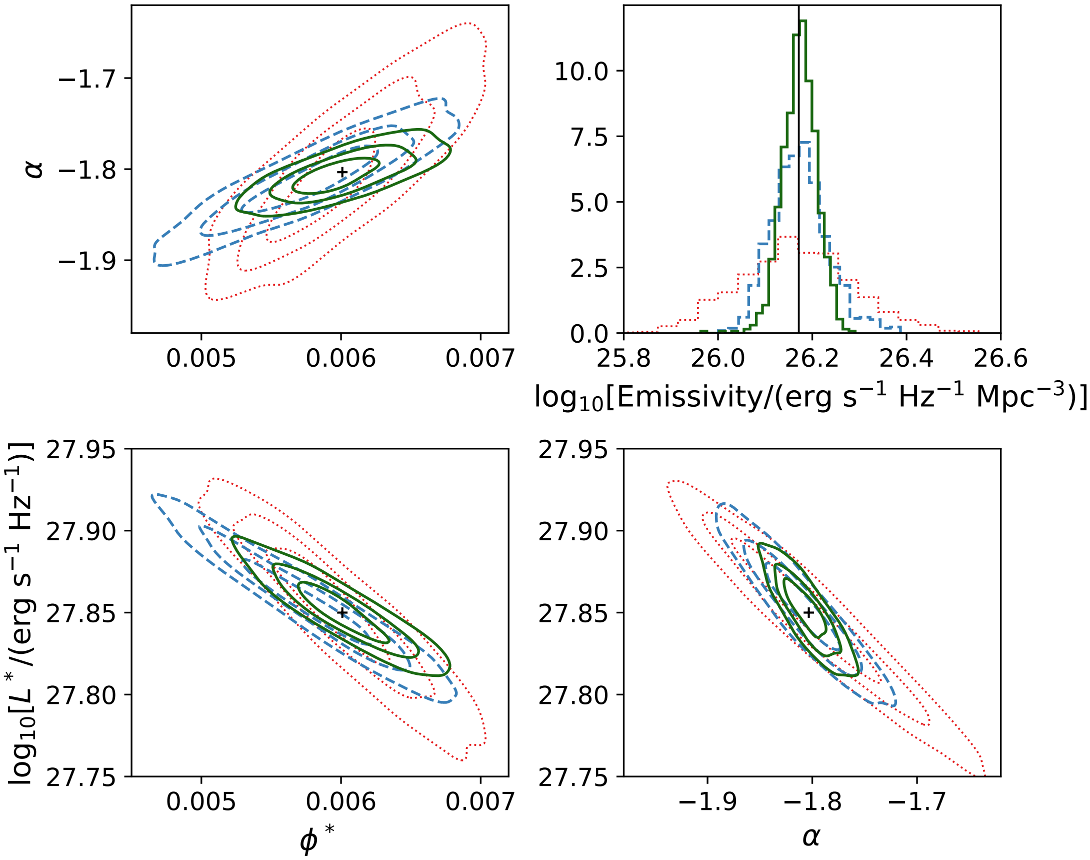

Figure 2.10: Recovered luminosity function parameters and uncertainties in the *Full* (*green solid contours*), *Standard* (*blue dashed contours*), and *Naive* (*red dotted contours*) methods. The contours represent the distribution of best fit average UVLF parameters (see eq. 2.22) for 2000 simulated pairs of UD and SN surveys at $z = 9$. The *Full* and *Standard* methods are significant improvements over the *Naive* method, though the *Full* method does the best job recovering the "true" parameters (*black crosses*). The *Standard* method is also slightly biased towards recovering a high L^*, low ϕ^*, and steeper α in this case.

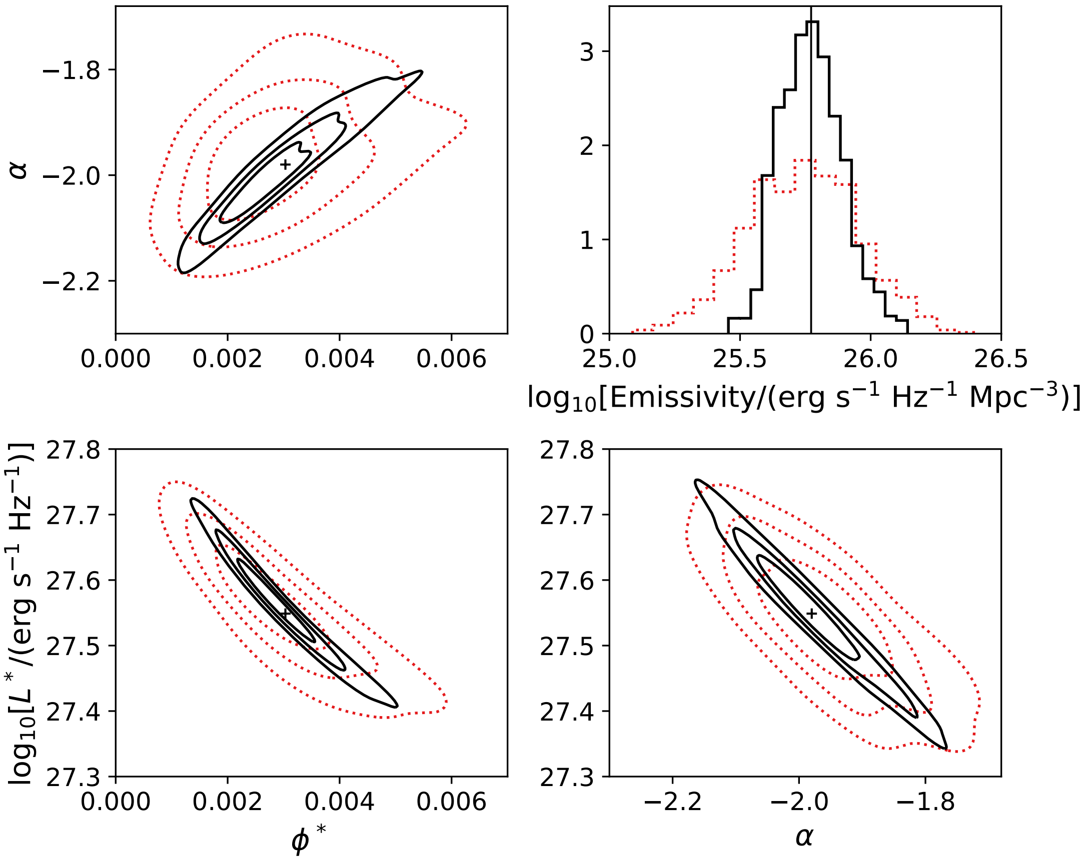

Figure 2.11: The same as Figure 2.9 but at $z = 12$.

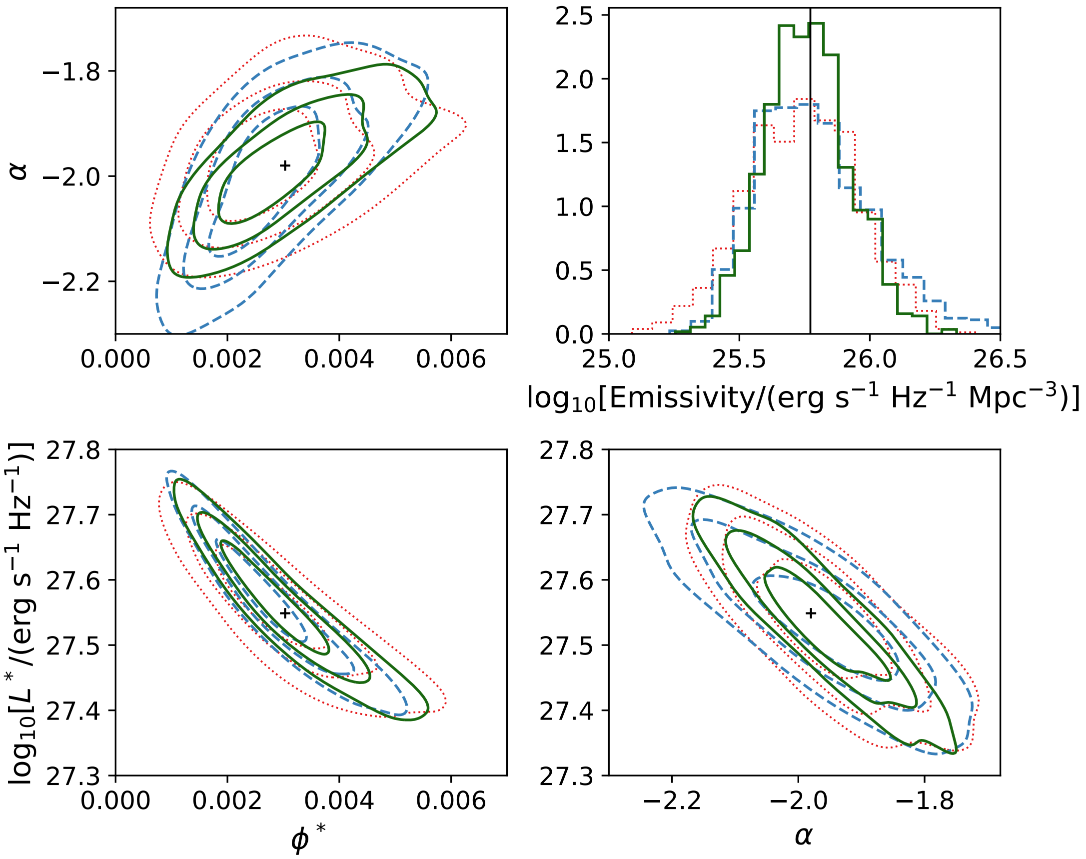

Figure 2.12: The same as Figure 2.10 but at $z = 12$. The difference between the *Full*, *Standard*, and *Naive* models is less pronounced at $z = 12$, though the *Full* method still performs best, and with the least amount of bias.

alone provides good constraints on ϕ^* and L^*, but the faint end slope α is constrained better by the joint fit than either survey alone.

At $z = 12$, shown in Figure 2.14, the combination of these two surveys is even more crucial, as neither survey can provide good constraints on any parameter by itself.

We also investigated the effects of splitting up the UD survey into four independent pointings and re-running the *Full* and *Standard* methods. This method gives a significantly better determination of the average number density of very faint sources. However, it only results in a slightly better determination of the average UVLF parameters, because the faint-end slope α is not very sensitive to cosmic variance, and the SN survey dominates the constraints of ϕ^* and L^*.

We see that tiered surveys, including both wide and deep strategies, will be essential for providing an accurate census of the high-z galaxy population.

2.6.4 Time allocation and survey design strategies

One important use for our results is to identify survey design strategies that result in the best constraints on the average UVLF parameters (Figs. 2.10 and 2.12). One could use simulations of our model to optimize the design given constraints on telescope time, survey depth, and area, but here we provide a strategy for a good initial guess.

Given a single magnitude bin and an error requirement $\epsilon_{\rm req}$ in measuring the average number counts in that bin, there is a minimum survey area below which cosmic variance will exceed the error requirement. For example, say we wish to design a survey that measures the average UVLF at $z = 9$ (and $\Delta z = 1$) at apparent magnitudes of 30 and 26 with contributions from cosmic variance at those magnitudes below 15% and 10%, respectively. Reading off Figure 2.4 or using our python package `galcv`, we find that these would require $\sim$300 arcmin^2 and $\sim$2.8 deg^2, respectively.

Alternatively, if we were willing to split each survey into independent pointings, we could satisfy the same error requirements with four $\sim$3 arcmin^2 surveys down to $m_{\rm AB} = 30$ and nine $\sim$0.11 deg^2 surveys down to $m_{\rm AB} = 26$. This observing plan requires $\sim$25× less telescope

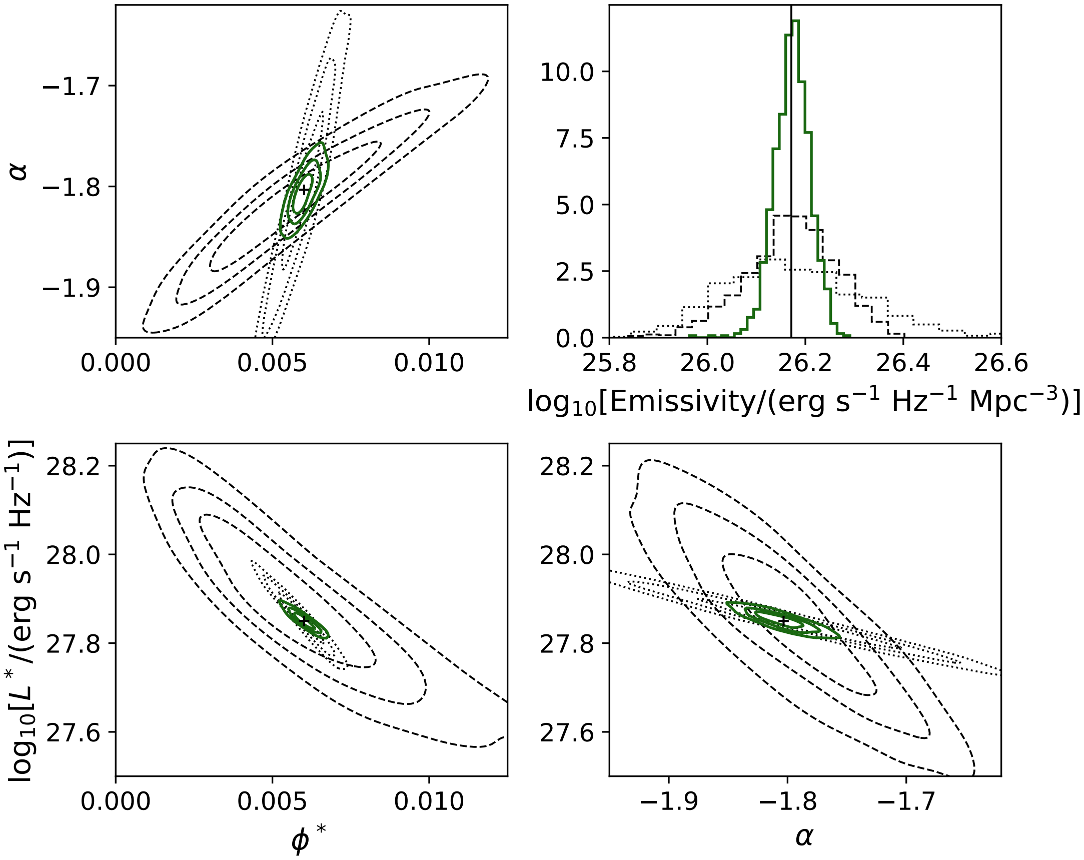

Figure 2.13: The range of parameters obtained with the SN survey alone (*dotted contours*), the UD survey alone (*dashed contours*), and with both simultaneously fit (*green solid contours*, same as those in Figure 2.10). The upper right panel shows the distribution of emissivities calculated from these distributions.

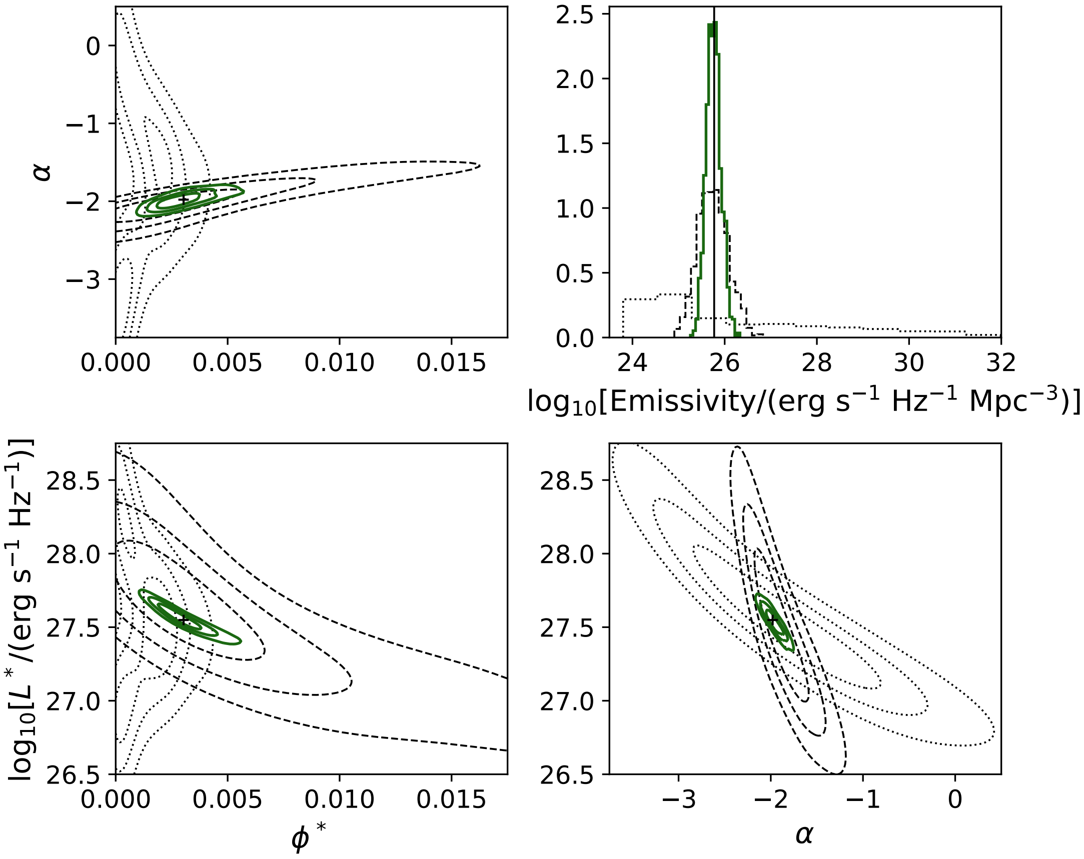

Figure 2.14: Same as Figure 2.13 but at $z = 12$. Each individual survey has much lower constraining power alone.

time for the deep survey and ∼3× less telescope time for the wide-field survey. Splitting up surveys is an especially efficient way to mitigate cosmic variance for narrow surveys because the curves in Fig. 2.4 are flattest at small survey area, so there is little penalty for moving to even narrower independent pointings. However, we do note that our model does not include nonlinear clustering that may become more important in such narrow survey fields.

Unfortunately, splitting up a survey into smaller and smaller sub-pointings is not without its drawbacks. Large mosaics can be used to measure clustering of galaxies; splitting up a survey leaves many spatial scales inaccessible, and clustering is typically more difficult to measure in the radial direction. Also, multiple small fields of view can miss interesting large-scale structures such as proto-clusters. Splitting surveys also increases observing overhead and survey design complexity. An efficient compromise would be a tiered approach: the majority of a survey's area is in one contiguous location, while a smaller fraction is split into a few independent pointings to calibrate for cosmic variance.

2.7 Conclusions

Cosmic variance will be an unavoidable source of error for next generation telescopes when measuring average properties of the Universe, especially at higher redshifts. Cosmic variance will dominate over Poisson noise for all but the brightest sources. This study integrates cosmic variance into the galaxy model developed in Furlanetto et al. (2017). We first consider how star formation rates vary with environment in the model. Next, we construct a conditional UVLF and provide its linear approximation for a wide variety of survey parameters with the parameter ε_{cv} via equation (2.20). We then study what parts of our model are most important in determining ε_{cv}. Finally, we propose a method for using these estimates as a prior on cosmic variance to improve fitting luminosity functions to high-z data.

In our model, the choice of star formation and feedback prescriptions has little effect on the relative strength of cosmic variance, and haloes of fixed mass are similar in all environments. Therefore, the main driver of cosmic variance in the UVLF is cosmic variance in the underlying dark matter halo population. The halo mass function is also the main driver in the uncertainty in our model; a more accurate conditional mass function would allow for a better prediction of cosmic variance.

A simple dark matter halo bias function along with an average halo mass to UV luminosity relation can adequately describe the relative effects of cosmic variance in the UVLF, except for density excursions exceeding $\sim 2\sigma$. In those regions, cosmic variance becomes non-Gaussian, and a full treatment is required.

We provide linear approximations of cosmic variance via ε_{cv} in terms of apparent (rest-UV) AB magnitude, survey area, and redshift. This approximation may be easily applied to any average UVLF via equation (2.20). We provide a public python package `galcv` for easy access to our results. This package provides values of ε_{cv} over a wide range of redshifts, magnitudes, survey areas, and redshift bin widths. It also includes two options for the conditional mass function used, which can be used as an estimate of the model uncertainty in the value of ε_{cv} (see Data Availability section for more details). We compare our results with cosmic variance predictions from simulations (Bhowmick et al., 2020; Ucci et al., 2021) and find good general agreement except at the smallest survey volumes (where Poisson noise begins to dominate and non-linear halo bias could be significant), or at volumes that are sizable fractions of their simulations' box size.

We also present a method for using our model as a prior on cosmic variance when fitting a UVLF to galaxy survey data. This method can inform our understanding of cosmic variance while also improving the quality of and reducing the bias in fitting the UVLF. It allows us to quantify the gains from splitting surveys into independent pointings and combining independent observations. In particular, we have shown that the combination of a shallow wide survey and a deep narrow survey is essential for fully constraining the UVLF. We also show that splitting up a survey can be an effective way to reduce the effects of cosmic variance.

Our model treats galaxy formation in a very simple manner. The primary simplification is in modelling only the *average* galaxy population in a given environment. We also ignore the effects of dust, mergers, scatter in the halo mass to UV luminosity relation, the evolution of the IMF, and the spatial distribution of star formation within a dark matter halo. Fortunately, these shortcomings pertain to (1) the details of star formation, which we have shown hardly affect the *relative* cosmic variance results ε_{cv}; and (2) individual galaxies, which are likely averaged out (to an extent) when considering cosmic variance in an ensemble of galaxies.

An understanding of cosmic variance is essential for quantifying the uncertainty in future surveys with observatories like JWST and the Roman Space Telescope. We hope that our flexible model, and the method we have introduced to incorporate cosmic variance explicitly into fitting multiple fields, can offer better constraints not just on the galaxy luminosity function but also on cosmic variance itself.

2.8 Appendix: The Eulerian Volume Correction

In this appendix we construct $p(\delta_b|R_{\rm e}, z)d\delta_b$, the fraction of volume in the Universe with linear density between (δ_b,$\delta_b + d\delta_b$) when averaged over the Eulerian scale $R_{\rm e}$. This distribution is in contrast to $p(\delta_b|R, z)d\delta_b$, the fraction of *mass* in the Universe with linear density between (δ_b,$\delta_b + d\delta_b$) when averaging over the *Lagrangian* scale R.

As described in section 2.3.2, $p(\delta_b|R, z)$ is by definition equal to a zero-mean Gaussian with variance $\sigma^2(M, z)$, where M is the mass of a region of radius R and average density. Unfortunately, $p(\delta_b|R, z)$ considers a fixed mass scale R[15], which corresponds to a density-dependent range of different volumes.

Let us choose one fixed Eulerian scale $R_{\rm e}$. We consider that scale's corresponding Lagrangian radii

$$R^3 = R_{\rm e}^3(1 + \delta_r), \tag{2.23}$$

where δ_r is the true, nonlinear density of the region. Following Mo & White (1996), the *real* density of a region may be related to the linear density via the following approximation

---

[15] R is a mass (Lagrangian) scale because it is defined as the radius of a region of mass M *if* that region were at average density. In reality, regions of mass M can have different physical volumes depending on their densities, as we will show in this Appendix.

(assuming spherical collapse):

$$\begin{aligned}\delta_b =& -1.35(1+\delta_r)^{-2/3} + 0.78785(1+\delta_r)^{-0.58661} \\ & -1.12431(1+\delta_r)^{-1/2} + 1.68647.\end{aligned} \tag{2.24}$$

Inserting this value into equation (2.23), we now have $R(\delta_b|R_\mathrm{e})$, a relation between linear density and Lagrangian radius at fixed Eulerian radius. We convert R to σ via $\sigma(M = 4\pi/3\ \bar{\rho}R^3)$ and convert δ_b to δ_0 via the growth function $\delta_0 = \delta_b/F_g(z)$. That process provides $\sigma^2(\delta_0|R_\mathrm{e}, z)$, a locus in ($\sigma^2$, δ_0) space of constant Eulerian radius R_e.

With $\sigma^2(\delta_0|R_\mathrm{e}, z)$, we can use the excursion set formalism to solve for $f_R(\sigma^2|R_\mathrm{e}, z)$, the distribution of mass in the Universe that is associated with a region with σ^2 (and thus corresponding R and δ_0) at fixed R_e. The excursion set formalism describes a random walk in dark matter density δ_0 as one averages over first a very large volume (small σ), and then successively smaller volumes (larger σ) centered at a single point in space (Bond et al., 1991; Lacey & Cole, 1993). The distribution of random walks that first cross the barrier $\sigma^2(\delta_0|R_\mathrm{e}, z)$ defines $f_R(\sigma^2|R_\mathrm{e}, z)$.

$f_R(\sigma^2|R_\mathrm{e}, z)$ has no analytic solution for an arbitrary barrier shape, so we approximate $\sigma^2(\delta_0|R_\mathrm{e}, z)$ as a straight line,[16]

$$B(\sigma^2|R_\mathrm{e}, z) = B_0 + B_1\sigma^2, \tag{2.25}$$

where B is the density δ_0, and B_0 and B_1 are fit parameters (corresponding to the y-intercept and slope, respectively).

Fortunately, the first-crossing distribution f_R for a linear barrier in (σ^2, δ_0) space has been solved analytically by Sheth (1998):

$$\begin{aligned}& f_R(\sigma^2, |R_\mathrm{e}, z)\mathrm{d}\sigma^2 = \\ & \frac{B(0|R_\mathrm{e}, z)}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{B^2(\sigma^2|R_\mathrm{e}, z)}{2\sigma^2}\right)\frac{\mathrm{d}\sigma^2}{\sigma^2}.\end{aligned} \tag{2.26}$$

This is an Inverse Gaussian distribution.

---

[16] We approximate $\sigma^2(\delta_0|R_\mathrm{e}, z)$ as a line by fitting it to the barrier near where most trajectories cross the barrier: $\delta_0 = 0$.

We convert $f_R(\sigma^2|R_\mathrm{e}, z)$, a mass fraction distribution in σ^2, to $p(\delta_b|R_\mathrm{e}, z)$, a volume fraction distribution in δ_b, following equation (16) of Sheth (1998):

$$p(\delta_b|R_\mathrm{e}, z)\mathrm{d}\delta_b = \frac{1}{(1+\delta_r)} f_R(\sigma^2|R_\mathrm{e}, z)\mathrm{d}\sigma^2. \tag{2.27}$$

In principle, dividing by the non-linear function $(1+\delta_r)$ can result in a $p(\delta_b|R_\mathrm{e}, z)$ that is not normalized. In practice, $p(\delta_b|R_\mathrm{e}, z)$ remains normalized within 1% for all cases we consider.

For the range of redshifts and scales considered in this paper, $p(\delta_b|R_\mathrm{e}, z)$ is near to a Gaussian with standard deviation $\sigma(M = 4\pi/3\ \bar{\rho}R_\mathrm{e}^3)$. However, the distribution is skewed towards negative densities, resulting in a boost in the negative wing and suppression in the positive wing, an effect that is most significant for volumes with radii less than ~10 Mpc (see Fig. 2.15).

At $z = 9$, we find that for regions with scales $R_\mathrm{e} = 5$, 10, and 50 Mpc, the fraction of volume in the Universe that is below average cosmic density is 56, 54, and 51%, respectively. These fractions increase slightly at lower redshifts as underdense regions continue to expand relative to overdense regions. This result indicates that surveys will be slightly more likely to probe underdense regions. Using different methods, Muñoz et al. (2010) also found that surveys are more likely to probe an underdense region because of those regions' more rapid cosmic expansion.

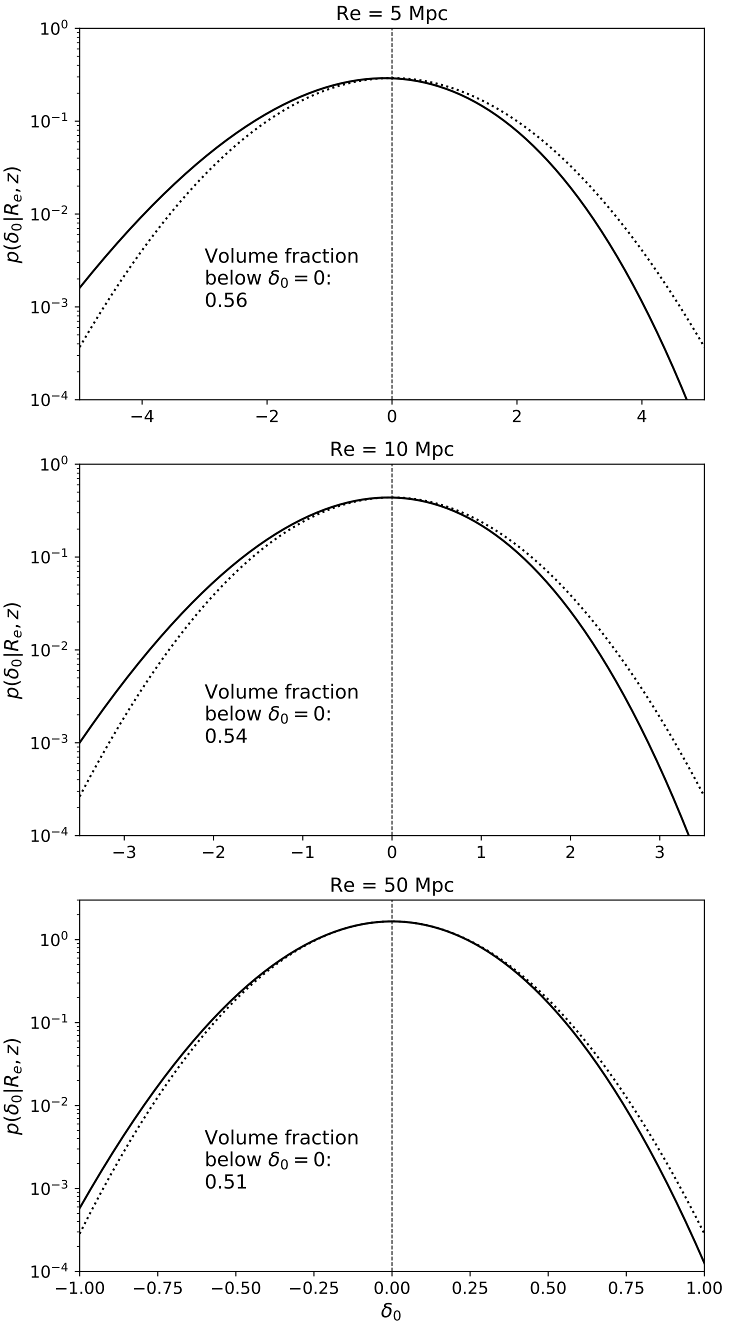

Figure 2.15: Effects of the Eulerian correction at $z = 9$. The *solid* lines show $p(\delta_0|R_\mathrm{e}, z)$, the distribution of linear densities at fixed scale (densities are extrapolated to $z = 0$ via the growth function). The *dotted* lines show the Lagrangian distribution of densities $p(\delta_0|R, z)$: Gaussian distributions with standard deviation $\sigma(R_\mathrm{e}, z)$. The volume of the Universe that is below average density at each scale is indicated in each panels.

CHAPTER 3

A Framework for Simultaneously Measuring Field Densities and the High-z Luminosity Function

This chapter can be seen in its published form (Trapp et al., 2022) here:
https://ui.adsabs.harvard.edu/abs/2022MNRAS.510.4844T/abstract

3.1 Abstract

Cosmic variance from large-scale structure will be a major source of uncertainty for galaxy surveys at $z \gtrsim 6$, but that same structure will also provide an opportunity to identify and study dense environments in the early Universe. Using a robust model for galaxy clustering, we directly incorporate large-scale densities into an inference framework that simultaneously measures the high-z ($z \gtrsim 6$) UV luminosity function and the average matter density of each distinct volume in a survey. Through this framework, we forecast the performance of several major upcoming James Webb Space Telescope (JWST) galaxy surveys. We find that they can constrain field matter densities down to the theoretical limit imposed by Poisson noise and unambiguously identify over-dense (and under-dense) regions on transverse scales of tens of comoving Mpc. We also predict JWST will measure the luminosity function with a precision at $z = 12$ comparable to existing Hubble Space Telescope's constraints at $z = 8$ (and even better for the faint-end slope). We also find that wide-field surveys are especially important in distinguishing luminosity function models.

3.2 Introduction

The first billion years of galaxy formation is about to be explored as never before. The James Webb Space Telescope (JWST), the Nancy Grace Roman Space Telescope, the Thirty Meter Telescope, the European Extremely Large Telescope, the Giant Magellan telescope, and many other next-generation observatories will open a new frontier at the beginning of structure formation in our Universe.

Perhaps the most basic observable of this era is the galaxy luminosity function, which tracks the growth of the galaxy population as a whole. The evolution of both its shape and normalization have important implications for galaxy formation scenarios, so it has been intensely studied by existing facilities (see e.g., Bouwens et al., 2015; Finkelstein et al., 2015; Bowler et al., 2015; Livermore et al., 2017; Atek et al., 2018; Oesch et al., 2018; Behroozi et al., 2019; Bouwens et al., 2021; Finkelstein et al., 2022). To date, these measurements have pinned down the abundance of relatively bright galaxies at $z \lesssim 8$ to a reasonable precision. The results are largely consistent with simple extrapolations of galaxy physics at lower redshifts (see e.g., Tacchella et al., 2013; Mason et al., 2015; Furlanetto et al., 2017; Mirocha et al., 2017). A few bright galaxies have also been discovered at $z \gtrsim 9$, but they are currently too rare for robust estimates of their abundance (Oesch et al., 2013, 2015; Bouwens et al., 2015; Ishigaki et al., 2015; McLeod et al., 2015, 2016; Bouwens et al., 2019).

Despite the substantial progress in understanding these galaxies over the last several years, the field is poised for a revolution with the launch of JWST and, beyond that, the Roman Telescope. The extraordinary sensitivity of these facilities will allow galaxy searches to extend both to significantly fainter sources and to higher redshifts (e.g., Behroozi & Silk 2015; Furlanetto et al. 2017; Kauffmann et al. 2020). It will be crucial to optimize these efforts in order to constrain the luminosity function.

Going to higher redshifts comes with bigger challenges but also more opportunities. Among the obstacles is the increase in uncertainty due to cosmic variance:[1] large-scale

---

[1] In this paper, we follow common usage and apply the term "cosmic variance" to describe the fluctuations in dark matter density between different volumes in our Universe and the consequences of that variance for

dark matter densities affect the formation of dark matter haloes and thus change the normalization and shape of the luminosity function for different volumes (see Figure 3.1, and Trapp & Furlanetto, 2020). But this cosmic variance is not simply a nuisance, because it reflects the large-scale structure that is itself a key driver of both galaxy formation and reionization during the Cosmic Dawn. As such, different large-scale densities (if they can be distinguished!) can serve as laboratories for discovery:

1. Because ionizing sources trace the large-scale density field, the reionization process is heavily dependent on those densities (Furlanetto et al., 2004). In order to understand that era, there is great interest in identifying over-dense regions that may host large ionized bubbles or other unusual ionization environments (see e.g. Zitrin et al., 2015; Jung et al., 2020; Tilvi et al., 2020; Hu et al., 2021; Endsley et al., 2021) as well as large-scale underdensities that may be the last regions to be reionized (Becker et al., 2018; Davies et al., 2018; Christenson et al., 2021).
2. The large-scale environment of galaxies appears to play a role in their evolution at later times (e.g., assembly bias; Gao & White, 2007), and it may be even more important at $z \gtrsim 6$. This is because feedback from large-scale radiation backgrounds can have enormous effects on star formation at these times. Most importantly, photoheating from reionization will suppress the formation of small galaxies inside ionized regions (Thoul & Weinberg, 1996; Iliev et al., 2007; Noh & McQuinn, 2014). Measurements of the large-scale density will allow detailed exploration of these mechanisms.
3. A region's density provides information on its past and future – whether it will become a galaxy cluster, its reionization history, etc. Understanding the assembly history of unusual objects like galaxy clusters is therefore facilitated by measuring large-scale densities at early times (e.g., Chiang et al. 2017).
4. Finally, while cosmologists understand the underlying dark matter density field fairly well, the transformation from those densities to galaxy observables is more uncertain.

---

the galaxy population. More precisely, this is a particular case of sample variance, with cosmic variance sometimes reserved for the errors intrinsic to having only one Universe to observe.

Comparing large-scale density estimates from surveys with the cosmological predictions can help identify any problems in theoretical models of galaxy formation (Trapp & Furlanetto, 2020).

In order to deal with the uncertainties of and gain insights from cosmic variance, we need a comprehensive way of modelling its effects on the galaxy population, starting with the luminosity function and its measurement. In Trapp & Furlanetto (2020), we developed a model of cosmic variance's effects on the luminosity function, including important non-linear effects and a correction for the elongated 'pencil-beam' shape of many survey volumes. Recent simulations have also estimated the cosmic variance of galaxies (Bhowmick et al., 2020; Ucci et al., 2021), and those results are comparable to our analytical estimates. However, these estimates have not yet been fully integrated into luminosity function fitting methods.

A 'standard method' to fit the luminosity function assumes its shape is constant across all fields, while ignoring the normalization of each one. At the end of the fitting process, a normalization is chosen such that the correct total number of sources is recovered (see e.g., Finkelstein et al., 2015; Bouwens et al., 2015). This method ignores the dependence of the shape of the luminosity function on the density of a region at higher redshifts (see Figure 3.1), and it discards potentially useful information about field densities. As a result, the standard method is effective in mitigating cosmic variance, but it leads to a sub-optimal (and slightly biased) fit (Trapp & Furlanetto, 2020). Here, we develop a method that integrates priors on cosmic variance into luminosity function inference, improving constraints and providing new information on the large-scale densities.

In section 3.3, we use Bayesian statistics to incorporate cosmic variance into a luminosity function fitting framework from the ground up. This framework simultaneously fits luminosity function parameters and individual field densities to data from multiple fields. We also develop a method to combine multiple fields into a single 'effective' field to cut down on computation time. We use our framework to forecast the performance of a selection of upcoming galaxy surveys in section 3.4. We explore how well different combinations of the surveys can measure the luminosity function and individual field densities, finding a fundamental limit to measuring a field's density that is due to Poisson noise. We present our conclusions in section 3.5.

We take the following cosmological parameters: $\Omega_m = 0.308$, $\Omega_\Lambda = 0.692$, $\Omega_b = 0.0484$, $h = 0.678$, $\sigma_8 = 0.815$, and $n_s = 0.968$, consistent with recent Planck Collaboration XIII results (Planck Collaboration et al., 2016). We give all distances in comoving units. All luminosities are rest-frame ultra-violet ($1500 - 2800$ Å)[2] luminosities, and all magnitudes are AB magnitudes.

3.3 Methods

In this section, we (i) define the luminosity function, (ii) create a posterior that includes the luminosity function parameters as well as the densities of each field, (iii) develop a novel way of combining the cosmic variance functions of separate fields, and (iv) discuss the challenge of dealing with systematic offsets between datasets.

3.3.1 The luminosity function of galaxies

Let us assume that the average number of galaxies with absolute magnitudes between $(M_{\rm AB}, M_{\rm AB} + dM_{\rm AB})$ is described by $\Phi_{\rm avg}(M_{\rm AB}, z)dM_{\rm AB}$, which has parameters $\vec{\phi}(z)$ that are only functions of redshift. We organize the parameters of the luminosity function into a vector for notational convenience; for a typical Schechter function, these would be the normalization ϕ^*, the characteristic magnitude M^*, and the faint-end power-law slope, α:

$$\begin{aligned}&\Phi_{\rm avg}(M_{\rm AB}, z)dM_{\rm AB} = \\ &(0.4\ \ \ln 10)\phi^*[10^{0.4(M^*_{\rm AB}-M_{\rm AB})}]^{\alpha+1}\exp[-10^{0.4(M^*_{\rm AB}-M_{\rm AB})}]dM_{\rm AB}.\end{aligned} \tag{3.1}$$

In a given volume V, the dark matter density may differ from the average, which then makes the galaxy luminosity function dependent on that large-scale density (Lacey & Cole, 1993). This effect is called cosmic variance or sample variance. For large volumes, a linear approach is sufficient to describe the conditional luminosity function (Mo & White, 1996;

---

[2] This wavelength range corresponds to H-band in the redshift range of $z \approx 5$–9 and K-band for $z \approx 8$–12.

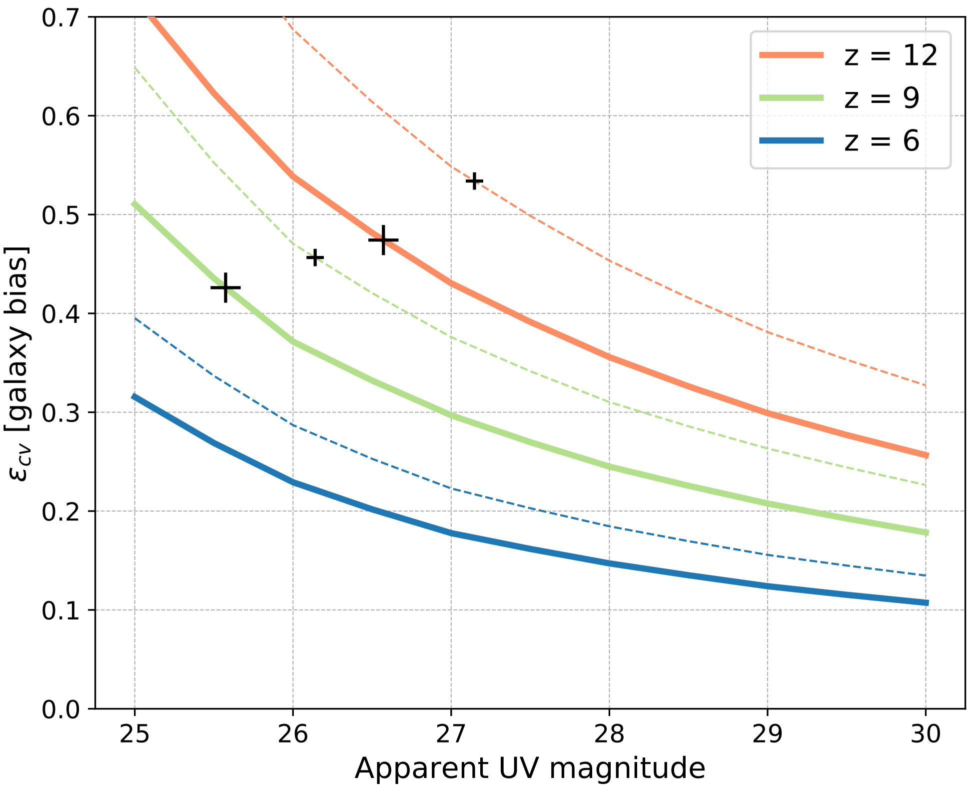

Figure 3.1: Cosmic variance as a function of apparent magnitude for redshifts 6, 9, 12 (bottom, middle, top *thick solid* lines). A cosmic variance value of e.g., $\varepsilon_{\rm cv} = 0.3$ means a region with a 1-σ over-density of dark matter will have 30% more galaxies than the Universe average for such a volume (not accounting for Poisson noise). The *thick solid* lines are for a 100 arcmin^2 survey with $\Delta z = 1$, and the *thin dashed* lines are for a 25 arcmin^2 survey. Cosmic variance becomes a stronger function of magnitude at higher redshift and smaller volume, meaning cosmic variance can significantly affect the shape of the luminosity function as well as the normalization. The '+' markers indicate where we would expect to find only $\sim$1 source at the indicated magnitude in such a survey, an indicator of where Poisson noise is clearly dominating.

Trapp & Furlanetto, 2020):

$$\Phi_{\mathrm{loc}}(M_{\mathrm{AB}}, z, \tilde{\delta}) = \Phi_{\mathrm{avg}}(M_{\mathrm{AB}}, z)(1 + \tilde{\delta}\varepsilon_{\mathrm{cv}}(M_{\mathrm{AB}}, z)), \tag{3.2}$$

where $\tilde{\delta}$ is a normally distributed random variable signifying the *normalized* relative dark matter density in the volume V. The actual relative density $\delta_b = (\rho - \bar{\rho})/\bar{\rho}$ can be found with $\delta_b = \tilde{\delta}\sigma_{\mathrm{PB}}$, where σ_{PB} is the rms fluctuation of δ_b across many volumes of shape V at redshift z. Finally, $\varepsilon_{\mathrm{cv}}(M_{\mathrm{AB}}, V, z)$ is the cosmic variance from Trapp & Furlanetto (2020).

The cosmic variance function $\varepsilon_{\mathrm{cv}}(M_{\mathrm{AB}}, V, z)$ is constructed by implementing non-linear halo clustering theory into a self-consistent analytical galaxy model. The function also includes a correction for the elongated 'pencil-beam' shape of many survey volumes and for non-linear expansion/contraction of a region. The largest uncertainty in $\varepsilon_{\mathrm{cv}}$ comes from non-linear halo clustering theory; $\varepsilon_{\mathrm{cv}}$ can vary $\sim$25% between different models. The choice of galaxy model can also affect $\varepsilon_{\mathrm{cv}}$, but to a lesser extent ($\varepsilon_{\mathrm{cv}}$ is a relatively weak function of magnitude, although it gets stronger at higher redshifts, see Figure 3.1). See Trapp & Furlanetto (2020) for more information on $\varepsilon_{\mathrm{cv}}$ and Bhowmick et al. (2020) and Ucci et al. (2021) for recent cosmic variance estimations calculated from large volume galaxy simulations.

The luminosity function that is actually measured depends on observational features like the completeness and contamination functions, which we combine into a single function $f(M_{\mathrm{AB}}, z)$ that is unique to each survey volume. The luminosity function becomes:

$$\Phi_{\mathrm{obs}}(M_{\mathrm{AB}}, z, \tilde{\delta}) = f(M_{\mathrm{AB}}, z) \times \Phi_{\mathrm{loc}}(M_{\mathrm{AB}}, z, \tilde{\delta}). \tag{3.3}$$

In the next section, we use a Bayesian framework to fit this observed luminosity function to data.

3.3.2 The posterior

Given data $\vec{D}$ from a large galaxy survey composed of N_{f} fields each with their own volume, observational effects $f(M_{\mathrm{AB}}, z)$, and local density $\tilde{\delta}$, we would like to learn the parameters $\vec{\phi}(z)$ of the average luminosity function. The data $\vec{D}$ is comprised of many galaxies with measured redshifts and magnitudes. Let us assume a sufficiently narrow redshift range that we may ignore any redshift dependence within the measured volume (see Appendix 3.6).

We would like to determine the posterior $p(\vec{\phi}|\vec{D})$: the probability density of the luminosity function parameters given the data. We are also interested in $p(\vec{\delta}|\vec{D})$: the probability density of the normalized dark matter densities[3] from the N_f fields given the data. A third set of parameters arise from the halo-galaxy mapping performed by a galaxy evolution model (as explored in Trapp & Furlanetto 2020). In this paper, we focus on understanding the large-scale density field so do not include such parameters, but we do estimate their impact on the constraints in section 3.4.6.

We then start then with the joint posterior $p(\vec{\phi}, \vec{\delta}|\vec{D})$. Using Bayes' theorem:

$$p(\vec{\phi}, \vec{\delta}|\vec{D}) \propto p(\vec{D}|\vec{\phi}, \vec{\delta}) \times p(\vec{\phi}) \times p(\vec{\delta}) \tag{3.4}$$

where $p(\vec{D}|\vec{\phi}, \vec{\delta})$ is the likelihood of the data given the parameters of the average luminosity function and the dark matter densities, and $p(\vec{\phi})$ and $p(\vec{\delta})$ are their respective independent priors. We assume flat priors for the luminosity function parameters in this work. When analysing real data, one should use a prior based on a previous and independent measurement of the luminosity function, which would result in a tighter final posterior. We assume the survey fields are separated on the sky such that their environments are independent of one another. Under that assumption, the density prior $p(\vec{\delta})$ is simply a multivariate Gaussian centered at the origin with an identity covariance matrix of size $N_\mathrm{f} \times N_\mathrm{f}$. We explore in section 3.3.3.2 how to deal with adjacent survey fields with correlated densities.

We write the total likelihood as a product of each independent field's likelihood:

$$\mathcal{L} = p(\vec{D}|\vec{\phi}, \vec{\delta}) = \prod_i^{N_\mathrm{f}} p(D_i|\vec{\phi}, \tilde{\delta}_i), \tag{3.5}$$

with $p(D_i|\vec{\phi}, \tilde{\delta}_i)$ as the likelihood that one has obtained data D_i from a field with density $\tilde{\delta}_i$

---

[3] $\vec{\delta}$ is a vector containing $\tilde{\delta}$ for each survey. We omit the tilde when writing the vector for convenience.

and global parameters $\vec{\phi}$, which can be described by

$$\begin{aligned} p(D_i|\vec{\phi},\tilde{\delta}_i) &= \quad (3.6) \\ & P(n_i \text{ galaxies are found in volume } V_i|\vec{\phi},\tilde{\delta}_i) \quad (3.7) \\ & \times p(\text{distribution of } M_{\rm AB}|\vec{\phi},\tilde{\delta}_i). \quad (3.8) \end{aligned}$$

The probability that n_i galaxies are found in a volume follows the Poisson distribution with expected value $n_{i,\rm exp}$,

$$P(n_i \text{ galaxies are found in volume } V_i|\vec{\phi},\tilde{\delta}_i) = \frac{n_{i,\rm exp}^{n_i}}{n_i!}e^{-n_{i,\rm exp}} \quad (3.9)$$

where

$$n_{i,\rm exp} = V_i \times \int_{-\infty}^{M_{\rm lim}} f_i(M')\Phi_{\rm avg}(M',\vec{\phi})(1+\tilde{\delta}_i\varepsilon_{{\rm cv},i}(M',V_i)){\rm d}M'. \quad (3.10)$$

The probability to find a particular distribution of magnitudes is found by multiplying the relative probabilities that each individual source is found at that specific magnitude,

$$\begin{aligned} & p(\text{distribution of } M_{\rm AB}|\vec{\phi},\tilde{\delta}_i) = \\ & \prod_j^{n_i} \frac{f_i(M_j)\Phi_{\rm avg}(M_j,\vec{\phi})(1+\tilde{\delta}_i\varepsilon_{{\rm cv},i}(M_j,V_i))}{\int_{-\infty}^{M_{\rm lim}} f_i(M')\Phi_{\rm avg}(M',\vec{\phi})(1+\tilde{\delta}_i\varepsilon_{{\rm cv},i}(M',V_i)){\rm d}M'}, \end{aligned} \quad (3.11)$$

where $M_{\rm lim}$ is the magnitude limit of the survey. Substituting in $n_{i,\rm exp}/V_i$ for the integral and plugging into equation (3.6), we find

$$\begin{aligned} & p(D_i|\vec{\phi},\tilde{\delta}_i) = \\ & \frac{n_{i,\rm exp}^{n_i}}{n_i!}e^{-n_{i,\rm exp}} \times \\ & \left(\frac{n_{i,\rm exp}}{V_i}\right)^{-n_i} \prod_j^{n_i} f_i(M_j)\Phi_{\rm avg}(M_j,\vec{\phi})(1+\tilde{\delta}_i\varepsilon_{{\rm cv},i}(M_j,V_i)). \end{aligned} \quad (3.12)$$

Inserting this into the full likelihood, taking the natural logarithm, and dropping terms that don't depend on $\vec{\phi}$ or $\vec{\delta}$ gives

$$\begin{aligned} & \ln\mathcal{L} \propto \\ & \sum_i^{N_{\rm f}} \left[-n_{i,\rm exp} + \sum_j^{n_i} \Big(\ln\Phi_{\rm avg}(M_j,\vec{\phi}) + \ln(1+\tilde{\delta}_i\varepsilon_{{\rm cv},i}(M_j,V_i))\Big)\right]. \end{aligned} \quad (3.13)$$

We can then write the posterior as

$$p(\vec{\phi}, \vec{\delta}|\vec{D}) \propto \mathcal{L} \times p(\vec{\delta}) \times p(\vec{\phi}). \tag{3.14}$$

Finally, we can marginalize over $\vec{\phi}$ or $\vec{\delta}$ to get $p(\vec{\delta}|\vec{D})$ or $p(\vec{\phi}|\vec{D})$, respectively.

A downside of considering the density of each field is the expanded dimensionality of the parameter space. For example, given a luminosity function form that has m parameters and a survey that has N_{f} sub-fields, the posterior has a dimensionality of $m \times N_{\mathrm{f}}$. However, each sub-field's $\tilde{\delta}_i$ parameter is assumed to be independent, reducing the problem to N_{f} different likelihoods each with dimensionality $m + 1$, a drastic reduction. These likelihoods can then quickly be combined to create the full posterior.

3.3.3 Combining fields

Unfortunately, sampling the posterior with many sub-fields can still be costly. For example, imagine a parallel survey with dozens of independent pointings, each of which provides a separate density and local luminosity function. To alleviate this limitation, we next describe a method of combining multiple fields into a single 'effective' field with $\tilde{\delta}_{\mathrm{eff}}$ and $\varepsilon_{\mathrm{cv,\,eff}}$. Combining fields is a trade-off. We reduce the dimensionality of the posterior, but we also lose the ability to measure individual field densities. This trade-off is well worth it when we only care about the density of a select number of fields in a survey, or in the following two cases: (i) the fields may be contiguous but at different depths (so that it might be useful to combine them in order to improve the environment measurement), or (ii) they may be widely separated and independent (but presumably shallow enough that no individual field will provide a robust environment anyway). We begin with the latter case.

3.3.3.1 Independent fields

We begin with a method that is especially useful in combining fields from a large parallel program, or from a handful of mosaics. Take N_{f} fields all at the same redshift, each with a different effective volume curve $V_{\mathrm{eff}}(M_{\mathrm{AB}})$ (effective volume is the combined completeness/-

contamination function $f(M_{\rm AB})$ times total volume V). The fields are far apart from one another on the sky so are in independent environments.

Choosing a single field, the number of sources in a small magnitude bin is

$$dN = V_{\rm eff}(M_{\rm AB})\Phi_{\rm avg}(M_{\rm AB})(1+\tilde{\delta}\varepsilon_{\rm cv}(M_{\rm AB},V))dM_{\rm AB}. \tag{3.15}$$

At fixed $M_{\rm AB}$, dN is essentially a 'measurable' with average value $dN_{\rm avg} = V_{\rm eff}\Phi_{\rm avg}dM_{\rm AB}$ and error in the measurement $\sigma_{\rm err} = V_{\rm eff}\Phi_{\rm avg}\varepsilon_{\rm cv}dM_{\rm AB}$. Combining the $N_{\rm f}$ different fields, we would have $dN_{\rm tot} = \sum_i^{N_s} dN_i$ with the average value as $dN_{{\rm tot},avg} = dM_{\rm AB}V_{\rm eff,\ tot}\Phi_{\rm avg}$ and $V_{\rm eff,\ tot}(M_{\rm AB}) = \sum_i^{N_s} V_{{\rm eff},i}(M_{\rm AB})$. Via standard propagation of errors,

$$\sigma_{\rm err,\ tot} = dM_{\rm AB}V_{\rm eff,\ tot}\Phi_{\rm avg}\frac{\sqrt{(V_{\rm eff,1}\varepsilon_{\rm cv,1})^2+(V_{\rm eff,2}\varepsilon_{\rm cv,2})^2+...}}{V_{\rm eff,\ tot}}. \tag{3.16}$$

We can then write the total expected number number of galaxies in this bin as

$$dN_{\rm tot} = V_{\rm eff,\ tot}(M_{\rm AB})\Phi_{\rm avg}(M_{\rm AB})(1+\tilde{\delta}_{\rm eff}\cdot\varepsilon_{\rm cv,eff}(M_{\rm AB}))dM_{\rm AB}, \tag{3.17}$$

with

$$\varepsilon_{\rm cv,\ eff}(M_{\rm AB}) = \frac{\sqrt{\sum_i^{N_s}[V_{{\rm eff},i}(M_{\rm AB})\varepsilon_{{\rm cv},i}(M_{\rm AB})]^2}}{V_{\rm eff,\ tot}(M_{\rm AB})} \tag{3.18}$$

as the effective cosmic variance function of the combined fields and $\tilde{\delta}_{\rm eff}$ a normally-distributed random variable. Not unexpectedly, the effective cosmic variance function is just the individual functions added in quadrature and weighted by volume. Now, we may combine the data from $N_{\rm f}$ independent fields and treat it as a single field using the 'effective' cosmic variance function and the total effective volume curve (which is just all of the individual curves added together).

Unfortunately, this method is technically correct only if all fields' effective volume curves $V_{\rm eff}$ (or completeness functions $f(M_{\rm AB})$) have the exact same shape (though they may have different normalization). This is because eqs. 3.15-3.18 are evaluated at a single magnitude. At some other magnitude, the surveys' fractional contribution to the total effective volume may be different, which would mean a different effective environment $\tilde{\delta}_{\rm eff}$. Fields that have very different effective volume curves should not be combined in this way. However, we find that fields that have consistent effective volume curves over the majority of their magnitude coverage can be combined this way with accurate results, even if one field goes ∼1 magnitude deeper than others.

3.3.3.2 Contiguous fields

Now we turn to the case in which volumes are contiguous with each other (as may occur with "wedding cake" surveys with embedded deep fields). Because the fields are contiguous, their densities are similar. Again, we take N_f fields, each with a different effective volume curve $V_\mathrm{eff}(M_\mathrm{AB})$. For one field, the number of sources in a small magnitude range is

$$dN = V_\mathrm{eff}(M_\mathrm{AB})\Phi_\mathrm{avg}(M_\mathrm{AB})\left[1 + \frac{\delta_b}{\sigma_\mathrm{PB}}\varepsilon_\mathrm{cv}(M_\mathrm{AB})\right] dM_\mathrm{AB} \tag{3.19}$$

where this time, we have explicitly written the density $\tilde{\delta} = \delta_b/\sigma_\mathrm{PB}$ (see section 3.3.1). Putting the fields together,

$$dN_\mathrm{tot} = dM_\mathrm{AB}\Phi_\mathrm{avg}\left(V_\mathrm{eff,\ tot} + \sum_i^{N_s} V_{\mathrm{eff},i}\frac{\delta_{b,i}}{\sigma_\mathrm{PB,i}}\varepsilon_{\mathrm{cv},i}\right). \tag{3.20}$$

We now make the assumption that $\delta_{b,i} = \delta_\mathrm{tot}$, i.e., the dark matter density is the same in all of the contiguous fields because they are near each other. This is only valid if $\sigma_\mathrm{PB}[\mathrm{within}] \ll 1$, where $\sigma^2_\mathrm{PB}[\mathrm{within}] = \sigma^2_\mathrm{PB}[\mathrm{smallest\ field}] - \sigma^2_\mathrm{PB}[\mathrm{tot}al\ survey\ volume]$ is the 1-σ fluctuation of dark matter density when zooming in from the entire contiguous survey to the smallest field. Of course, in reality the different fields will not have the same density, though they will be correlated with each other. The assumption of a uniform density is a simplification useful for forecasting results; the inferred density will then be a weighted combination of the different parts of the field. If the true densities are of interest, our method can be extended to include these correlations (see section 3.4.4).

Re-writing again, we find

$$dN_\mathrm{tot} = dM_\mathrm{AB}V_\mathrm{eff,\ tot}\Phi_\mathrm{avg}(1 + \frac{\delta_\mathrm{tot}}{\sigma_\mathrm{PB,tot}}\sum_i^{N_s}\frac{V_{\mathrm{eff},i}}{V_\mathrm{eff,\ tot}}\frac{\sigma_\mathrm{PB,tot}}{\sigma_\mathrm{PB,i}}\varepsilon_{\mathrm{cv},i}). \tag{3.21}$$

In other words,

$$dN_\mathrm{tot}(M_\mathrm{AB}) = dM_\mathrm{AB}V_\mathrm{eff,\ tot}(M_\mathrm{AB})\Phi_\mathrm{avg}(M_\mathrm{AB})[1 + \tilde{\delta}_\mathrm{eff}\varepsilon_\mathrm{cv,\ eff}(M_\mathrm{AB})], \tag{3.22}$$

with

$$\varepsilon_{\text{cv, eff}}(M_{\text{AB}}) = \sum_{i}^{N_s} \frac{V_{\text{eff},i}(M_{\text{AB}})}{V_{\text{eff, tot}}(M_{\text{AB}})} \frac{\sigma_{\text{PB,tot}}}{\sigma_{\text{PB,i}}} \varepsilon_{\text{cv},i}(M_{\text{AB}}). \tag{3.23}$$

and $\tilde{\delta}_{\text{eff}} = \delta_{\text{tot}}/\sigma_{\text{PB,tot}}$ is just a normal random variable. This time, the effective cosmic variance function is weighted by both the effective volume and the individual fields' rms dark matter variation.

Again, we may combine the data from N_{f} contiguous fields and treat them as a single field using the 'effective' cosmic variance function and the total effective volume curve. However, we must ensure our assumption of $\sigma_{\text{PB}}[\text{within}] \ll 1$ is valid. Through testing with our public cosmic variance calculator `galcv` (Trapp & Furlanetto, 2020), $\sigma_{\text{PB}}[\text{within}] < 0.2$ gives an effective cosmic variance function within 10% of the "correct" answer.

Finally, a mixture of independent and contiguous fields may be combined into an effective field by applying these methods one after another, always starting with the contiguous combinations. The value of $\varepsilon_{\text{cv, eff}}$ may also be used to forecast the aggregate effects of cosmic variance on any set of surveys before they are observed (given an estimate of their effective volumes).

To summarize, in this section we have introduced simple ways to combine several fields for a joint analysis. We emphasize that these simplifications are largely for computational convenience; given the approximations already inherent to forecasting future surveys, they are certainly useful at this stage, but they may not be useful once the data are in hand.

3.3.4 Systematic offsets between fields

Unknown systematic offsets between fields complicates the measurement of their densities. Our framework determines a field's density by comparison with other fields. If data from all fields are reduced in a consistent manner, any systematic errors in normalization would not affect this comparison. However, when applying our framework to systematically distinct datasets, a normalization offset could result in erroneous density measurements.

The parameter that holds these systematics is the effective volume V_{eff}, which unfortunately cannot be measured, only modeled. Most often it is modeled by inserting artificial sources into simulations, taking a virtual observation of those simulations, and recording the number of correctly and incorrectly recovered sources (e.g. Finkelstein et al., 2015; Bouwens et al., 2015). This method is robust and well developed, but the versions from different groups would ideally be tested on identical datasets to estimate their systematics and therefore uncertainties, which could then be accounted for in fitting.

Another source of potential systematic bias involves the galaxy selection process. In the same field, different selection criteria can identify different galaxy populations, which may have separate associated luminosity functions that should not be fit to with a single model. We do not attempt to model such differences in our forecasts below, but they will be important to understand in future data. For simplicity, we assume that surveys are "perfect" in that they can reliably identify all galaxies in the survey volume, aside from incompleteness at the faint end (modeled by $V_{\rm eff}$). In other words, we assume that the galaxy selection criteria used in the surveys do not "miss" galaxy populations due to dust, old stars, etc. Kauffmann et al. (2020) find that this assumption is a reasonable one for predicting JWST survey results.

3.4 Applications for JWST and beyond

In this section, we test our framework by simulating a set of upcoming JWST surveys and one Nancy Grace Roman Space Telescope survey in the range $6 \leq z \leq 12$. These simulations are designed to predict the performance of these surveys in regards to measuring the average luminosity function and field densities. It also serves as an example of how to apply our framework.

3.4.1 Simulating surveys

We simulate upcoming surveys in a very simple manner, and as such, our forecasts cannot be seen as authoritative. However, they should provide reasonably accurate estimates of the expected precision of future measurements, and the general trends and qualitative relationships we find are robust and can be used as a tool to better plan future surveys. Kauffmann et al. (2020) contains a more detailed forecast for a smaller set of surveys.

We simulate 6 upcoming JWST surveys:

1. the CEERS survey (∼100 arcmin^2)
2. the JADES survey (∼236 arcmin^2)
3. the PRIMER survey (∼695 arcmin^2)
4. the PANORAMIC parallel survey (∼1500 arcmin^2)
5. the WDEEP deep-field survey (∼10 arcmin^2)

We also simulate one Nancy Grace Roman Space Telescope survey: the Roman Supernova (RSN) survey (∼32400 arcmin^2). All survey features are described in Table 3.1, though note that the RSN parameters are only estimates as the survey has not been finalized. We do not consider a comprehensive list of JWST galaxy surveys, choosing instead a representative sample. Other surveys will be useful (such as COSMOS-Webb) and more will be scheduled for later cycles.

We simulate these surveys by first choosing a "true" luminosity function $\Phi_{\rm avg}$ to be a Schechter function with redshift-dependent parameters from Finkelstein et al. (2015) (see Table 3.2). This choice does not strongly affect our results; we examine other "true" luminosity functions in section 3.4.6. For each field, we define a local luminosity function $\Phi_{\rm loc}$ (see eq. 3.2) and draw a random value from a normal function for $\tilde{\delta}$; $\varepsilon_{\rm cv}$ for each field is calculated with the python package `galcv` (Trapp & Furlanetto, 2020), which uses the energy-regulated galaxy model of Furlanetto et al. (2017) and the halo mass function of Trac et al. (2015) converted into a conditional halo mass function with a scaling method from Tramonte et al. (2017). Finally, we draw sources randomly from each field's "observed" luminosity function $\Phi_{\rm obs}$ (see eq. 3.3) with completeness curve $f(M_{\rm AB})$ by splitting $\Phi_{\rm obs}$ into small magnitude bins (0.05 magnitudes per bin) and drawing galaxies from a Poisson distribution around the expected value for each bin.

We create completeness curves $f(M_{\rm AB})$ by scaling the HST completeness curves from Finkelstein et al. (2015) to the appropriate magnitude limits. This procedure is only approximate but should serve as a reasonable stand-in for our purposes. All sources in the $z = 6$

Table 3.1: Simulated survey parameters. The rms fluctuation of the dark matter density field at redshift 6 ($\sigma_{\rm PB}(z=6)$) assumes $\Delta z = 1$; this value will differ at higher redshifts both due to the growth of structure and to the smaller volume at higher redshift for fixed angular area and Δz. The value of $\sigma_{\rm PB}(z=6)$ and $\varepsilon_{\rm cv}(m=26, z=6)$ for these surveys illustrate that we are in the linear regime. Contrary to what one might infer from the growth of structure over cosmic time, $\varepsilon_{\rm cv}(m=26, z=6)$ *increases* with redshift, as the volume decreases (with fixed area and Δz) and becomes less elongated (see Figure 3.1).

Survey	Sub-Fields	App. mag lim. [mag]	Area [arcmin^2]	$\sigma_{\rm PB}(z=6)$ rms den. fluc.	$\varepsilon_{\rm cv}(m=26, z=6)$ cosmic variance
WDEEP	-	30.75	10	0.055	0.32
CEERS	-	28.97	100	0.038	0.25
JADES			236	-	-
	North medium	28.8	95	0.039	0.23
	South medium	28.8	95	0.039	0.23
	South deep	29.8	46	0.045	0.26
PRIMER			695	-	-
	COSMOS shallow	28.89	144	0.036	0.21
	COSMOS medium	29.11	108	0.038	0.23
	COSMOS deep	29.51	33	0.047	0.28
	UDS shallow	28.48	234	0.032	0.19
	UDS medium	28.89	175	0.034	0.21
PANORAMIC	150x parallel	<29.5	1,500	0.055 each	0.32 each
RSN	-	28.8	32,400	0.005	0.03

Table 3.2: Schechter parameters from Finkelstein et al. (2015).

Parameter	Redshift Dependence
$\log_{10}[\phi^*]$	$-1.58 \pm 0.3 - (0.31 \pm 0.07)z$
α	$-0.79 \pm 0.21 - (0.19 \pm 0.04)z$
M^*	$-20.27 \pm 0.42 - (0.12 \pm 0.09)z$

bin are drawn from the $z = 6$ luminosity function, and likewise for the other redshifts. We use a redshift bin width of $\Delta z = 1$ for defining the volume of the surveys. In reality, the luminosity function is changing throughout that range, but that effect is small in practice (see Appendix 3.6).

We consider many forthcoming surveys and combine them in various subsets (in the manner described in sections 3.3.3.1 and 3.3.3.2). Many of the surveys themselves have subcomponents. For the purposes of computational efficiency in making our forecasts, we treat these surveys and their subcomponents in the following manner:

(i) The WDEEP, CEERS, and Roman SN fields are each treated as single independent volumes.

(ii) For the JADES survey, we combine the JADES South medium and deep fields contiguously. We combine this independently with the Jades North medium field.

(iii) For the PRIMER survey, we combine the COSMOS shallow, medium, and deep fields contiguously, and we combine the UDS shallow and medium fields contiguously.

(iv) PANORAMIC has approximately 150 separate pointings. We separate these into two groups by depth and combine each subset independently.

When we combine many surveys, we make some additional combinations for computational convenience. In particular, we combine the PRIMER COSMOS, PRIMER UDS, JADES North, and JADES South fields together independently when all four fields are present in one survey combination. When the PANORAMIC survey is combined with other surveys, the two separate depth groups are combined independently. Figure 3.2 shows the simulated data for one of the combinations of surveys.

We apply the framework described in section 3.3 to many simulated survey combinations to obtain Schechter parameter $\vec{\phi}$ posteriors for the average luminosity function and density $\vec{\delta}$ posteriors for the individual fields. We are interested in the width of these posteriors as a measure of how well that survey combination constrains the luminosity function and field densities. However, each instance of simulated surveys is affected by both cosmic variance and Poisson noise; different realizations of the same surveys will have differently-shaped posteriors. In order to get an understanding of how well a survey combination does on average, we repeat this whole process N times for each survey. We choose N such that the standard error in determining the average posterior width for each parameter is less than 10%.

In Table 3.1, the 2 right-most columns shows the rms linear density fluctuation σ_{PB} (at $z = 6$) and cosmic variance (at apparent magnitude 26) for each of the fields. The values of σ_{PB} are all much less than unity, indicating we are within the linear regime of structure formation. $\varepsilon_{\mathrm{cv}}$ also remains well below unity, indicating our assumption of a gaussian local luminosity function (eq. 3.2) is valid in these cases. However, as redshift increases, so does $\varepsilon_{\mathrm{cv}}$ (for fixed area and Δz); when $\varepsilon_{\mathrm{cv}}$ becomes larger than ∼0.5, the assumption of a gaussianity becomes worse. Fortunately, when $\varepsilon_{\mathrm{cv}}$ becomes larger than 0.5, Poisson noise typically takes over as the dominant source of uncertainty (see Figure 3.1).

Finally, for a real-data comparison, we have applied our framework to the multi-field Hubble Space Telescope data set from Finkelstein et al. (2015) ($6 \leq z \leq 9$ only) and the wide-field ground-based data from Bowler et al. (2015). We describe the procedure fully in Trapp et al. (in prep.); here we include the results only for context in understanding the improvements JWST will provide.

3.4.2 Survey posteriors and validation

In this section we verify that our framework recovers the input “true” luminosity function parameters and randomly drawn field densities. In the next sections, we explore how well the framework constrains these parameters.

Figure 3.3 shows the posterior of the Schechter parameters (marginalized over field densities) for the average luminosity function at $z = 9$ for one of the simulated survey combinations. The “correct” input values (from Finkelstein et al., 2015) are shown as vertical black lines. For this survey combination, there is a strong degeneracy between the normalization and the location of the faint-end cut-off (bottom left panel). In general, our framework successfully recovers the correct input Schechter values except for cases where the surveys sample only the faint-end of the luminosity function, making the exponential cutoff impossible to localize. In these cases, the faint-end slope is still well recovered, but a two-parameter power-law model should be used instead of a full Schechter function. In the cases where the exponential cutoff is just barely covered by data, the fitting framework is biased towards choosing a larger normalization, shallower faint-end slope, and lower-luminosity exponential cutoff. This happens for the smaller-area ($A \lesssim 500$ arcmin^2) survey combinations at $z \gtrsim 9$.

In this work, we do not explicitly compare our inference framework to other fitting methods. However, in Trapp & Furlanetto (2020), we found fitting with cosmic variance in mind can result in a less biased recovery and better overall fit of the luminosity function than the standard method used by e.g., Bouwens et al. (2015), Finkelstein et al. (2015), which assumes the luminosity function has the same shape in all environments. This result becomes more important at higher redshifts, which is reflected in Figure 3.1: the luminosity function's shape changes with environment, and more-so as redshift increases (see Trapp & Furlanetto, 2020, for more details).

Figure 3.4 shows the posterior of the field densities at $z = 9$ for one of the simulated surveys. The "correct" input values (drawn randomly for each field) are shown as vertical black lines for fields that have not been combined with others. Our framework's density posteriors consistently recover the correct input density values. The resulting density posterior for PANORAMIC is not very informative. This is because PANORAMIC's density is an *effective* density, a combination of multiple independent fields as described in section 3.3.3. PANORAMIC's density is less well constrained than the other fields because it has a very small $\varepsilon_{\text{cv, eff}}$ value, meaning a change in density is very difficult to distinguish from Poisson noise.

3.4.3 Measuring the luminosity function

A primary goal of JWST's extragalactic mission is to understand the physics and formation of high-z galaxies. The luminosity function is the cornerstone of this effort and the key observable that models and simulations compare to. We now examine the constraining

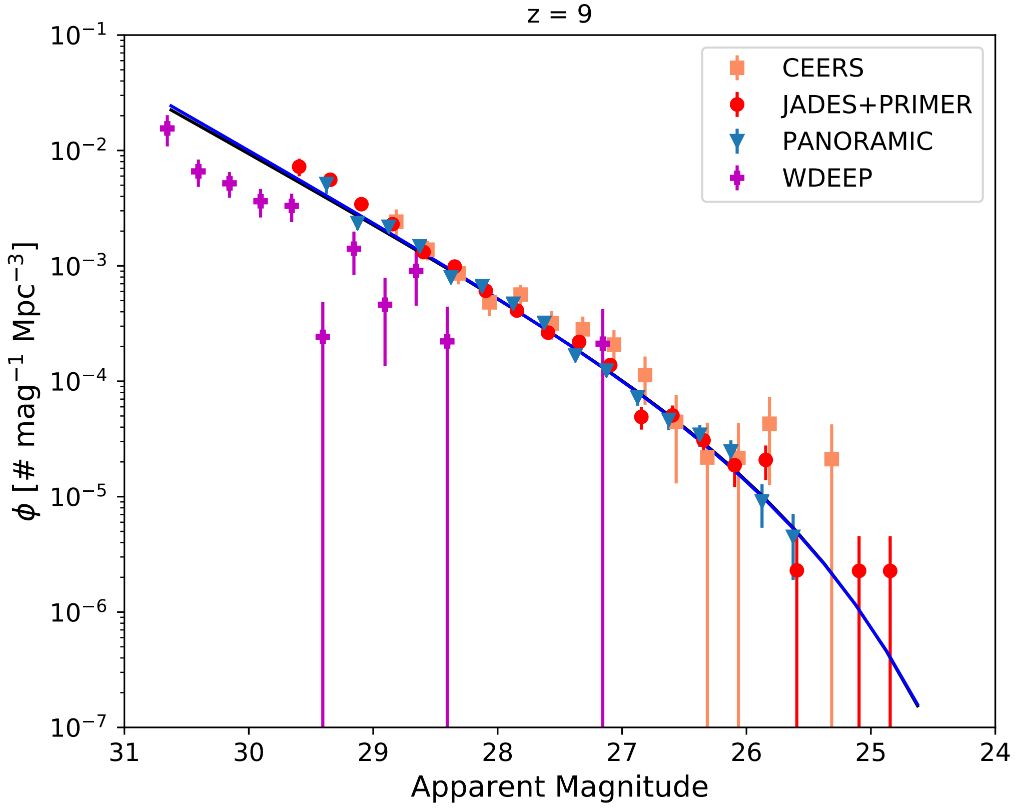

Figure 3.2: Data from one simulation of the CEERS+JADES+PRIMER+PANORAMIC+WDEEP surveys at $z = 9$. *Black* line: input luminosity function; *blue* line: best fit luminosity function. In this case, the WDEEP pointing happens to be a very under-dense region (see Figure 3.4), but the "true" input luminosity function is still recovered.

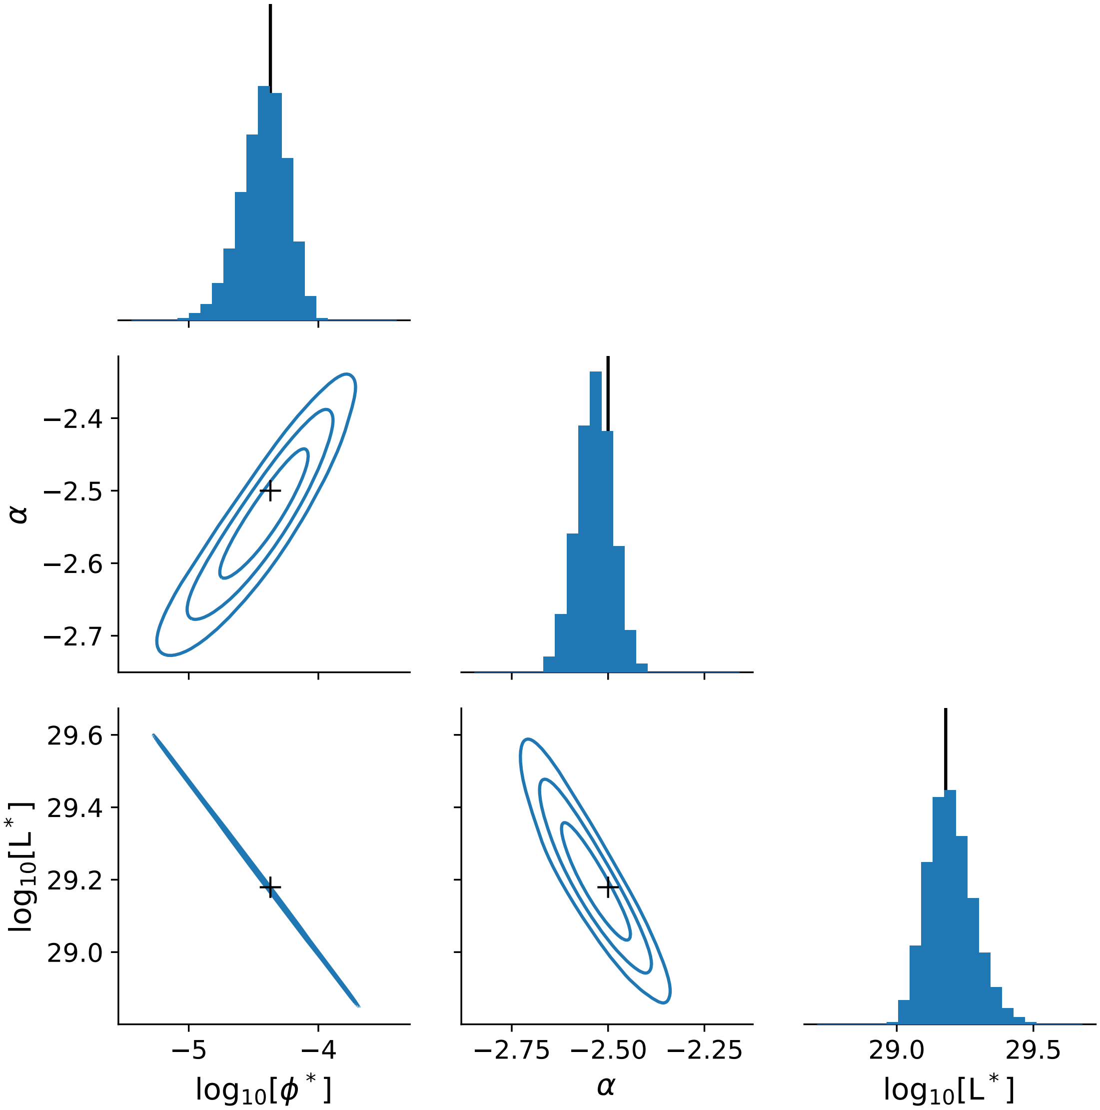

Figure 3.3: The Schechter parameter posterior (marginalized over all densities) for one simulation of the CEERS+JADES+PRIMER+PANORAMIC+WDEEP surveys at $z = 9$. The vertical black lines and the '+' marks show the "true" input values. At this redshift, the bright-end cutoff is barely covered by the data (see Figure 3.2), leaving a large (and tight) degeneracy between the normalization and characteristic luminosity.

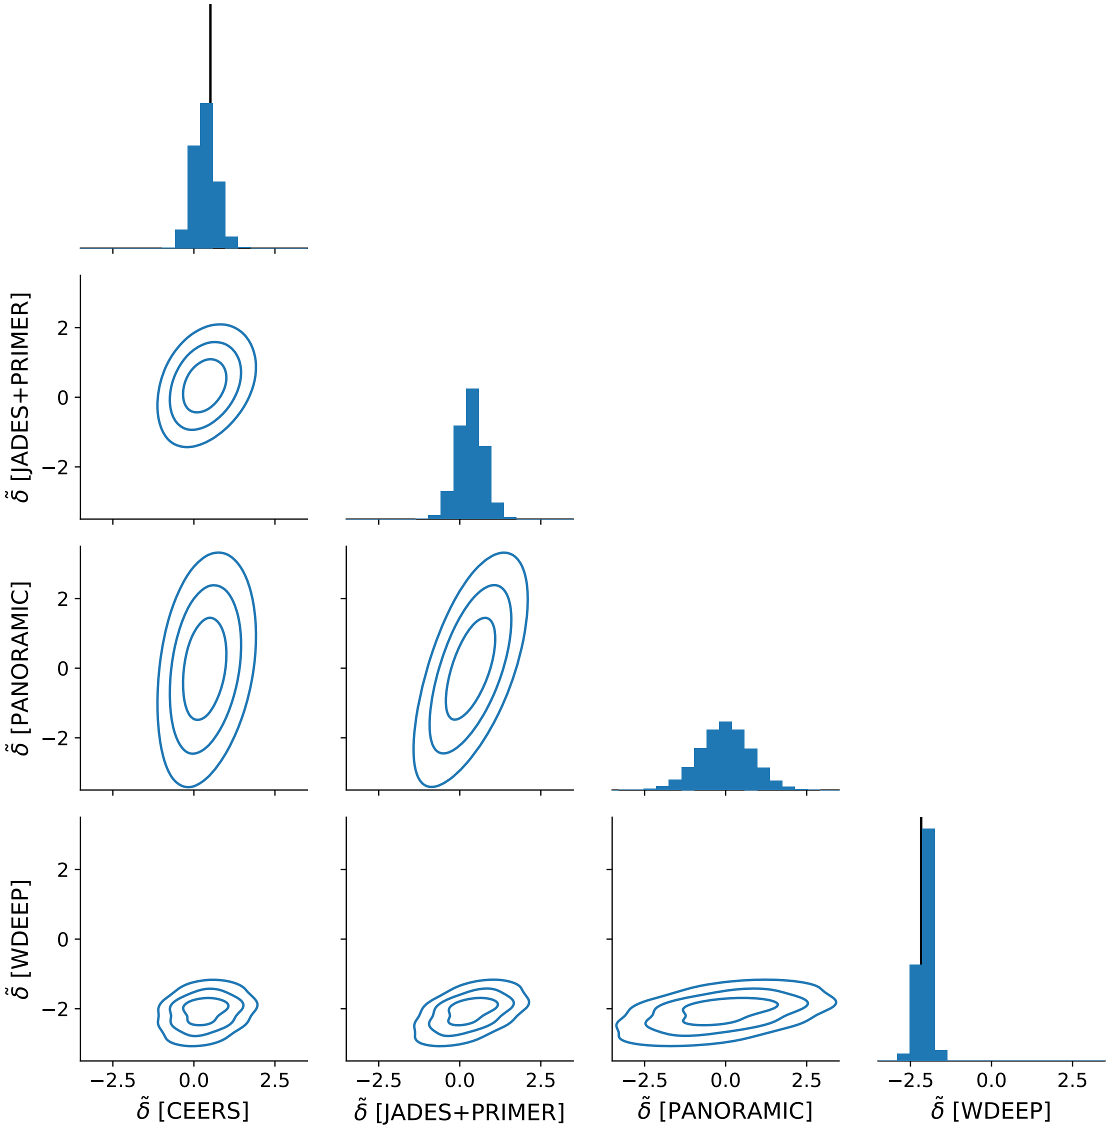

Figure 3.4: The density posterior (marginalized over all Schechter parameters) for one simulation of the CEERS+JADES+PRIMER+PANORAMIC+WDEEP surveys at $z = 9$. The vertical black lines on the histograms show the "correct" input density for CEERS and WDEEP. The "correct" densities for JADES+PRIMER and PANORAMIC are not shown as they are effective densities, not physical ones.

power of the simulated JWST surveys described in the previous sections. Individual JWST surveys will be effective tools for measuring the luminosity function, but combining surveys gives even better constraints. We quantify those improvements by comparing the widths for the luminosity function parameters. We first compare the PRIMER, CEERS+JADES, and PANORAMIC individual performances. Then, we test various combinations of surveys using the PRIMER+CEERS+JADES survey as a baseline.

Figure 3.5 displays the 68.27% confidence interval widths of the posterior of the luminosity function parameters for redshifts $z = 6-12$ for the simulated PRIMER, CEERS+JADES, and PANORAMIC surveys. PRIMER and CEERS+JADES perform similarly, with PRIMER providing a slightly better measurement of the exponential cutoff due to its larger volume. Both PRIMER and CEERS+JADES would halve the uncertainty in the faint-end slope compared to HST at $z = 8$ and significantly improve the normalization measurement. PANORAMIC by itself offers a significant improvement over PRIMER or CEERS+JADES (see Figure 3.5). The combination of all JWST surveys is shown for comparison.

Figure 3.6 shows several additional survey combinations. Combining CEERS+JADES with PRIMER (C+J+P) offers a significant improvement (green line). Adding WDEEP to C+J+P improves the measurement of the faint-end slope much more than the other two parameters, because it goes significantly deeper than any other survey. Adding just PANORAMIC to C+J+P improves the constraints on all parameters. When combining all JWST surveys, we (unsurprisingly) get the best constraints (black line). With all surveys combined, JWST will measure the faint-end slope at $z = 12$ to similar precision as HST measures at $z = 6$–7. JWST will measure the normalization and exponential cutoff at $z = 12$ to similar precision as HST measures at $z = 8$.

We add one additional further-future survey, the Roman SN survey, which will vastly improve the measurement of the exponential cutoff and normalization due to its large volume. Perhaps surprisingly, the faint-end slope is also much better measured. With the degeneracy between the exponential cutoff and normalization broken, the faint-slope is easier to pin down with the JWST surveys. This effect highlights the importance of wide-field searches for the brightest objects (see also Kauffmann et al., 2020).

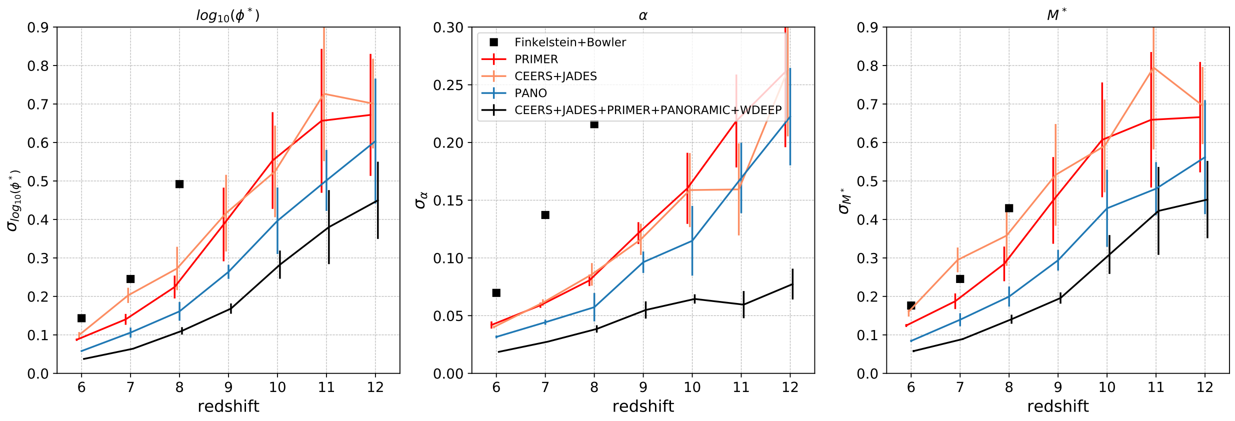

Figure 3.5: The 68.27% confidence interval width of the Schechter parameters (averaged over many simulations) for different survey combinations as a function of redshift. The error bars represent the standard deviation of the posterior widths between simulations, which is due to Poisson noise as well as different random draws of the individual fields' densities. The black points are results from applying our framework to data from Finkelstein et al. (2015) and Bowler et al. (2015) for a comparison to HST's capabilities. Curves are shifted slightly in the x-direction to avoid overlapping error bars.

One may wonder how much the parallel nature of PANORAMIC helps in the constraints. To test this, we also considered an artificial survey that is identical to the PANORAMIC survey but in two contiguous fields (one incorporating the deeper pointings and one the shallower ones). The constraints are nearly identical to those from the full PANORAMIC, likely due to the fact that the area is so large (1500 arcmin^2) that cosmic variance's effects are small even if it were a contiguous field. However, the parallel nature of PANORAMIC provides other unique opportunities (see section 3.4.5).

3.4.4 Measuring field densities

We have now seen that, in the JWST era, galaxy abundances throughout the Cosmic Dawn will be measured to unprecedented precision. But our new method also opens an entirely new window onto that era, by providing measurements of the large-scale environments of survey fields. At least in principle, such environmental measurements will allow us to associate

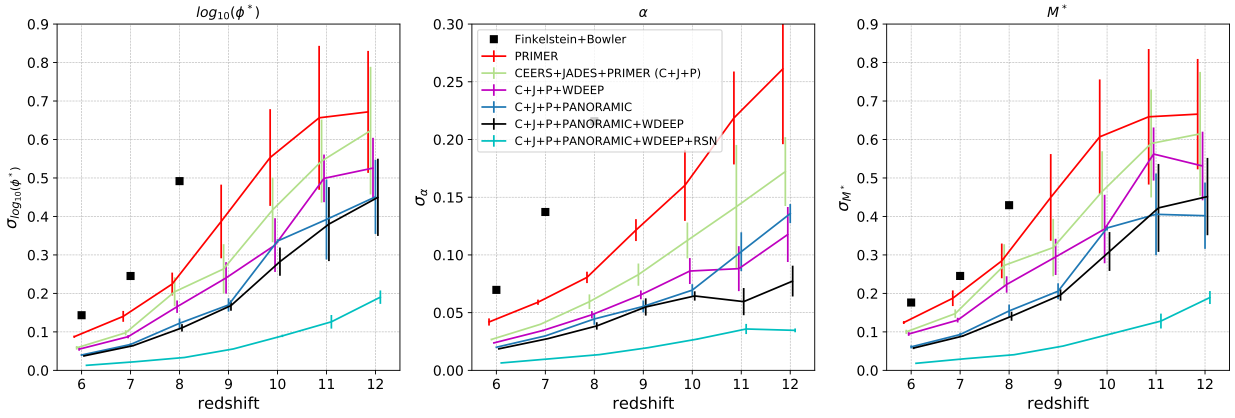

Figure 3.6: Same as Figure 3.5 but showing different combinations of surveys. The order of the lines in the legend (top to bottom) is the same as the order of the lines at $z = 8$ (top to bottom) for all three panels.

galaxies with their ionization environments, their future descendant populations, and to conduct searches for how galaxy evolution depends on those environments.

The density of a single field can only be determined in comparison to other fields. As more fields are added, the measurement of the average luminosity function's normalization gets better, and so too does the measurement of an individual field's density. This effect can be seen in the density posteriors for our fields. Figure 3.7 shows the 68.27% confidence interval width of the marginalized posterior for the CEERS and WDEEP densities as a function of redshift for various combinations of surveys. The errorbars reflect the natural variance in the posterior widths from the effects of sample variance and Poisson noise.

With CEERS and JADES alone, the CEERS environment can be measured with a posterior width[4] of 0.6–0.7. Combining CEERS+JADES with PRIMER improves this measurement significantly. Adding a large parallel survey like PANORAMIC does an even better job at improving the density measurement of the CEERS field. Adding a deep field like WDEEP to C+J+P does not greatly increase the measurement of the environment as WDEEP is a small volume with large cosmic variance and thus does not help constrain the normalization

---

[4] $\tilde{\delta}$ is in units of standard deviations away from mean density. An over-density of $\tilde{\delta}$ means there are a factor of $(1 + \varepsilon_{\rm cv}\tilde{\delta})$ more/fewer galaxies than average. The prior on $\tilde{\delta}$ is a normal with width $\sigma_{\tilde{\delta}} = 1$.

as well as other surveys. When combining all surveys' data, JWST can measure the density of a 100 arcmin^2 field (like CEERS) to a precision of 0.3 and 0.6 at $z = 6$ and 12, respectively, and a deep 10 arcmin^2 field (like WDEEP) to a precision of 0.35 or 0.45 at $z = 6$ and 12, respectively.

Clearly, more data result in tighter constraints on the density of any given field. However, the results appear to asymptote as more and more surveys are added, indicating there is a fundamental limit to measuring a field's density. That fundamental limit is Poisson noise, plotted as the gray dashed line. This "Poisson floor" is calculated by simulating a survey field with Poisson noise and some density $\tilde{\delta}$ drawn from a normal distribution. Then, the posterior of $\tilde{\delta}$ is calculated in a similar manner to the posterior calculation in section 3.3.2, but assuming perfect knowledge of the luminosity function. The width of that posterior represents the best-case determination of the density, which we refer to as the 'Poisson floor'. This floor depends slightly on the field's actual density (which was drawn randomly), and the variance from that effect is represented by the error-bars in the Poisson floor. The Poisson floor also depends weakly on the assumed luminosity function and strongly on the completeness function of the survey.

We plot the Poisson floor for a variety of our fields in Figure 3.8. From testing with our framework, we find JWST will be able to measure all of its fields' densities to near this limit when multiple surveys are combined.

3.4.5 Testing galaxy physics in different environments

A large parallel survey like PANORAMIC will probe a wide range of densities (150 independent pointings). After measuring the global luminosity function as in section 3.4.3, we can go back and measure the density of each of the PANORAMIC fields using equations (3.13) and (3.14). This time, the luminosity function prior is determined by the earlier global fit, over whose errors we can marginalize.[5]

---

[5] An individual PANORAMIC field contributes so little to the global fit of the luminosity function that it can safely be considered independent, as this description is assuming.

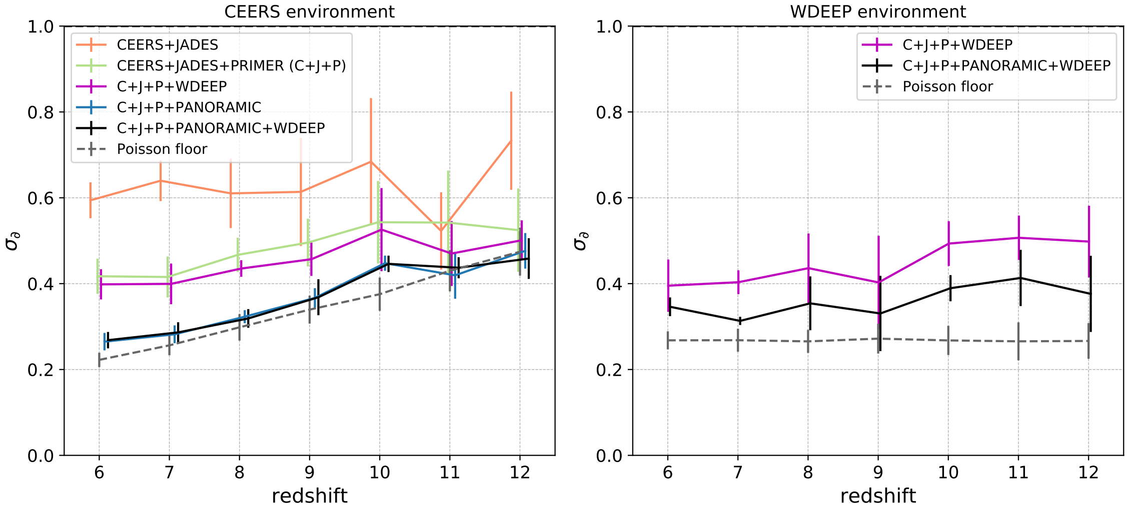

Figure 3.7: The 68.27% confidence interval widths of the CEERS and WDEEP field densities (averaged over many simulations) for different survey combinations as a function of redshift. The error bars represent the standard deviation of the width in the posterior between runs, which is due to Poisson noise as well as different random draws of the individual fields' densities. The dashed black line shows the maximum accuracy in determining the field density (see section 3.4.4). Curves are shifted slightly in the x-direction to avoid overlapping errorbars.

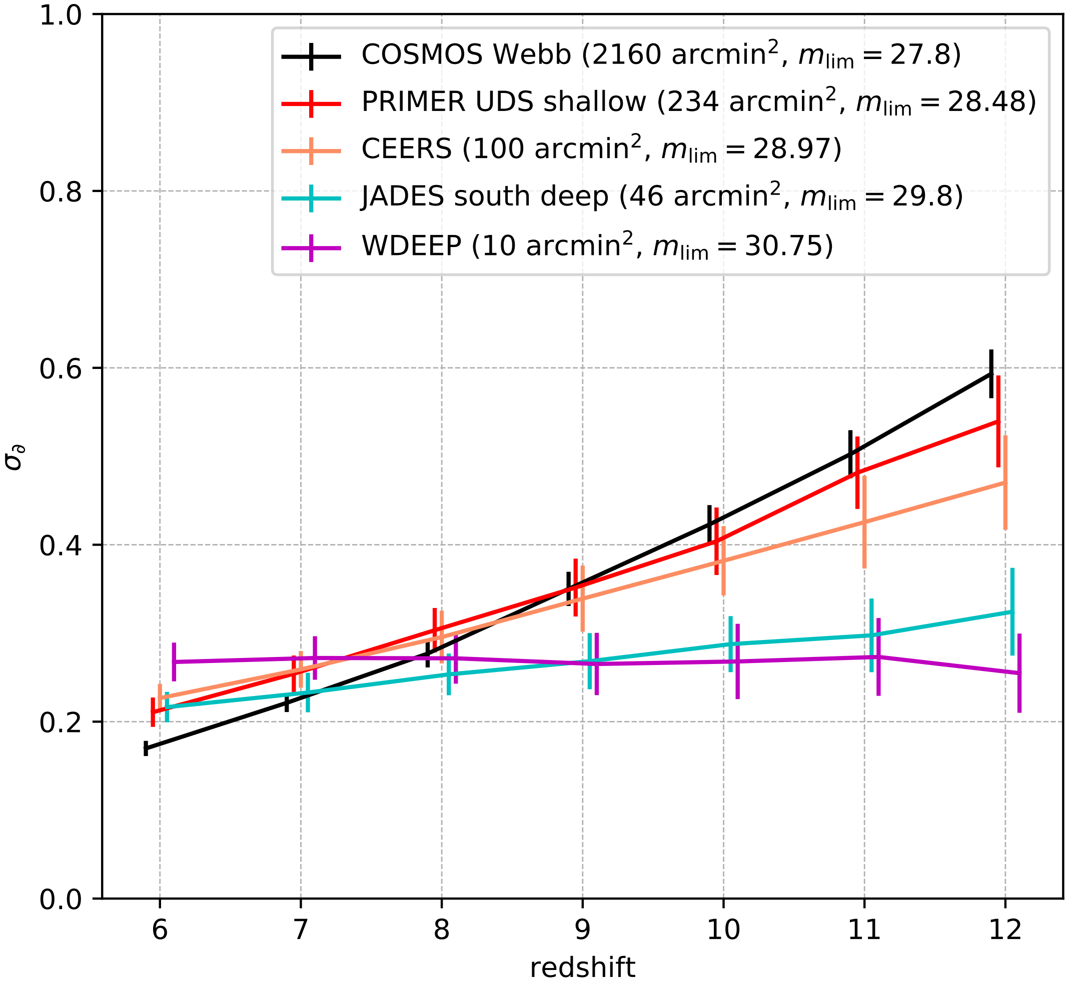

Figure 3.8: The best possible precision in measuring the density for various contiguous JWST fields (see section 3.4.4). This limit is due to Poisson noise only; the luminosity function is assumed to be perfectly known. The prior on the density is $\sigma_\partial = 1$. A deeper survey has a lower and flatter curve (at fixed area), and a larger area survey has a lower curve (at fixed depth). Curves are shifted slightly in the x-direction to avoid overlapping errorbars.

While an individual PANORAMIC pointing will have an uncertainty in its density ($\sigma_{\tilde{\delta}}$ ∼0.4-0.8, from testing), they can be sorted from least to most dense. One can then search for correlations between dark matter density and galaxy properties, ionization history, etc. A similar procedure could be applied to large contiguous fields like PRIMER, with the added complication that the field densities are correlated, being physically next to one-another.

We do caution the reader that we have not incorporated uncertainties in the galaxy models themselves into our inference framework; in particular, we assume a perfectly known cosmic variance function as expressed by $\varepsilon_{\rm cv}$. In the future, uncertainties in the galaxy model should be marginalized over. Fortunately, they are not large (see section 3.3.1), and at least in principle can be separated because the cosmic variance functions depend on halo mass (and hence galaxy luminosity).

3.4.6 Limitations and future improvements

In our simulations, we use a luminosity function calibrated at low redshift ($z \approx 5 - 9$) and extrapolated to higher redshift (Finkelstein et al., 2015, see our Table 3.2). We cannot know if this extrapolation holds beyond the redshift where we have data, and we will not know until JWST and other observatories make their measurements. However, this extrapolation has reported uncertainties, so we re-do our analysis for a more slowly-evolving and more rapidly-evolving case based on those uncertainties. The slowly-evolving case has a shallower (increased) slope with redshift and decreased y-intercept for all Schechter parameters, and the rapidly-evolving case has steeper (decreased) slope and increased y-intercept (see Table 3.2).

The slowly-evolving case is similar to the fiducial case when measuring $\log_{10}(\phi^*)$ and M^*, but it generally has a worse determination of the faint-end slope α with up to ∼30% wider posteriors, depending on the combination of surveys and the redshift considered. The rapidly-evolving case is similar to the fiducial case when measuring α, but it generally has a worse determination of $\log_{10}(\phi^*)$ and M^* with up to ∼40% wider posteriors, again depending on the survey combination and redshift. These variations are usually smaller in scale than the variations between individual trials of the surveys, meaning our predictions are relatively robust to different luminosity functions.

As discussed in section 3.3.1, the cosmic variance function $\varepsilon_{\rm cv}$ has an estimated uncertainty of ∼25%. To test the effects of a different "true" cosmic variance function, we re-do our analysis again with a 25% lower/higher value for $\varepsilon_{\rm cv}$ (when generating simulations). We find that we recover $\log_{10}(\phi^*)$ and M^* to similar precision as the fiducial case, but generally have a worse determination of the faint-end slope α with up to ∼25% wider posteriors depending on survey combination and redshift.

One of the benefits of our framework is its potential for constraining the cosmic variance function itself from the data by treating it as a free parameter with a prior based on its estimated uncertainty and marginalizing the posterior space over all other parameters. However, we do not explore that potential constraining power in this work.

In our calculation of the posterior, we assume that all measurements of galaxy magnitudes (real and simulated) have no uncertainty. Taking these uncertainties into account would make the measurement of the bright end of the luminosity function (where it is steepest) more realistic, but also more uncertain and require more computation time. We plan to explore this trade-off in the future.

We have described the conditional luminosity function (eq. 3.2) as being gaussian with respect to the density $\tilde{\delta}$. However, as $\varepsilon_{\rm cv}$ approaches unity, the conditional luminosity function (at fixed magnitude) is better described by a log-normal or gamma function with standard deviation equal to $\varepsilon_{\rm cv}$ (see Figure 6 in Trapp & Furlanetto, 2020). To correct for this, we switch from gaussian to log-normal when $(1 + \tilde{\delta}\varepsilon_{\rm cv}) \leq 0$ (when the number density becomes negative). This correction is approximate; a full treatment requires the full conditional luminosity function (see Trapp & Furlanetto, 2020). However, this problem will affect only the brightest sources in the smallest volumes where Poisson noise is typically dominating anyway.

We use a very simple estimate of the completeness function for the JWST surveys. Actual completeness functions will require detailed knowledge of the telescope's performance and sophisticated simulations (see Kauffmann et al., 2020, for detailed predictions of JWST's completeness). If our estimates are off in normalization, that will increase the effect of Poisson noise but should not change the qualitative results much. If the shape of our completeness function is wrong, it will likely only affect the faint end of the luminosity function where the completeness function deviates significantly from its maximum value. The completeness can also moderately affect the Poisson floor which is sensitive to the faint-end of the luminosity function where the largest number of sources are found.

3.5 Conclusions

We develop a framework that calculates the posteriors of the average galaxy luminosity function and field densities simultaneously. We also develop a method to combine independent and/or contiguous fields into "effective" fields with their own cosmic variance functions, reducing computation time and providing an estimate for the effects of cosmic variance on a complicated set of surveys.

We find that JWST will improve and extend measurements of the luminosity function, with precision at $z = 12$ roughly equal to what HST is capable of at $z = 8$. With all early-cycle JWST programs combined, we expect to measure the normalization of the luminosity function to a precision of 0.05 and 0.6 dex at $z = 6$ and 12, the faint-end power-law index to a precision of 0.03–0.10 over the same redshift interval, and the characteristic magnitude to a precision of 0.08 and 0.55 at $z = 6$ and 12 (see Figure 3.6).

Large-area surveys are most important for breaking the degeneracy between the normalization and the bright-end cutoff. For the highest redshifts ($z \gtrsim 9$), a single power-law fit will likely be better than a Schechter fit even if the underlying population is a Schechter function, as probing the bright-end cutoff will require Roman-sized survey areas. Parallel surveys are subject to less cosmic variance than contiguous surveys, but at sufficiently large area, contiguous surveys are mostly unaffected by cosmic variance.

When combined, early-cycle JWST galaxy surveys will be able to measure the normalized densities $\tilde{\delta}$ of individual survey fields with a 68.27% confidence interval of ~0.4 for deep fields such as WDEEP, and ∼0.5 for larger fields like CEERS, nearly their theoretical maximum precision (see Figures 3.7 and 3.8). For these purposes, consistent data reduction between different surveys will be crucial, as the determination of field density depends on the relative normalization between surveys.

Measurements of large-scale environments will open new opportunities to study galaxy evolution during the Cosmic Dawn. Most importantly, reionization is driven by the underlying matter distribution (through the biased formation of galaxies; Furlanetto et al. 2004), so density measurements can also provide estimates of the local ionization state (albeit in a model-dependent fashion). Interestingly, the transverse scales of many anticipated fields are comparable to the ionized bubbles that appear through reionization – for example, a comparison to Lin et al. (2016) and Davies & Furlanetto (2022) shows that WDEEP and PANORAMIC fields are comparable in extent to bubbles halfway through reionization, while CEERS and JADES fields are close to the sizes in the later stages. Studies across such fields will allow exploration of the effects of reionization on the local galaxy population – which has long been expected (e.g., Thoul & Weinberg 1996; Iliev et al. 2007; Noh & McQuinn 2014) but never observed directly. Density measurements will also allow studies of more conventional effects of environment on galaxy evolution and the quantitative association of high-z environments with their descendants.

Finally, density measurements allow targeted exploration of unusual environments (such as protoclusters) whose histories have long been studied. While our method provides only a first step toward such ambitious goals, it demonstrates that such inferences will soon be possible with forthcoming observational programs.

Our forecasts demonstrate the transformative potential of future galaxy surveys in understanding the Cosmic Dawn on both large and small scales. While harvesting this information will require improvements to the model (such as localizing galaxies in the radial direction, accounting for correlations between neighboring volumes, and incorporating uncertainties in the galaxy physics itself), the enormous potential provides strong motivation for such future efforts.

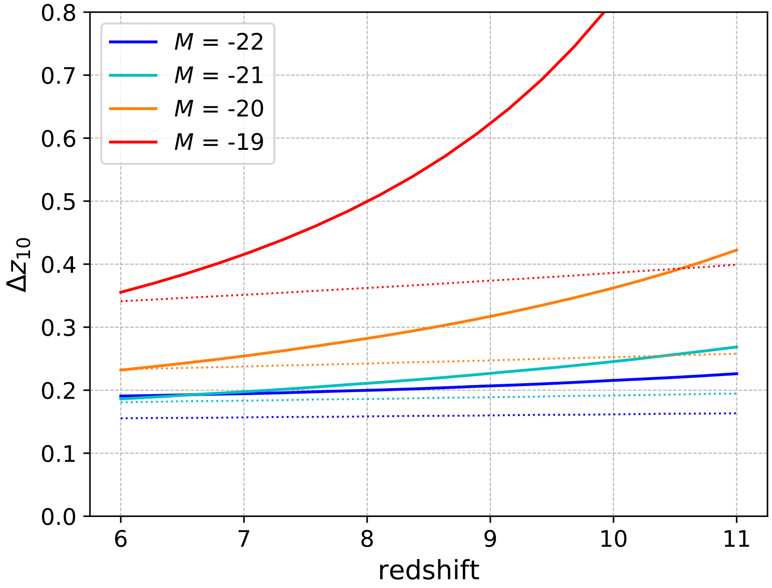

Figure 3.9: The change in redshift (towards lower redshift) that corresponds to a 10% growth in the luminosity function. The dotted lines are the same but for the more slowly-evolving luminosity function (see section 3.4.6).

3.6 Appendix: Choosing Optimal Redshift Bins

It is ideal to measure the luminosity function in as narrow a redshift range as possible, as the luminosity function is constantly evolving. However, it is often necessary to use large redshift bins to get a sufficient number of galaxies for fitting. How large then is too large? We define a function $\Delta z_{10}(z, M_{\mathrm{AB}})$: the redshift change that corresponds to a 10% growth in the luminosity function at magnitude M_{AB}. Figure 3.9 shows $\Delta z_{10}(z, M_{\mathrm{AB}})$ for a selection of magnitudes. Figure 3.9 would suggest that Δz should be below one, especially for the bright end of the luminosity function. However, this can be mitigated by averaging over a wider bin.

We calculate the volume averaged luminosity function over the redshift range Δz centered at z_{c} and then find the corresponding redshift with the same luminosity function value, defining that redshift as the 'effective' redshift of the range $z_{\mathrm{eff}}(M_{\mathrm{AB}}, \Delta z, z_{\mathrm{c}})$. The 'effective' redshift is a function of magnitude, and it is what *should* be used as the true center redshift,

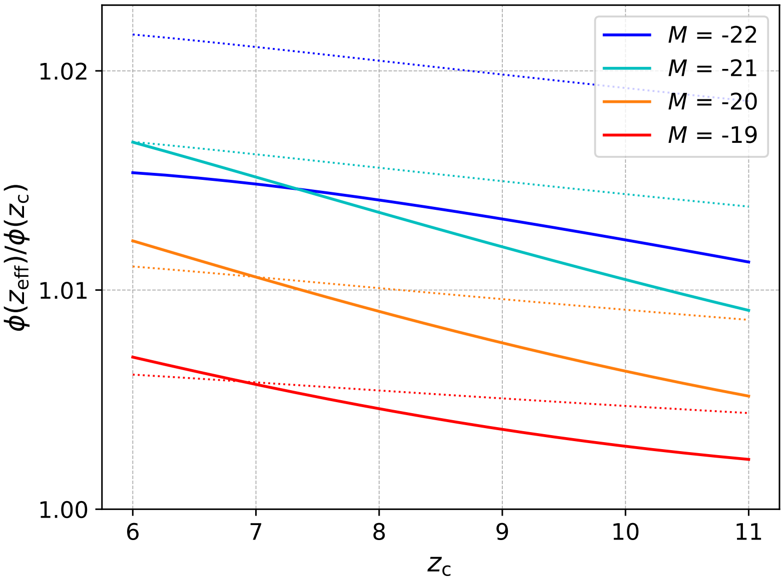

Figure 3.10: The ratio of the luminosity function at $z_{\rm eff}$ (the effective center of a bin centered at $z_{\rm c}$ with $\Delta z = 1$) and the luminosity function at $z_{\rm c}$. The dotted lines are the same but for the more slowly-evolving luminosity function (see section 3.4.6).

not $z_{\rm c}$. For survey estimates, the key question is the error introduced by using the central redshift instead of this (magnitude-dependent) effective value $z_{\rm eff}$.

We plot the ratio between the luminosity function at $z_{\rm eff}$ and $z_{\rm c}$ for a variety of magnitudes and redshifts in Figure 3.10. We find no significant difference ($\lesssim$2%) when using a redshift bin width of $\Delta z = 1$. Increasing Δz to 2 results in deviations increasing to $\sim$7% for the brightest sources.

These calculations depend on the shape and evolution of the luminosity function. By default, we use the Schechter function fit from Finkelstein et al. (2015). The dotted lines in Figures 3.9 and 3.10 are the same but for the more slowly-evolving luminosity function (see section 3.4.6); this effect does not strongly depend on luminosity function choice.

Because the errors introduced by the redshift binning are much smaller than our expected uncertainties, we use a constant z_c in the main text.

CHAPTER 4

The first joint measurement of galaxy luminosity functions and large-scale structure during the Epoch of Reionization

4.1 Abstract

One of the most exciting advances of the current generation of telescopes has been the detection of galaxies at the tail end of reionization, using deep fields that have pushed these instruments to their limits. It is essential to optimize our analyses of these fields in order to extract as much information as possible from them. In particular, standard methods of measuring the galaxy luminosity function discard information on large-scale dark matter density fluctuations, even though this large-scale structure drives galaxy formation and reionization during the Cosmic Dawn. Measuring these densities would provide a bedrock observable, connecting galaxy surveys to theoretical models of the reionization process and structure formation. Here, we use existing Hubble deep field data to simultaneously fit the universal luminosity function and measure large-scale densities for each Hubble deep field at $z = 6$–8 by directly incorporating priors on the large-scale density field and galaxy bias. Our fit of the universal luminosity function is consistent with previous methods but differs in the details. For the first time, we measure the underlying densities of the survey fields, including the most over/under-dense Hubble fields. We show that the distribution of densities is consistent with current predictions for cosmic variance. This analysis on just 17 fields is a small sample of what will be possible with the James Webb Space Telescope, which will measure hundreds of fields at comparable (or better) depths and at higher redshifts.

4.2 Introduction

Over the past three decades, astronomers have put enormous effort – and time with facilities like the Hubble Space Telescope (HST) – into observing the most distant galaxies. As we approach the era of the James Webb Space Telescope (JWST), we expect a revolution in our understanding of the early Universe. JWST's first observational campaigns will uncover many interesting and complex phenomena in the Cosmic Dawn, but interpreting these new observations will require a solid bedrock of survey analysis.

A key observable of the Cosmic Dawn is the galaxy luminosity function, which describes the galaxy population and its growth as a whole. Much effort has been put into its study, as its evolution in shape and normalization have important implications for the ways galaxies form and evolve (see e.g., Schenker et al., 2013; McLure et al., 2013; Bouwens et al., 2015; Finkelstein et al., 2015; Bowler et al., 2015; Livermore et al., 2017; Atek et al., 2018; Oesch et al., 2018; Behroozi et al., 2019; Bouwens et al., 2021; Finkelstein et al., 2022). These studies have pinned down the abundance of relatively bright galaxies at $z \lesssim 8$, with results largely consistent with models extrapolated from lower redshift (see e.g., Tacchella et al., 2013; Mason et al., 2015; Furlanetto et al., 2017; Mirocha et al., 2017). However, above $z \gtrsim 9$, galaxies are currently too rare to decisively measure their abundances, although the observations still provide important insights into early galaxies (Oesch et al., 2013, 2015; Bouwens et al., 2015; Ishigaki et al., 2015; McLeod et al., 2015, 2016; Bouwens et al., 2019; Roberts-Borsani et al., 2022). These measurements have been possible thanks to several large observing campaigns across several fields: only by combining many such efforts have astronomers managed to obtain the current constraints.

One of the (many) challenges in measuring the luminosity function is the uncertainty due to cosmic variance[1]: the normalization and shape of the luminosity function differ between distinct volumes due to fluctuations in the large-scale dark matter density field (see

---

[1] In this paper, we use the term "cosmic variance" to describe dark matter density fluctuations between volumes in our Universe and the subsequent consequences for the galaxy population. To be precise, this is a case of sample variance. The term cosmic variance is sometimes reserved for the errors stemming from having only one Universe to observe.

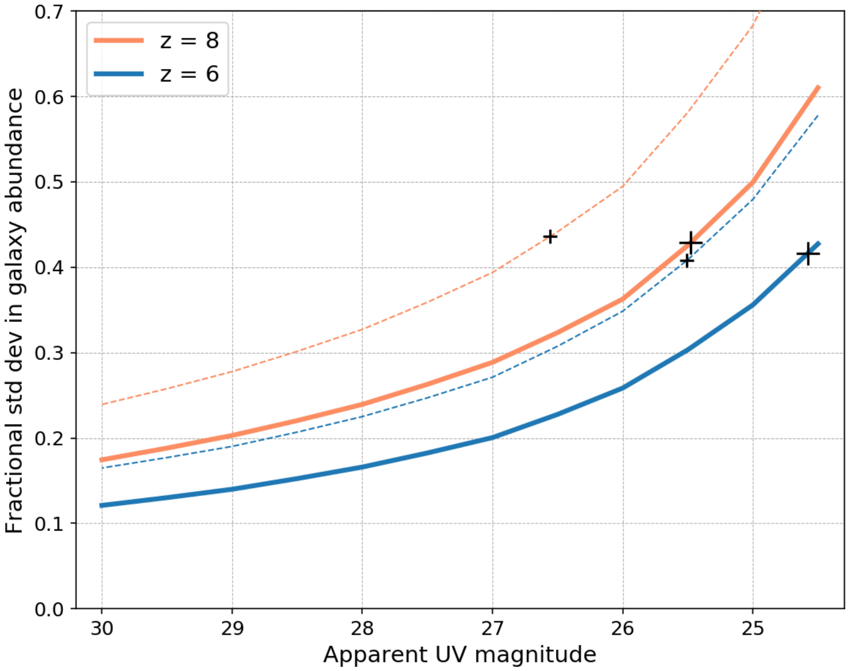

Figure 4.1: The strength of cosmic variance as a function of apparent magnitude for redshifts 6 and 8. A value of e.g., 0.3 corresponds to a galaxy over-density of 30% for a 1-σ over-dense region (not accounting for Poisson noise). *Thick solid* lines are for a 50 arcmin^2 survey and *thin dashed* lines are for a 5 arcmin^2 survey, both with $\Delta z = 1$. The strength of cosmic variance becomes more dependent on magnitude at higher redshift and for smaller volumes; cosmic variance significantly affects the shape of the luminosity function in these cases. The '+' markers show where we would expect to find ∼1 source at the indicated magnitude in a survey. This is to mark where Poisson noise dominates.

Figure 4.1, and Trapp & Furlanetto, 2020). However, to the extent that it reflects real large-scale structure in the Universe, cosmic variance is not just a nuisance; it is itself a key driver of both galaxy formation and reionization during the Cosmic Dawn. If large-scale densities can be measured, they can complement the luminosity function as another bedrock observable. The insights to be gained from such measurements include: *(i)* Reionization likely began in the densest parts of the Universe and ended in the largest voids. Identifying such over/under-densities is an area of great interest (see e.g. Zitrin et al., 2015; Jung et al., 2020; Tilvi et al., 2020; Hu et al., 2021; Endsley et al., 2021; Becker et al., 2018; Davies et al., 2018; Christenson et al., 2021). *(ii)* Large-scale feedback mechanisms, driven by large-scale structure, are likely to strongly affect the galaxy population before and during reionization (Thoul & Weinberg, 1996; Iliev et al., 2007; Noh & McQuinn, 2014). *(iii)* Measuring large-scale densities at early times facilitates the understanding of the assembly history of rare objects like galaxy clusters (e.g., Chiang et al. 2017), which form from the densest environments. *(iv)* Finally, comparing large-scale density measurements from surveys with theoretical predictions of cosmic variance can help test models of the galaxy–halo connection (Trapp & Furlanetto, 2020).

In Trapp et al. (2022), we developed a framework that simultaneously measures field densities *and* the high-z luminosity function given a set of galaxy surveys. Unlike the standard approach to estimating luminosity functions, which acknowledges the existence of cosmic variance but does not attempt to model it, our new framework uses Bayesian statistics to fold in a comprehensive model of cosmic variance (Trapp & Furlanetto, 2020) and its effect on the galaxy population (which changes the shape of the luminosity function, see Figure 4.1). As a result, our method also measures the large-scale densities of the survey fields. We then predicted the precision of various JWST cycle-1 surveys, finding that these surveys can measure field densities to the maximum precision allowed by Poisson noise. We also found these surveys can measure the luminosity function at $z = 12$ with comparable precision to HST's existing constraints at $z = 8$, but only if the data sets can be combined effectively.

In this paper, we apply that same framework to existing HST galaxy data from Bouwens et al. (2021)[2] and Finkelstein et al. (2015) (see Table 4.1). We obtain a new measurement of the galaxy luminosity function for $z = 6$–8 and, for the first time, measure the underlying large-scale density of every HST survey field in this data set. This work also demonstrates the power of measuring field densities, with an eye forward to the much larger, deeper, and more comprehensive data set that will be arriving in cycle-1 of JWST, allowing for many more densities to be measured at higher precision.

After analyzing the existing data with our framework, we compare to earlier luminosity estimates and present the environment measurements. We also develop a new method of calculating individual field densities *after* a global luminosity function fit. This method drastically reduces computation time required to obtain field densities with a minimal loss in precision.

In section 4.3, we describe the data sets we use, briefly summarize the framework from Trapp et al. (2022), and describe the new method of measuring environments. In sections 4.4 and 4.5, we present our new measurements of the $z = 6$–8 luminosity function and field densities. In section 4.6 we discuss our results.

We use the following cosmological parameters: $\Omega_m = 0.308$, $\Omega_\Lambda = 0.692$, $\Omega_b = 0.0484$, $h = 0.678$, $\sigma_8 = 0.815$, and $n_s = 0.968$, consistent with recent Planck Collaboration XIII results (Planck Collaboration et al., 2016). We provide all distances in comoving units. We present all luminosities as rest-frame ultra-violet $(1500 - 2800\ \text{Å})$[3] luminosities, and all magnitudes are AB magnitudes.

---

[2] For $z = 6$–8, the data-set from Bouwens et al. (2021) is the same as Bouwens et al. (2015) with the addition of the COSMOS, UDS, and EGS fields.

[3] This wavelength range corresponds to H-band in the redshift range of $z \approx 5$–9 and K-band for $z \approx 8$–12.

Table 4.1: The area, magnitude limit, and source counts of each field. For fields with both Bouwens et al. (2021) and Finkelstein et al. (2015) data, the latter number count is in parentheses. Finkelstein et al. (2015) uses redshift bins of size $\Delta z = 1$ centered at $z = 6$, 7, 8. Bouwens et al. (2021) use redshift intervals of $5.5 < z < 6.3$ for their $z \sim 6$ sample, $6.3 < z < 7.3$ for their $z \sim 7$ sample, and $7.3 < z < 8.4$ for their $z \sim 8$ sample.

Field	Area	approx depth		Number	
	[arcmin^2]	[rest-UV]	z = 6	z = 7	z = 8
Bouwens (Finkelstein)					
CANDELS-GS-DEEP	64.5	27.9	198 (142)	77 (48)	26 (16)
ERS	40.5	27.8	61 (80)	46 (48)	5 (6)
CANDELS-GS-WIDE	34.2	27.4	43 (40)	5 (4)	3 (1)
CANDELS-GN-DEEP	68.3	28.1	188 (180)	134 (92)	51 (18)
CANDELS-GN-WIDE	65.4	27.4	69 (63)	39 (24)	18 (14)
HUDF/XDF	4.7	29.7	97 (94)	57 (40)	29 (15)
HUDF09-1	4.7	28.8	38 (35)	22 (16)	18 (4)
HUDF09-2	4.7	28.9	32 (31)	23 (10)	15 (3)
MACS0416-Par	4.9	29.0	25 (24)	19 (8)	4 (1)
Abell 2744-Par	4.9	28.9	20 (13)	11 (8)	4 (1)
Bouwens only					
CANDELS-UDS	151.2	26.8	33	18	6
CANDELS-COSMOS	151.9	26.8	48	15	9
CANDELS-EGS	150.7	26.9	50	43	9
MACS0717-Par	4.9	28.8	41	21	10
MACS1149-Par	4.9	28.8	36	31	6
Abell S1063-Par	4.9	28.8	40	20	7
Abell 370-Par	4.9	28.8	47	20	3

4.3 Methods and Data

4.3.1 Fitting the luminosity function

We fit a Schechter luminosity function to galaxy catalogs using a Bayesian fitting framework. This framework is described in detail in section 2.1 and 2.2 of Trapp et al. (2022), but we summarize it here.

Let us assume that the average number density of galaxies with absolute magnitudes between $(M_{\rm AB}, M_{\rm AB}+dM_{\rm AB})$ is described by $\Phi_{\rm avg}(M_{\rm AB}, z)dM_{\rm AB}$, which is a Schechter function with the following redshift-dependent parameters $\vec{\phi}(z)$: *(i)* the normalization ϕ^*, *(ii)* the characteristic magnitude M^*, and *(iii)* the faint-end power-law slope α:

$$\begin{aligned}&\Phi_{\rm avg}(M_{\rm AB}, z)dM_{\rm AB} = \\ &(0.4\ \ln 10)\phi^*[10^{0.4(M^*_{\rm AB}-M_{\rm AB})}]^{\alpha+1}\exp[-10^{0.4(M^*_{\rm AB}-M_{\rm AB})}]dM_{\rm AB}.\end{aligned} \tag{4.1}$$

The luminosity function that can actually be observed also depends on: *(i)* the effects of cosmic variance and *(ii)* observational features like the completeness and contamination functions (which we combine into a single function $f(M_{\rm AB}, z)$) that are unique to each survey volume. The luminosity function in each survey volume becomes:

$$\begin{aligned}&\Phi_{\rm obs}(M_{\rm AB}, V, z, \delta) = \\ &f(M_{\rm AB}, z)\cdot\Phi_{\rm avg}(M_{\rm AB}, z)\left(1+\frac{\delta}{\sigma_{\rm PB}}\ \varepsilon_{\rm cv}(M_{\rm AB}, V, z)\right).\end{aligned} \tag{4.2}$$

where $\delta = (\rho-\bar{\rho})/\bar{\rho}$ is the relative linear dark matter density in the volume V, $\sigma_{\rm PB}$ is the rms fluctuation of the linear dark matter density field on the scale V (and 'PB' refers to the pencil-beam shape typical of real surveys), and $\varepsilon_{\rm cv}(M_{\rm AB}, V, z)$ parameterizes the luminosity-dependent cosmic variance using the model from Trapp & Furlanetto (2020).

The cosmic variance function $\varepsilon_{\rm cv}(M_{\rm AB}, V, z)$ combines non-linear halo clustering with a self-consistent analytical galaxy model. It also corrects for the 'pencil-beam' shape of survey volumes. The largest uncertainty in $\varepsilon_{\rm cv}$ comes from the models of non-linear halo clustering; $\varepsilon_{\rm cv}$ varies ∼25% between models. The galaxy model can also affect $\varepsilon_{\rm cv}$, although to a much lesser extent because $\varepsilon_{\rm cv}$ is not a strong function of magnitude (at increasing redshift, however, $\varepsilon_{\rm cv}$'s dependence on magnitude increases). For more information on $\varepsilon_{\rm cv}$, see Trapp & Furlanetto (2020). See also Bhowmick et al. (2020) and Ucci et al. (2021) for recent simulation-based estimations of cosmic variance.

Given data $\vec{D}$ from a large suite of galaxy surveys composed of $N_{\rm f}$ fields each with their own volume, $f(M_{\rm AB}, z)$, and density δ, we would like to determine the probability density of the luminosity function parameters given the data: $p(\vec{\phi}|\vec{D})$, where $\vec{D}$ contains many galaxies with measured magnitudes and redshifts.. We are also interested the probability density of the dark matter densities of the $N_{\rm f}$ fields given the data: $p(\vec{\delta}|\vec{D})$, where $\vec{\delta}$ is a vector containing δ for each survey. Starting with the joint posterior $p(\vec{\phi}, \vec{\delta}|\vec{D})$ and applying Bayes' theorem, we have:

$$p(\vec{\phi}, \vec{\delta}|\vec{D}) \propto p(\vec{D}|\vec{\phi}, \vec{\delta}) \times p(\vec{\phi}) \times p(\vec{\delta}) \tag{4.3}$$

where $p(\vec{D}|\vec{\phi}, \vec{\delta})$ is the likelihood $\mathcal{L}$ given the average luminosity function parameters and densities, and $p(\vec{\phi})$ and $p(\vec{\delta})$ are their priors. We assume flat priors for each luminosity function parameter. The prior for each density $p(\delta_i)$ is simply a normal function centered at zero with standard deviation equal to $\sigma_{\rm PB,i}$.

From Trapp et al. (2022), the log likelihood is

$$\begin{aligned} \ln\mathcal{L} \propto \sum_{i}^{N_{\rm f}} \Bigg\{ & - n_{i,\rm exp} + \\ & \sum_{j}^{n_i} \left[\ln\Phi_{\rm avg}(M_j, \vec{\phi}) + \ln\left(1 + \frac{\delta_i}{\sigma_{\rm PB,i}} \varepsilon_{{\rm cv},i}(M_j, V_i)\right)\right] \Bigg\}, \end{aligned} \tag{4.4}$$

where the first sum is over each field, and the second sum is over each source in the $i^{\rm th}$ field. Also, $n_{i,\rm exp}$ is the number of sources *expected* in the $i^{\rm th}$ field given the average luminosity function parameters $\vec{\phi}_{\rm avg}$ and the local density δ_i. We can then write the posterior as

$$p(\vec{\phi}, \vec{\delta}|\vec{D}) \propto \mathcal{L} \times p(\vec{\delta}) \times p(\vec{\phi}). \tag{4.5}$$

Finally, we can marginalize over $\vec{\phi}$ or $\vec{\delta}$ to get $p(\vec{\delta}|\vec{D})$ or $p(\vec{\phi}|\vec{D})$, respectively.

At these high redshifts, the exponential cutoff can be poorly sampled by data. If M^* is brighter than the brightest galaxy in the sample, the data would be better fit by a single power-law. This results in an extreme degeneracy between the normalization ϕ^* and cutoff

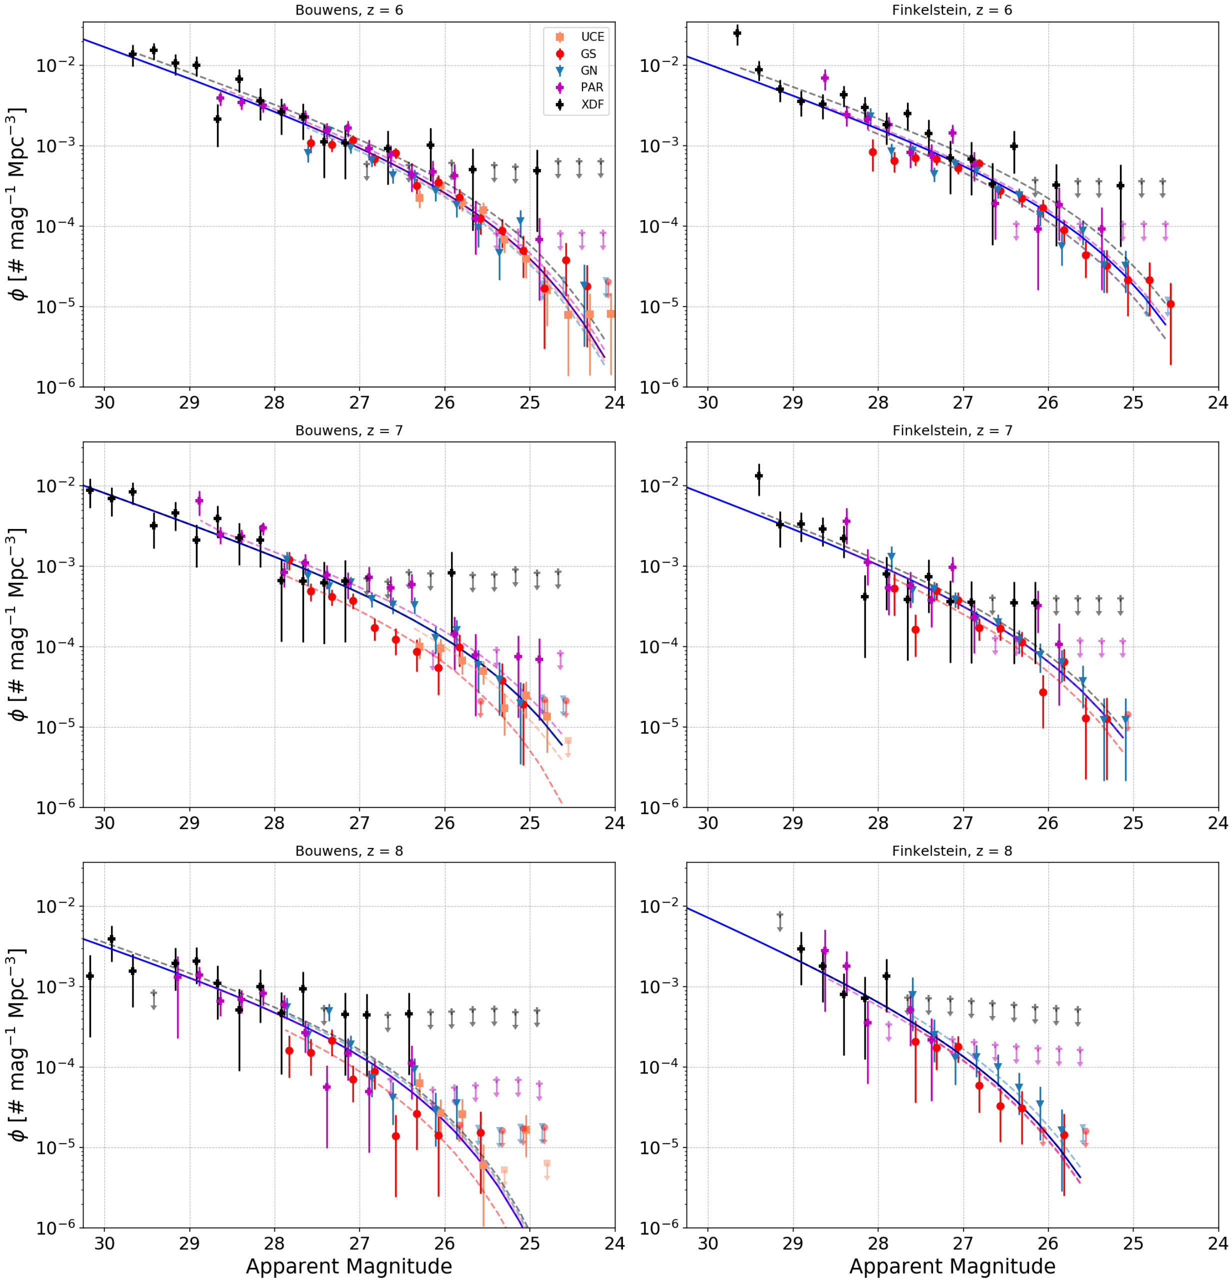

Figure 4.2: The best fit Schechter functions (*solid blue line*) for the Bouwens et al. (2021) (*left column*) and Finkelstein et al. (2015) (*right column*) data sets at $z = 6$–8 (*top-to-bottom*). The dashed lines are the best-fit local luminosity functions for each composite survey.

location M^*. To address this, we restrict the value of M^* in our fits to be fainter than the brightest galaxy in our sample.

When $\varepsilon_{\rm cv}$ becomes large, the factor $(1 + \delta/\sigma_{\rm PB} \cdot \varepsilon_{\rm cv})$ can become negative for moderate under-densities, implying a negative expectation value for the number of galaxies. This is a limitation of the Gaussian approximation to cosmic variance: in reality, we must have $\delta \geq -1$. The dark matter fluctuations on the relevant scales (of the survey fields) are much smaller than this, but the relative density of highly-biased, luminous galaxies can reach this limit. Cosmic variance's effects on a local luminosity function are better described by a distribution function in which the probability density vanishes at negative number counts but reduces to a Gaussian when $\varepsilon_{\rm cv}$ is small, like a log-normal or gamma distribution. In this paper, we use the latter, and the observed luminosity function becomes

$$\Phi_{\rm obs}(M_{\rm AB}, V, z, \delta) = f(M_{\rm AB}, z) \cdot f_\Gamma(x, k, \theta), \tag{4.6}$$

where the gamma distribution f_Γ is defined as

$$f_\Gamma(x) = \frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-x/\theta}. \tag{4.7}$$

The mean of this distribution is $k\theta$ and the variance is $k\theta^2$. We choose k and θ such that those values match the Gaussian case: $k\theta = \Phi_{\rm avg}(M, z)$ and $k\theta^2 = \Phi^2_{\rm avg}(M_{\rm AB}, z)\varepsilon^2_{\rm cv}(M_{\rm AB}, V, z)$. $\Gamma(k)$ is the gamma *function*. The variable x is chosen such that the the gamma *cumulative* distribution function at x is equal to the normal cumulative distribution function at $\delta/\sigma_{\rm PB}$. This switch also carries over to the likelihood function in the following way:

$$\ln\mathcal{L} \propto \sum_i^{N_{\rm f}} \left[-n_{i,\rm exp} + \sum_j^{n_i} \ln f_\Gamma(x, k, \theta) \right]. \tag{4.8}$$

4.3.2 Data Sets

Our analysis makes use of existing data in the redshift range $z = 6$–8. We use the public galaxy catalogs from Bouwens et al. (2021) and Finkelstein et al. (2015). Table 4.1 lists the fields used and the number counts of galaxies in those fields from each group. These catalogs contain photometric redshift and rest-frame UV magnitudes for each galaxy in their samples.

We also use the completeness and contamination functions calculated for each field by those groups. These functions are obtained by simulations and become uncertain for very faint galaxies, which can affect the results. For Finkelstein et al. (2015)/Bouwens et al. (2021), we discard all sources fainter than the magnitude at which the effective volume curve drops below 50%/33%, respectively.

The Bouwens et al. (2021) data set contains all the fields covered by Finkelstein et al. (2015), plus four additional HST parallel fields and three shallow wide field surveys[4]. Where they overlap, these two groups start with similar raw data, but use different selection criteria and reduction pipelines. We consider both final data sets in parallel to compare their results and to test the robustness of our method for calculating large-scale environments in regards to these systematic choices.

A challenge of considering the density of each field is the expanded dimensionality of the parameter space; each field introduces a new density parameter to the fit. That new parameter has a tight prior, however: a normal distribution centered at zero with standard deviation equal to $\sigma_{\mathrm{PB,i}}$, as all of these fields are large enough to be in the linear regime. Unfortunately, sampling the posterior with many sub-fields can still be costly. To alleviate this limitation, in Trapp et al. (2022) we developed a method to combine many different fields into "composite" fields with a single density parameter and an "effective" cosmic variance value, with the treatment of the combination depending on whether the fields are contiguous or treated as independent. In this work, we combine the following groups of fields into "composite" fields:

1. **GS**: CANDELS-GS-DEEP, ERS, and CANDELS-GS-WIDE are combined *contiguously.*
2. **GN**: CANDELS-GN-DEEP and CANDELS-GN-DEEP are combined *contiguously.*
3. **PAR**: HUDF09-1, HUDF09-2, MACS0416-Par, and Abell 2744-Par are combined *independently.* For fits with Bouwens et al. (2021) data, this grouping also includes

---

[4] We do not include BoRG sources (Roberts-Borsani et al., 2022) at $z = 8$, as their updated effective volume curves are not yet available.

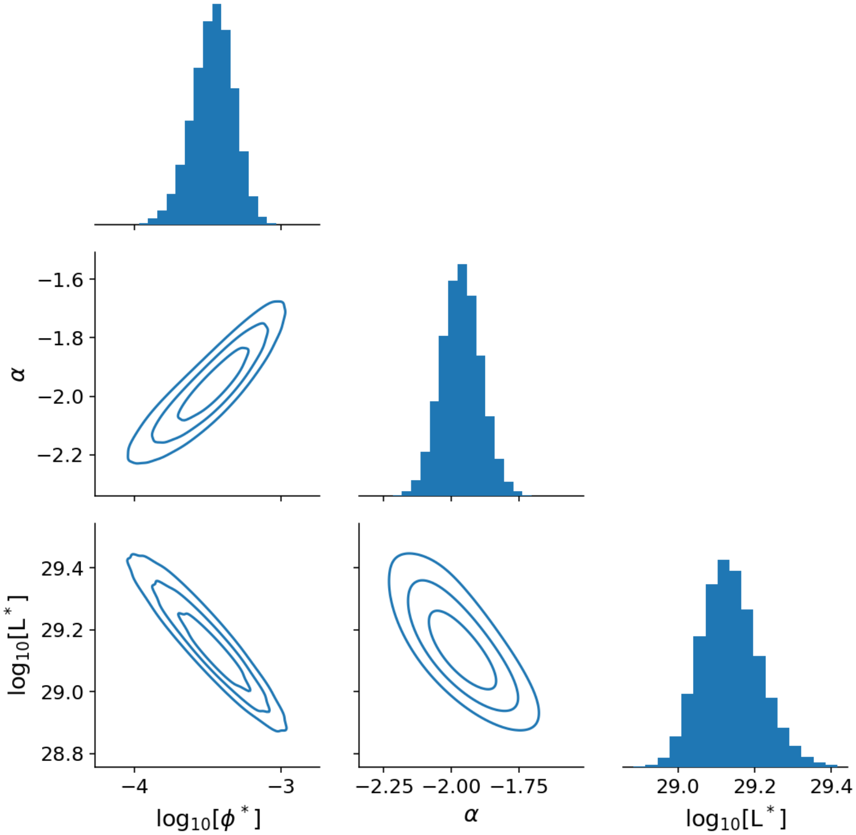

Figure 4.3: Our posterior of the Schechter parameters using the $z = 6$ Bouwens et al. (2021) data set.

MACS0717-Par, MACS1149-Par, Abell S1063-Par, and Abell 370-Par.

4. **XDF**: The HUDF/XDF field is not combined with any others.
5. **UCE**: For fits with Bouwens et al. (2021) data, the CANDELS-UDS, CANDELS-COSMOS, and CANDELS-EGS fields are combined *independently*.

We test the effects of these combinations in section 4.5.2. In short, we find that combining fields can have a non-negligible effect, but the changes are well within the current uncertainties, so the existing HST data do not demand a more intensive treatment.

4.3.3 Measuring environments

In the previous section, we combined survey fields in order to speed up the calculation of the luminosity function posterior. Unfortunately, the densities of the individual fields are lost when combining them in this way (except for HUDF/XDF, which is not combined with other fields). In this section, we introduce a new "post-processing" method to measure their densities efficiently.

We want $p(\delta_i|D_i)$ for the i^{th} field. The likelihood $\mathcal{L}_i = p(D_i|\delta_i, \vec{\phi})$ of the data is similar to equation (4.4), but applied only to one field:

$$\ln\mathcal{L}_i \propto -n_{i,\text{exp}} + \sum_{j}^{n_i} \ln f_\Gamma(x, k, \theta). \tag{4.9}$$

The posterior for the density of the ith field then becomes

$$p(\delta_i|\vec{D}) \propto p(\delta_i) \int \mathcal{L}_i \times p(\vec{\phi}) d\vec{\phi}, \tag{4.10}$$

where the prior on the Schechter parameters $p(\vec{\phi})$ is the *posterior* on those parameters that were found using section 4.3.1. Technically, the Schechter parameter prior $p(\vec{\phi})$ and the likelihood $\mathcal{L}_i$ are not completely independent, as the data from the i^{th} field was used to create $p(\vec{\phi})$. The correct thing to do is recalculate the $p(\vec{\phi})$ using all fields *except* the field for which we wish to measure the density. However, this procedure would require re-calculating $p(\vec{\phi})$ N_{f} times, which would be very computationally expensive. Further, each individual field is only a small part of the data set, having a small effect on the calculation of $p(\vec{\phi})$, making them only weakly correlated. We verify this claim in section 4.5.2.

4.4 Measurements of the Luminosity Function

We plot the Bouwens et al. (2021) and Finkelstein et al. (2015) data with our best-fit average luminosity functions and each composite field's luminosity function in Figure 4.2. Figure 4.3 is an example of the Schechter function parameter posteriors produced by our framework.

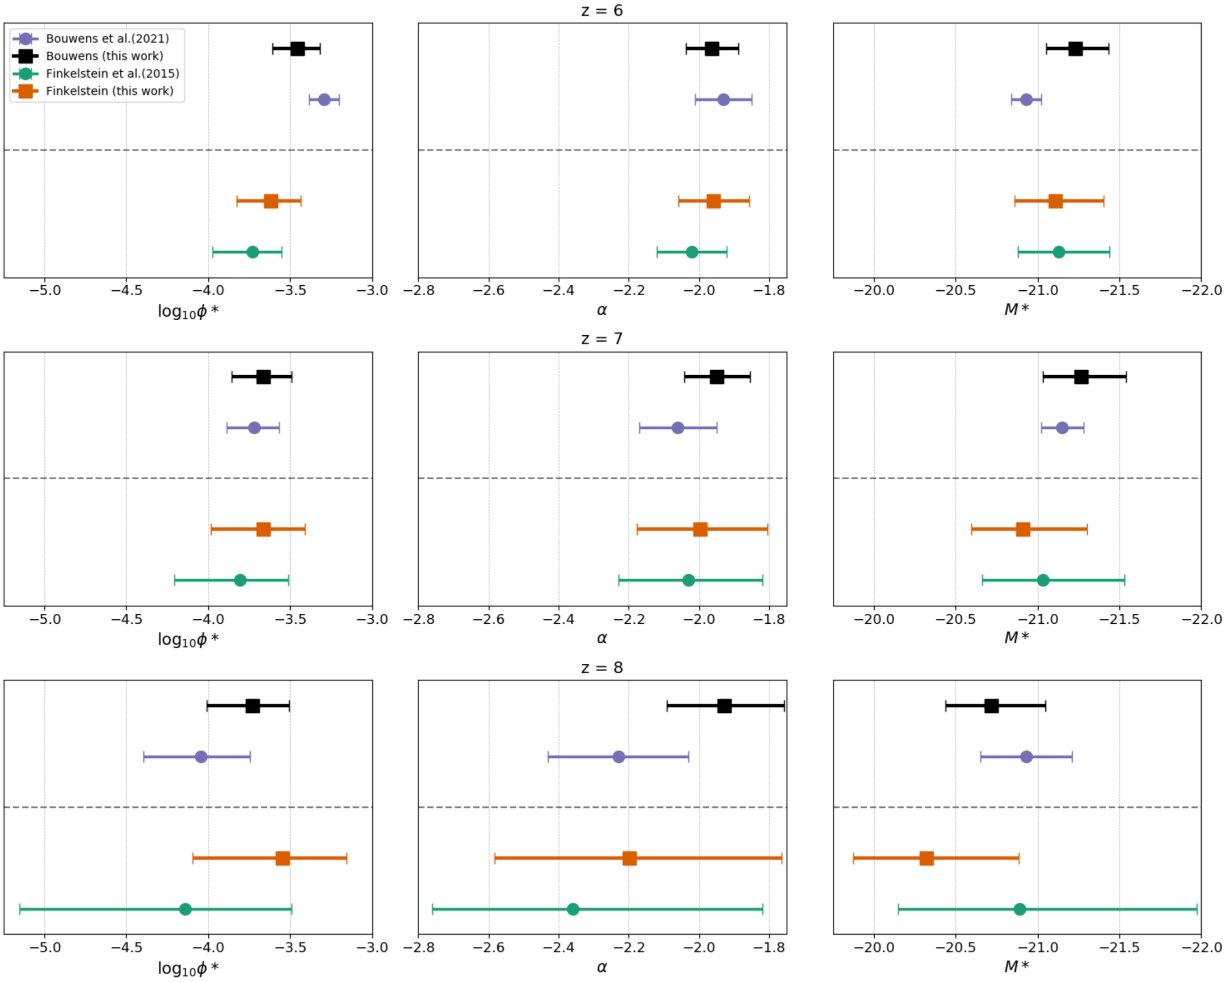

Figure 4.4: Marginalized posteriors of luminosity function parameters for $z = 6, 7, 8$ (*top, middle, bottom row*). The *black* and *orange squares* are our results using the Bouwens et al. (2021) and Finkelstein et al. (2015) data sets, respectively. The *purple* and *green* points are the results of Bouwens et al. (2021) and Finkelstein et al. (2015), respectively. These data points are available in table 4.2.

Table 4.2: Constraints on luminsosity functions from various survey combinations. All error bars are 68% confidence intervals

Redshift	Data Set	Parameter Posteriors		
		$\phi^* \times 10^4$	α	M^*
6				
	Bouwens et al. (2021)	$5.1^{+1.2}_{-1.0}$	$-1.93^{+0.08}_{-0.08}$	$-20.93^{+0.09}_{-0.09}$
	Bouwens	$3.5^{+1.3}_{-1.0}$	$-1.96^{+0.08}_{-0.07}$	$-21.23^{+0.18}_{-0.20}$
	Finkelstein et al. (2015)	$1.9^{+0.9}_{-0.8}$	$-2.02^{+0.1}_{-0.1}$	$-21.13^{+0.25}_{-0.31}$
	Finkelstein	$2.4^{+1.2}_{-0.9}$	$-1.96^{+0.10}_{-0.10}$	$-21.11^{+0.25}_{-0.29}$
7				
	Bouwens et al. (2021)	$1.9^{+0.8}_{-0.6}$	$-2.06^{+0.11}_{-0.11}$	$-21.15^{+0.13}_{-0.13}$
	Bouwens	$2.2^{+1.1}_{-0.8}$	$-1.95^{+0.10}_{-0.09}$	$-21.26^{+0.23}_{-0.28}$
	Finkelstein et al. (2015)	$1.6^{+1.5}_{-1.0}$	$-2.03^{+0.21}_{-0.20}$	$-21.03^{+0.37}_{-0.50}$
	Finkelstein	$2.2^{+1.7}_{-1.1}$	$-2.00^{+0.19}_{-0.18}$	$-20.91^{+0.32}_{-0.39}$
8				
	Bouwens et al. (2021)	$0.9^{+0.9}_{-0.5}$	$-2.23^{+0.20}_{-0.20}$	$-20.93^{+0.28}_{-0.28}$
	Bouwens	$1.9^{+1.3}_{-0.9}$	$-1.93^{+0.17}_{-0.16}$	$-20.72^{+0.28}_{-0.33}$
	Finkelstein et al. (2015)	$0.7^{+2.5}_{-0.7}$	$-2.36^{+0.54}_{-0.40}$	$-20.89^{+0.74}_{-1.08}$
	Finkelstein	$2.8^{+4.2}_{-2.0}$	$-2.20^{+0.44}_{-0.38}$	$-20.32^{+0.45}_{-0.56}$

4.4.1 Comparison of luminosity function parameters

Figure 4.4 compares the Schechter function parameter measurements using our method with the results of Finkelstein et al. (2015) and Bouwens et al. (2021). While our results agree broadly with these works, we do differ in the details. This is not surprising, as our method is more constrained when it comes to the normalization of each individual field and allows for slightly different luminosity function shapes through cosmic variance (Trapp & Furlanetto, 2020; Trapp et al., 2022). In particular, compared to Bouwens et al. (2021), our framework prefers a 1-σ lower ϕ^* and 1.5-σ lower $M^*_{\rm AB}$ at $z = 6$. At $z = 7$, our framework prefers a 1-σ higher ϕ^* and 1.5-σ shallower α. At z = 8, our framework prefers a 1-σ higher ϕ^* and 1-σ shallower α. In their work, Bouwens et al. (2021) includes a treatment of the uncertainty in measured source luminosity, an effect not considered in this work. This would have the strongest effect at the bright end of the luminosity function, and could contribute to the different findings for the highly-correlated $M^*_{\rm AB}$ and ϕ^*. Despite these differences, our best-fit luminosity function matches the total number density of sources measured by Bouwens et al. (2021) between $m_{\rm app} = 26-29$ within 10% at all redshifts. Our framework recovers the Finkelstein et al. (2015) results within 1-σ across the board. However, our best-fit luminosity function predicts 10%, 20%, and 35% more sources than Finkelstein et al. (2015)'s at $z = 6$, 7, and 8, respectively (between $m_{\rm app} = 26 - 29$).

We expect a systematic offset to occur when combining any source catalogs that use different selection and reduction techniques. For example, the results reported in Bouwens et al. (2021) and Finkelstein et al. (2015) agree with one another at the level of 1-2-σ, with Finkelstein et al. (2015) preferring a lower normalization parameter and steeper faint-end slope. At $z = 6$, 7, and 8, the best fit luminosity functions from Finkelstein et al. (2015) predict 45%, 32%, and 3% fewer sources than the best fit luminosity functions from Bouwens et al. (2021) in the range $m_{\rm app} = 26 - 29$. Similarly, at $z = 6$, 7, and 8, the best fit luminosity function from Finkelstein et al. (2015) predicts 37% less, 36% less, and 45% *more* total star formation rate density than the best fit luminosity function from Bouwens et al. (2021) (integrating down to $M_{\rm abs} = -13$). The exact reasons for these differences

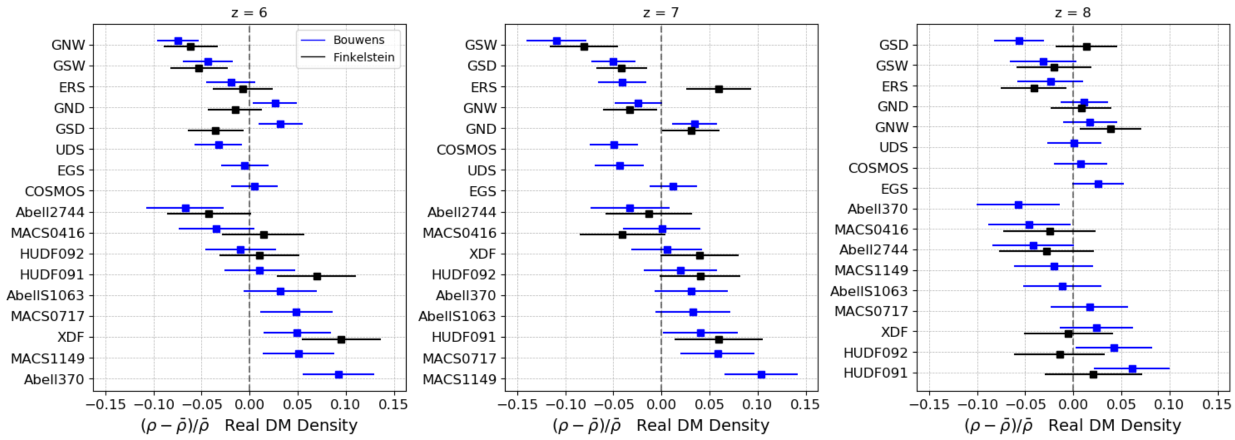

Figure 4.5: The relative overdensities of each field and their 68% confidence intervals.

are not clear, and may be a combination of one or more of the following: *(i)* differences in methodology for generating effective volume curves (which is most important at the faint-end, see Finkelstein et al., 2015; Bouwens et al., 2021, for details about these methodologies); *(ii)* differences in redshift intervals being probed, leading to sampling different physical volumes (see Table 4.1); *(iii)* groups may be sampling slightly different galaxy population due to different selection criteria; and *(iv)* other systematics. In the next section, we analyze the individual environments of each field using both groups' data sets. Discrepancies in these densities could help illuminate differences between the groups.

Some studies have found an excess of bright sources, making the data inconsistent with a Schechter function's exponential cutoff (see e.g., Bowler et al., 2014, 2015, 2020). To test this, we fit the data again using a modified Schechter function where we change the exponential factor to $e^{(L/L^{*})^{\Gamma}}$, with Γ a constant parameter. When Γ is less than one, this has the effect of flattening out the exponential cutoff, resulting in more bright galaxies. We find adding this extra parameter is disfavored given the data-set we use, increasing the Bayesian Information Criteria (BIC) by ∼5.

4.4.2 Exploring systematics

There are many subtleties that affect the measurement of the luminosity function. For example, in our framework, the choice of where to cut off the faintest sources when fitting the luminosity function can have a $\sim$1-σ effect on the resulting parameters, especially for the faint-end slope. This systematic will be immediately improved in the first cycle of JWST data, with multiple large and deep galaxy surveys like PRIMER, CEERS, JADES, PANORAMIC, WDEEP, and COSMOS-Webb. At present, differing treatments of the faint end contribute to the differences between methods shown here.

Different groups also have different selection criteria and analysis pipelines. They may therefore be probing slightly different populations of galaxies or physical locations. Empirically, this appears to have had a $\sim$1-σ effect on the luminosity function fit and a corresponding $\sim$30% difference in total number density and $\sim$40% difference in total star formation rate density (see section 4.4). This reasons behind these differences are difficult to disentangle, thanks to the many steps in galaxy observations and data reduction, and they are not something we can explore fully in this work. However, measuring the density of individual fields and comparing between groups can help illuminate differences, as we show next.

Some studies of the high-z galaxy distribution have found evidence that a Schechter function does not provide a good fit to the bright end of the luminosity function (e.g, Bowler et al. 2015). Using our method, the existing data do not justify adding another parameter to flatten out the bright end of galaxies. JWST alone will provide strong constraints on the shape of the luminosity function at the bright-end, especially at $z \sim 6$. At higher redshifts, the Roman Telescope will be crucial in measuring the bright-end of the luminosity function, especially because its data will be easily comparable to JWST due to the considerable overlap in their observable magnitudes. Wide-field ground-based surveys will also help in measuring the brightest galaxies. However, these ground-based surveys are limited in the redshifts they can reach, and when combining deep but narrow space-based images with shallow but wide ground-based images, one must consider carefully potential systematic normalization offsets between space- and ground-based measurements.

4.5 Measurements of Large-Scale Structure

4.5.1 Survey Field Environments

Figure 4.5 shows the physical dark matter densities of each Hubble field ρ relative to the average dark matter density of the Universe $\bar{\rho}$. We list the numerical results in Table 4.3. We display the results when using the Bouwens et al. (2021) data-set as well as when using the Finkelstein et al. (2015) data-set. In general, the results are consistent between groups, with a few exceptions that will be discussed below. The *normalized* relative densities (in units of standard deviations σ_{PB} from average) are also given in Table 4.3.

We convert from a dark matter over/under-density to a theoretical M^* galaxy number over/under-density using the bias value calculated using our public Python package `galcv` (Trapp & Furlanetto, 2020). Typical bias values for an M^* galaxy these surveys are 6–9. At $z = 6$, the densest Bouwens et al. (2021) field is Abell 370, with 9% more dark matter than the Universe average for that volume, corresponding to a 60% overdensity in M^* galaxies. The least dense fields are Abell 2744 and GOODS-North Wide, with 7%/7% less matter than average and 43%/48% fewer M^* galaxies. At $z = 7$, the densest field is MACS1149 with 10% more matter and 80% more M^* galaxies than average, and the least dense field is GOODS-South Wide with 11% less matter and 90% fewer M^* galaxies than average. Finally, at $z = 8$, the densest field is HUDF091 with 6% more matter and 45% more M^* galaxies than average, and the least dense field is Abell 370 with 6% less matter and 47% fewer M^* galaxies than average. At $z = 8$ however, the uncertainty is much higher in these measurements and there is more disagreement between the data-sets.

We find, in general, our measurements of the dark matter density are consistent when using data from Finkelstein et al. (2015) or Bouwens et al. (2021). However, GSD at $z =$ 6 and 8, and ERS at $z = 7$, have very discrepant density measurements between the two analyses. The exact reasons for these differences are unknown; the potential culprits are likely the same as those for the differences in the determination of the luminosity function discussed in section 4.4.1.

Table 4.3: The *real* field densities, $\delta = (\rho - \bar{\rho})/\bar{\rho}$, and the *normalized* field densities, $\delta/\sigma_{\mathrm{PB}}$ with their uncertainties.

Field & Data-set	z = 6		z = 7		z = 8	
[*B*]ouwens or [*F*]inkelstein	[norm]	[real]	[norm]	[real]	[norm]	[real]
CANDELS-GS-DEEP [*B*]	0.8±0.5	0.03±0.02	-1.2±0.6	-0.05±0.02	-1.5±0.7	-0.06±0.03
CANDELS-GS-DEEP [*F*]	-0.8±0.7	-0.04±0.03	-1.0±0.7	-0.04±0.03	0.4±0.8	0.01±0.03
ERS [*B*]	-0.4±0.6	-0.02±0.03	-0.9±0.6	-0.04±0.02	-0.6±0.8	-0.02±0.03
ERS [*F*]	-0.2±0.7	-0.01±0.03	1.4±0.8	0.06±0.03	-1.0±0.8	-0.04±0.03
CANDELS-GS-WIDE [*B*]	-0.9±0.6	-0.04±0.03	-2.5±0.7	-0.11±0.03	-0.7±0.8	-0.03±0.03
CANDELS-GS-WIDE [*F*]	-1.1±0.6	-0.05±0.03	-1.8±0.8	-0.08±0.04	-0.5±0.9	-0.02±0.04
CANDELS-GN-DEEP [*B*]	0.6±0.5	0.03±0.02	0.9±0.6	0.03±0.02	0.3±0.7	0.01±0.02
CANDELS-GN-DEEP [*F*]	-0.4±0.7	-0.02±0.03	0.8±0.8	0.03±0.03	0.2±0.8	0.01±0.03
CANDELS-GN-WIDE [*B*]	-1.8±0.5	-0.07±0.02	-0.6±0.6	-0.02±0.02	0.5±0.8	0.02±0.03
CANDELS-GN-WIDE [*F*]	-1.5±0.7	-0.06±0.03	-0.8±0.7	-0.03±0.03	1.0±0.9	0.04±0.03
HUDF/XDF [*B*]	0.8±0.6	0.05±0.03	0.1±0.7	0.01±0.04	0.5±0.7	0.02±0.04
HUDF/XDF [*F*]	1.6±0.7	0.10±0.04	0.7±0.7	0.04±0.04	-0.1±0.9	-0.00±0.05
HUDF09-1 [*B*]	0.2±0.6	0.01±0.04	0.7±0.7	0.04±0.04	1.1±0.7	0.06±0.04
HUDF09-1 [*F*]	1.2±0.7	0.07±0.04	1.1±0.8	0.06±0.05	0.4±0.9	0.02±0.05
HUDF09-2 [*B*]	-0.2±0.6	-0.01±0.04	0.4±0.7	0.02±0.04	0.8±0.7	0.04±0.04
HUDF09-2 [*F*]	0.2±0.7	0.01±0.04	0.7±0.7	0.04±0.04	-0.3±0.9	-0.01±0.05
MACS0416-Par [*B*]	-0.6±0.7	-0.03±0.04	0.0±0.7	0.00±0.04	-0.8±0.8	-0.05±0.04
MACS0416-Par [*F*]	0.2±0.7	0.01±0.04	-0.7±0.8	-0.04±0.04	-0.5±0.9	-0.02±0.05
Abell 2744-Par [*B*]	-1.1±0.7	-0.07±0.04	-0.6±0.7	-0.03±0.04	-0.8±0.8	-0.04±0.04
Abell 2744-Par [*F*]	-0.7±0.7	-0.04±0.04	-0.2±0.8	-0.01±0.04	-0.5±0.9	-0.03±0.05
CANDELS-UDS [*B*]	-0.9±0.7	-0.03±0.02	-1.3±0.7	-0.04±0.03	0.0±0.9	0.00±0.03
CANDELS-COSMOS [*B*]	0.1±0.7	0.01±0.02	-1.5±0.7	-0.05±0.02	0.2±0.9	0.01±0.03
CANDELS-EGS [*B*]	-0.1±0.7	-0.01±0.02	0.4±0.7	0.01±0.02	0.8±0.8	0.03±0.03
MACS0717-Par [*B*]	0.8±0.6	0.05±0.04	1.0±0.7	0.06±0.04	0.3±0.7	0.02±0.04
MACS1149-Par [*B*]	0.9±0.6	0.05±0.04	1.9±0.7	0.10±0.04	-0.4±0.8	-0.02±0.04
Abell S1063-Par [*B*]	0.6±0.6	0.03±0.04	0.6±0.7	0.03±0.04	-0.2±0.8	-0.01±0.04
Abell 370-Par [*B*]	1.6±0.6	0.09±0.04	0.6±0.7	0.03±0.04	-1.1±0.8	-0.06±0.04

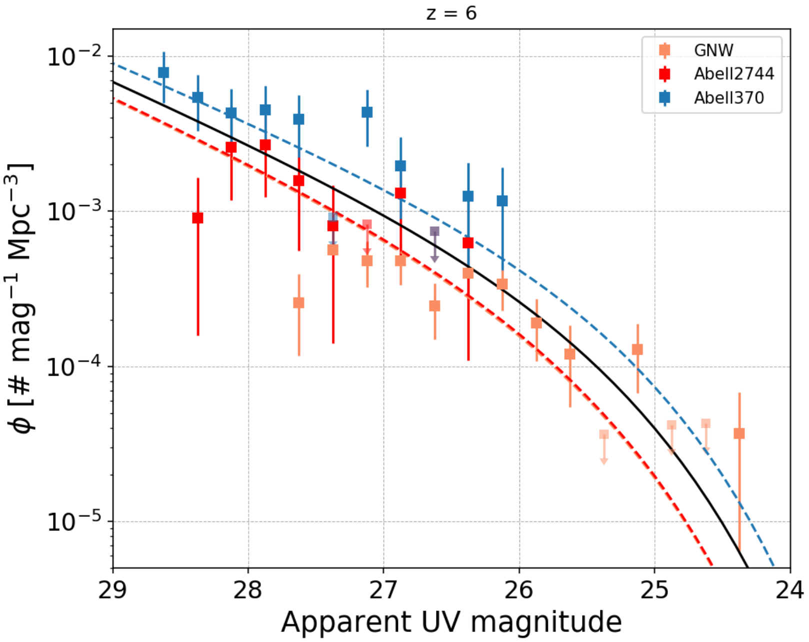

Figure 4.6: Data and best-fit luminosity functions Φ_{obs} for the three fields with the most extreme environments at $z = 6$ in the Bouwens et al. (2021) data set. The *black* curve is the best-fit average luminosity function. Abell370 and Abell2744 are clearly offset in normalization, but have very similar shapes in the magnitudes they cover. GNW is underdense (*orange-dashed* curve lies directly beneath *red-dashed* curve), but its data points differ significantly from its best-fit luminosity function due to Poisson noise on the bright and faint ends.

Finally, we find the most and least-dense fields have luminosity functions that are distinguishable by their normalization, but not by their shape (see Fig. 4.6). JWST will be able to measure individual fields to much higher precision, potentially being able to distinguish over/under-densities by the shapes of their luminosity functions.

4.5.2 Validation

4.5.2.1 Compositing surveys

In section 4.3.2, we described creating composite surveys that cut down on the time it takes to fit the luminosity function. We tested the effect of combining surveys in this way by using an alternate grouping: folding the HUDF/XDF field into the **PAR** composite group. At all redshifts, this re-grouping does not affect the width of the resulting posterior, but it does shift its position in the following way. At $z = 6$, the parameters ϕ^*, α, and M^* decrease by 0.10σ, 0.18σ, and 0.15σ. At $z = 7$, the same parameters decrease by 0.51σ, 0.76σ, and 0.48σ, respectively. Finally, at $z = 8$, the same parameters *increase* by 0.25σ, 0.47σ, and 0.30σ, respectively. This difference most likely occurs because the XDF is significantly deeper than the other **PAR** fields, so its completeness function is reasonably different from the other fields as well. Fields should only combined into composite fields if they have similarly-shaped effective volume curves and cover similar magnitude ranges, as is the case for the rest of our composite fields. Even so, the modest changes resulting from the different treatment of the XDF demonstrate that our results are robust to the details of the composite fields.

4.5.2.2 Post-processed environments

In section 4.3.3, we describe a process to measure the densities of each individual field, marginalizing over the posterior of the luminosity function $p(\vec{\phi}|\vec{D})$. Technically, this “double-counts” the fields, as each field was used to create that posterior in the first place. We test this by comparing the density of XDF generated by the full fitting framework (not double-counted) with the density of XDF using the “post-processing” described in section 4.3.3 (double-counting). The densities are the same within 0.02σ at all redshifts, showing the process is robust despite the double-counting. As described previously, the XDF has a relatively strong effect on the determination of $p(\vec{\phi}|\vec{D})$, so the double-counting effect is even smaller for less-influential fields.

4.5.2.3 Measuring cosmic variance

We can evaluate our model of cosmic variance with our results by calculating the standard deviation of all of the *normalized* density measurements ($\delta/\sigma_{\rm PB}$) weighted by the sizes of the error bars. By definition, the standard deviation of the normalized densities should equal unity. If our model for cosmic variance gave values that were globally too large, we would expect the standard deviation of our measured normalized densities to be less than one, or greater than one if our model for cosmic variance was too small. The standard deviation of the measured normalized densities is 0.92 ± 0.09. Therefore, the data do not prefer globally larger or smaller cosmic variance model. Doing the same calculation but splitting out into the different redshifts gives us standard deviation values of 0.90 ± 0.16, 1.06 ± 0.19, and 0.74 ± 0.13 for redshifts 6, 7, and 8, respectively. Redshifts 6 and 7 do not prefer larger or smaller values of cosmic variance, but redshift 8 prefers a 2-σ smaller value for cosmic variance.

4.6 Conclusions

We measure the universal luminosity function of galaxies for $z = 6$–8 using existing data Finkelstein et al. (2015); Bouwens et al. (2021). We use a new fitting method that is more constrained than existing methods and also incorporates the density of individual fields or composited groups of fields (Trapp et al., 2022). Our results are consistent with existing studies (Finkelstein et al., 2015; Bouwens et al., 2021), but differ in the details (see Fig. 4.4). Our method has the benefit of considering the shape change of the luminosity function for different densities, an effect that will become more pronounced at higher redshift (Trapp & Furlanetto, 2020).

We measure the dark matter density of most deep Hubble galaxy fields from $z = 6$–8. We find the least/most-dense Hubble deep fields at $z = 6$, 7, 8 are GNW/Abell370, GSW/-MACS1149, and Abell370/HUDF091, respectively. These fields have expected dearths/excesses of M^* galaxies of -48%/60%, -90%/80%, and -47%/45%, respectively. We find dark matter densities are distributed in a way that is consistent with current estimations of cosmic variance (Trapp & Furlanetto, 2020; Bhowmick et al., 2020; Ucci et al., 2021). JWST will obtain many more dark-matter measurements of survey fields and at a higher precision than currently possible. These densities can be sorted and used to compare many statistical aspects of galaxies in under/over-dense environments, from the shape of the luminosity function, to the star-formation histories of galaxies, to the number of LAEs or QSOs in a region.

The pencil-beam shape of these volumes make interpreting a *high* density complicated, as galaxies are likely clustered radially within the pencil-beam. If galaxies can be sorted more precisely in redshift space (through accurate photometric redshifts, for example), it will be possible to make field density estimates on smaller, more spherical volumes, which will allow for closer comparison to observations. In particular, our method could be improved to incorporate the probability distributions of photometric redshifts (and luminosities) as measurements improve.

We do not analyze the density sub-structure of fields, something that would be especially useful for large contiguous fields like COSMOS-Webb. This is in-principle doable, as we will have 3-D positions of each source in each field. We do not attempt this here because of the substantial radial widths of the HST fields, but it should be possible with JWST.

CHAPTER 5

Lyman-α Emitters in Ionized Bubbles: Constraining Environment and Ionization Fraction

5.1 Abstract

Lyman-alpha emitters (LAEs) are excellent probes of early ionized regions, as they must be surrounded by large ionized bubbles in order to be visible. Large ionized regions are thought to be very over-dense, and may be (or contain) protoclusters, making them interesting test-beds for early massive structures. For these reasons, close associations containing several LAEs are often assumed to mark ionized bubbles. Here, we develop the first framework to quantify the ionization and density fields of high-z galaxy associations, and we study how these galaxies can be used to learn about the progress of reionization. We explore the interplay between *(i)* the large-scale density of a survey field, *(ii)* Poisson noise due to the small number density of sources at high redshifts ($z \sim 7$), and *(iii)* the effects of the ionized fraction on the observation of LAEs. We use Bayesian statistics, a simple model of re-ionization, and a Monte-Carlo simulation to construct a more comprehensive method for calculating the large-scale density of LAE regions than previous works. We find that Poisson noise has a strong effect on the inferred density of a region if there are a small number of sources used to calculate the density and show how the ionized fraction can be inferred. We then apply our framework to the strongest association yet identified: Hu et al. (2021) found 14 LAEs in a volume of ~50,000 cMpc^3 inside the COSMOS field at $z \sim 7$. We show that this is most likely a 2σ over-density (despite there being a 5× over-density in the number of visible LAEs in that region) inside of an ionized or nearly ionized bubble. We also show that this LAE association implies that the global ionized fraction is $\bar{Q} = 0.70^{+0.14}_{-0.09}$, within

the context of our reionization model. Our approach provides the first step in quantifying the implications of LAE association for reionization.

5.2 Introduction

The Cosmic Dawn is about to be explored as never before, and a spotlight shines on reionization. The search for rare, large ionized regions and the first large-scale objects like protoclusters is of particular interest (Ouchi et al., 2005; Calvi et al., 2019; Tilvi et al., 2020; Jung et al., 2020; Hu et al., 2021; Endsley et al., 2021). These extreme regions may serve as test-beds for the physics of early galaxy and cluster formation. By measuring the large-scale densities of these regions, we can learn about the assembly history of galaxy clusters and other large, rare objects (e.g. Chiang et al., 2017). Also, feedback on the scale of these extreme regions, driven by underlying large-scale structure, is likely to affect the galaxy population strongly during re-ionization (Thoul & Weinberg, 1996; Iliev et al., 2007; Noh & McQuinn, 2014).

Lyman-alpha emitters (LAEs) are an exciting probe of reionization (as well as galaxy formation). The young galaxies present in the reionization era are likely to have large intrinsic Lyman-α luminosities (Partridge & Peebles, 1967; Ouchi et al., 2020), but those photons are subject to substantial absorption by the intergalactic medium (IGM) before reionization is complete, thanks to the enormous optical depth of remaining neutral islands (Gunn & Peterson, 1965). We therefore expect to see a decline in the abundance of LAEs as we penetrate further into the reionization era, making these galaxies an effective probe of the ionization state (Madau & Rees, 2000; Haiman, 2002; Santos, 2004).

The number counts of LAEs evolve rapidly above $z \sim 5$, which may in part be attributable to reionization (Malhotra & Rhoads, 2004; Kashikawa et al., 2011; Itoh et al., 2018). However, the evolution of the overall number density of LAEs is difficult to disentangle from the overall evolution of the galaxy abundances. For that reason, attention has shifted (when possible) to more sophisticated ways to leverage LAEs. One possibility is to study the fraction of photometrically-selected galaxies with Lyman-α lines (Pentericci et al., 2011; Stark et al., 2011; Ono et al., 2012). These studies have typically found that the Universe must have a substantial neutral fraction at $z \sim 7$, though the inferences depend on the reionization model (Mason et al., 2018).

Another exciting prospect is that the inhomogeneous ionization field will modify the spatial distribution of LAEs. In order for their Lyman-α lines to survive, the host galaxy must be embedded in a large ionized region, which, in most reionization models, corresponds to an over-dense region (Furlanetto et al., 2004; Wyithe & Loeb, 2005). Galaxies outside of such large regions may produce Lyman-α photons, but they will not escape. Thus the ionization field modulates the clustering of LAEs (Furlanetto et al., 2006; McQuinn et al., 2007). Unfortunately, these clustering measurements are challenging and require a large number of sources (e.g., Yoshioka et al. 2022).

A particularly interesting aspect of this modulation is that it exaggerates the existing clustering: galaxies in over-dense regions (which host large ionized bubbles) will remain visible, while even galaxies inside modest over-densities will disappear early enough in reionization. This suggests that focusing on identifying rare ionized regions can be a powerful probe of reionization (e.g., Mesinger & Furlanetto 2008).

Meanwhile, surveys for LAEs, typically with narrowband filters, are now common, with large enough areas that unusual regions can be found. Of particular interest for us, several teams have discovered candidate ionized bubbles hosting apparently significant over-densities of LAEs (e.g., Tilvi et al. 2020; Jung et al. 2020; Hu et al. 2021; Endsley et al. 2021).

Perhaps the most compelling such region has recently been identified by Hu et al. (2021), who found 14 LAEs in a region with volume $V \sim 50,000$ cMpc^3 at $z = 6.93$ (~ 200 pMpc^3), across a larger field with approximately 20 volumes of the same size. The expected number of LAEs to be found in this volume is just $N_{\rm exp,avg} = 2.9 \pm 1.1$. This is thus a clear over-density in the LAE counts, and hence very likely a large ionized bubble.

However, to date such inferences are purely qualitative, with only simple efforts to transform observed features in the galaxy distribution to a quantitative constraint on the underlying ionization and density fields.

Finding and analysing these regions, and then connecting the visible sources to the total population of galaxies that are ionizing the Universe will be challenging on many fronts. The bulk of ionizing photons are coming from galaxies that we cannot see with HST *nor* JWST (see e.g., Behroozi & Silk, 2015; Furlanetto et al., 2017). In a partially-ionized Universe, even strong Lyα lines can be obscured by intervening neutral Hydrogen, leaving only a small number of the brightest LAEs in the most ionized regions visible. The observed number density of these LAEs is then strongly affected by Poisson noise. Poisson noise is uncorrelated between magnitude bins, meaning an over-density of a few bright LAEs does not necessarily mean an over-density in the much-larger underlying population of faint or obscured sources. In other words, determining the large-scale density of a region from a small number of LAEs requires a simultaneous treatment of *(i)* cosmic variance and the corresponding galaxy bias of those LAEs, *(ii)* Poisson noise, and *(iii)* the effects of a partially ionized Universe on the visibility of LAEs.

The first two of these points are more general than LAEs: efforts to find “protoclusters" or other unusual environments from galaxy distributions suffer from Poisson noise as well. Such efforts to identify protoclusters are useful for tracing the history of the most massive structures in the Universe today (Ouchi et al., 2005; Trenti et al., 2012; Chiang et al., 2017; Calvi et al., 2019; Hu et al., 2021), but the associations have also been largely qualitative.

In this paper, we develop the first quantitative, statistically robust framework to infer the underlying density and ionization environment of observed galaxy associations during the reionization era. This framework can be applied both to LAEs (which probe both the ionization state and the density) and to other surveys (which probe only the density).

In section 5.3, we construct an analytic form of the posterior for the large-scale density of a region, ignoring any effects of inhomogeneous reionization. We apply this method to the Hu et al. (2021) region in section 5.4. We then expand our method to the partially-ionized case using a simple model of re-ionization and a Monte-Carlo (MC) simulation in section 5.5. This MC simulation also has the ability to constrain the ionization fraction of the Universe (in the context of our simple model of re-ionization). In section 5.6, we apply the MC simulation method to the same region from Hu et al. (2021), obtaining a new measurement of its density and local ionization field as well as a constraint on the ionization fraction of the Universe.

We use the following cosmological parameters: $\Omega_m = 0.308$, $\Omega_\Lambda = 0.692$, $\Omega_b = 0.0484$, $h = 0.678$, $\sigma_8 = 0.815$, and $n_s = 0.968$, consistent with recent Planck Collaboration XIII results (Planck Collaboration et al., 2016).

5.3 Inference of the local density field from galaxy associations

We begin by ignoring reionization and just imagining inferring the underlying dark matter density of a field with some set of observed sources. Take the average number density of bright LAEs at a given redshift to be $n_{\rm avg}$. In a volume V, one would expect to find $N_{\rm \,exp,avg} = V \cdot n_{\rm avg}$ bright LAEs. Now, assume one completes an observational campaign of size $N_{\rm vols} \times V$, finding that the most dense region of size V has $N_{\rm obs} > N_{\rm \,exp,avg}$. What can we infer about that region? Does the ratio of $N_{\rm obs}/N_{\rm \,exp,avg}$ carry through to those sources below the magnitude limit of the survey? Can we measure the excess amount of dark matter in this region using a bias function and infer whether the region will collapse into a cluster by $z = 0$? To answer these questions, one must consider Poisson noise and cosmic variance jointly.

In the gaussian approximation of cosmic variance, the expected number of sources in some region V with *linearized* relative density $\delta = (\rho - \bar{\rho})/\bar{\rho}$, is

$$N_{\rm exp} = N_{\rm \,exp,avg}(1 + \delta \cdot b(V)), \tag{5.1}$$

where $b(V)$ is the *bias* of those sources[1] (Mo & White, 1996). However, when one observes a region of density δ and corresponding $N_{\rm exp}$, one does not always find $N_{\rm obs} = N_{\rm exp}$. The observed number is drawn from the Poisson distribution with $\lambda = N_{\rm exp}$; explicitly,

$$P(N_{\rm obs}|\lambda = N_{\rm exp}) = \frac{N_{\rm exp}^{N_{\rm obs}} e^{-N_{\rm exp}}}{N_{\rm obs}!}. \tag{5.2}$$

---

[1] The bias also typically depends on the luminosity or mass of the objects being considered. However, for a set population defined by a magnitude range, there is an effective bias value for that population.

One can infer the value of $N_{\rm exp}$ given $N_{\rm obs}$ using Bayes' theorem,

$$p(N_{\rm exp}|N_{\rm obs}) \propto P(N_{\rm obs}|\lambda = N_{\rm exp})p(\lambda = N_{\rm exp}), \tag{5.3}$$

where $p(\lambda = N_{\rm exp})$ is the prior on $N_{\rm exp}$. From equation (5.1) and the fact that the cosmological density field is a Gaussian (at least in the linear approximation appropriate on large scales), that prior is a Gaussian centered at $N_{\rm \ exp,avg}$ with standard deviation $\sigma_{N_{\rm exp}} = N_{\rm \ exp,avg} \cdot b \cdot \sigma_\delta$,

$$p(N_{\rm exp}) = \frac{1}{\sigma_{N_{\rm exp}}\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{N_{\rm exp} - N_{\rm \ exp,avg}}{\sigma_{N_{\rm exp}}}\right)^2\right], \tag{5.4}$$

where σ_δ is the r.m.s. fluctuation of δ for the geometry matching the survey volume (see Newman & Davis, 2002; Stark et al., 2007; Muñoz et al., 2010; Robertson, 2010; Trapp & Furlanetto, 2020, for calculations of non-spherical density fluctuations). Multiplying equations (5.2) and (5.4), taking the natural log, and dropping terms that do not depend on $N_{\rm exp}$ gives the log likelihood. Changing the inferred quantity to the underlying dark matter density δ via equation (5.1), we find

$$\ln p(\delta|N_{\rm obs}) \propto N_{\rm obs}\ln(1 + \delta \cdot b) - N_{\rm \ exp,avg} \cdot \delta \cdot b - \frac{1}{2}\left(\frac{\delta}{\sigma_\delta}\right)^2. \tag{5.5}$$

The above equation gives the full posterior of the region's dark matter density given an observed number of sources $N_{\rm obs}$, an average expected number of sources $N_{\rm \ exp,avg}$, a bias value b for those sources, and knowledge of the r.m.s. fluctuation in the linear dark matter density field σ_δ.

If there is uncertainty in one of those parameters – say the bias value b – its probability distribution $p(b)$ can be marginalized over in the following way:

$$p(\delta|N_{\rm obs}) = \int p(\delta, b|N_{\rm obs})db = \int p(\delta|b, N_{\rm obs})p(b)db. \tag{5.6}$$

So far, we have imagined a simple experiment in which only the total number of sources is known. We will restrict ourselves to this simple case in this paper, but it is easy to extend the formalism to more sophisticated experiments. In a more ideal case, the average luminosity function of these sources $\Phi_{\rm avg}$ is known, with corresponding local luminosity function $\Phi_{\rm loc}(M) = \Phi_{\rm avg}(M)(1 + \delta \cdot b(M, V))$, where M is the absolute magnitude of the sources. In this case, with $N_{\rm obs}$ observed sources, each with M_j, the posterior becomes

$$\begin{aligned} &\ln p(\delta|N_{\rm obs}) \propto -N_{\rm \,exp,avg} \cdot \delta \cdot b - \frac{1}{2}\left(\frac{\delta}{\sigma_\delta}\right)^2 + \\ &\sum_j^{N_{\rm obs}} \Big[\ln\Phi_{\rm avg}(M_j) + \ln\left[1 + \delta \cdot b(M_j, V)\right]\Big]. \end{aligned} \tag{5.7}$$

5.4 Applying the density inference framework

In Hu et al. (2021), 14 LAEs were found in a region with volume $V = 66\times30\times26$ cMpc^3 at $z = 6.93$ (~ 200 pMpc^3), embedded within a larger survey field with approximately 20 volumes of the same size. Henceforth, for convenience we will refer to this specific over-dense volume of 14 sources as the "Hu et al. (2021) volume", and in general we will call a group of LAEs found at high-z an "association". The expected number of LAEs to be found in this volume – using the number density in the larger COSMOS field – is $N_{\rm \,exp,avg} = 2.9 \pm 1.1$ (boosted to account for redshift-space distortions, see below). Hu et al. (2021) calculates a galaxy over-density of $\delta_g = 5.11^{+2.06}_{-1.70}$ for the region with 14 sources. Combining this with a bias value of $b_{\rm real} = 4.54 \pm 0.63$ and a redshift-space distortion correction $C = 0.79$ via the equation $1+b_{\rm real}\delta_{\rm real} = C(1+\delta_g)$, they find the dark-matter over-density to be $\delta_{\rm real} = 0.87$. Converting this real over-density to a linear over-density via Mo & White (1996) gives $\delta = 0.54$. Given the r.m.s. fluctuation in the linear density field for the Hu et al. (2021) volume σ_R =0.074, this corresponds to a 7.3 σ_R over-density, a very large excursion.

We adopt the bias value used in Hu et al. (2021), though convert it to a linear bias value $b = 7.31 \pm 1.02$ via $(1 + \delta_{\rm real} \cdot b_{\rm real}) = (1 + \delta \cdot b)$. With these parameters, using equation (5.5) and marginalizing over the uncertainty in b and $N_{\rm \,exp,avg}$, we plot the posterior of the density of this region in Figure 5.1 in *magenta*, finding a density and 68.27% confidence interval of $\delta = 0.16 \pm 0.05$, much lower than estimated in Hu et al. (2021) (shown by the vertical dashed *orange* line). Given the r.m.s. density fluctuation in the linear density field for the Hu et al. (2021) volume σ_R =0.074, this corresponds to a 2.2 ± 0.7 σ_R over-density. The Figure also shows the likelihood in *red* (eq. 5.5 without the final term, still marginalized over

b and $N_{\text{exp,avg}}$) and the prior on the density in *blue*.

Our posterior is prior-dominated, as the likelihood is much broader than the prior, even with 14 observed sources. The vertical dashed *black* line indicates the approximate density required for this region to collapse into a cluster by $z = 0$, $\delta_{\text{pc}} = 0.27$ (assuming the spherical collapse model), a situation that is disfavored by the posterior. If we do not marginalize over the uncertainty in b and $N_{\text{exp,avg}}$, the posterior changes very slightly, still giving $\delta = 0.16 \pm 0.05$. This is not surprising, given we are in the prior-dominated regime. If we significantly change the bias value to $b = 5.0 \pm 0.7$ – the bias for a $\sim 4 \times 10^{11} M_\odot$ halo – the density inference still shifts very little, to $\delta = 0.15 \pm 0.06$

In Figure 5.1, the likelihood and prior are very far apart, indicating that this volume is a rare find. In fact, the probability of a single region of this volume to host 14 *or more* sources given the estimates of $N_{\text{exp,avg}}$, b, σ_R (including the uncertainties in these quantities), and Poisson noise, is 0.3%. The probability of finding at least one such region in a larger survey of 20 similar volumes is 5%, about a $\sim 2\sigma$ excursion. This is surprising enough that one might ask if another factor might change the probability – such as a patchwork of ionized bubbles modulating the LAE surface density.

Using the best-fit luminosity function from Trapp & Furlanetto 2022b (in prep.) and the UV galaxy bias function from Trapp & Furlanetto (2020), integrating down to $M_{\text{UV}} = -13$, we estimate that this region has an excess in UV light production of $40 \pm 13\%$ over average.

5.5 Inference of the ionization field from LAE associations

The last section assumed that *no* LAEs in the COSMOS field were blocked by intervening neutral hydrogen, or in other words that the Universe was completely ionized. In reality, at $z = 6.93$, the Universe is thought to be only partially ionized (e.g., Davies et al. 2018). In this case, we would expect only a fraction $\bar{F}_{\text{surv}}$ of LAEs to actually be visible. Thus, we would expect the true underlying $N_{\text{exp,avg}}$ to actually be *larger* than the visible density of LAEs, $N_{\text{exp,avg}} = N_{\text{exp,avg,vis}}/\bar{F}_{\text{surv}}$ with $N_{\text{exp,avg,vis}} = 2.9 \pm 1.1$. By itself, this would have the effect of shifting the density posterior of the Hu et al. (2021) volume to smaller over-densities.

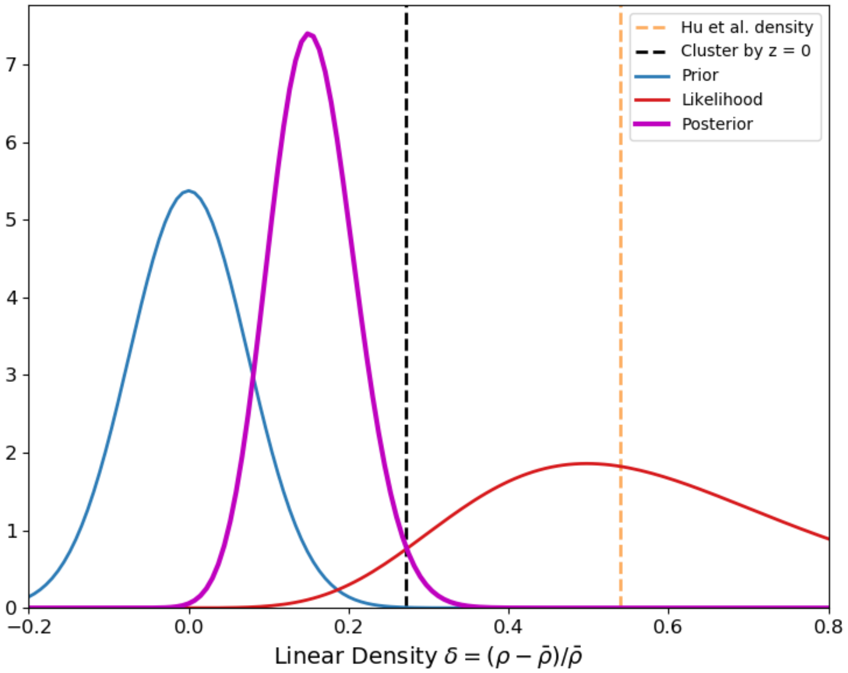

Figure 5.1: The inferred density for the Hu et al. (2021) volume, assuming no effect from reionization. The *magenta curve* shows posterior of the linear density, which has 68% confidence interval $\delta = 0.16 \pm 0.05$. The vertical dashed *black* line indicates the approximate density required for this region to collapse into a cluster by $z = 0$, $\delta_{\mathrm{pc}} = 0.27$, which is disfavored by the posterior. The vertical dashed *orange* line shows the density estimate of Hu et al. (2021). The *red and blue* curves show the likelihood and prior, respectively; we are in the prior-dominated regime, given the likelihood's large breadth. The likelihood and prior are very far apart, indicating that this volume is a rare find; we show in the text that the probability of finding at least one such region in a larger survey of 20 similar volumes (the COSMOS field) is 5%.

However, in this context a region with an over-density of visible LAEs must also be mostly ionized in order to see the sources inside it. In the standard picture of reionization (e.g., Furlanetto et al. 2004) it must have a *high* density in order to be mostly ionized. More generally, in a fully ionized Universe, the clustering of LAEs is due to cosmic variance and Poisson noise alone. In a mostly neutral Universe, ($\bar{Q} \ll 1$), we would expect to find nearly all of the visible LAEs (if indeed any could be found!) in a small number of large ionized regions – or, in other words, more highly clustered. We found in section 5.4 that the volume found in Hu et al. (2021) is highly clustered when compared to its surroundings, so that the probability of finding such a region among 20 same-sized regions in a fully ionized Universe is only 5%. If ionized regions exaggerate the clustering, there must be a value for $\bar{Q} < 1$ that is most likely to produce a comparable LAE association once within a larger survey volume.

Incorporating these considerations into the density posterior requires a mapping between between $\bar{Q}$ and $\bar{F}_{\rm surv}$ (sec. 5.5.1) and a treatment of $F_{\rm surv}$ for an individual region the size of the Hu et al. (2021) volume (sec. 5.5.2).

5.5.1 LAEs in a simple reionization model

For a mapping between $\bar{Q}$ and $\bar{F}_{\rm surv}$, one first needs a model of reionization. In this section, we construct a very simple such model. Our prescription can be made more rigorous by comparing to more detailed reionization models, such as those generated by `21cmFAST` (Mesinger et al., 2011; Davies et al., 2021), but we focus here on a very simple prescription to make the inference framework as transparent as possible. This model was inspired by the measurement of a short mean free path for ionizing photons, $\lambda_i = 0.75^{+0.65}_{-0.45}$ pMpc at $z = 6.0$ (Becker et al., 2021).

Let us assume that for an LAE to be visible, it must be inside an ionized bubble with $R > R_\alpha$ where $R_\alpha = 1$ pMpc, allowing for its Lyα photons to redshift out of resonance (Miralda-Escudé, 1998). Let us further assume that the ionizing photon mean free path is smaller than R_α. This means that for an ionized bubble to grow large enough to allow transmission of Lyα photons, all of the ionizing photons that generate the bubble have to come from sources inside of it. Also, every region of size R_α is independent of its neighbors. In reality, some high-frequency photons will be shared between regions, but for the sake of simplicity we ignore them here.

We make another simplifying assumption that the ionized fraction of hydrogen Q in an independent region depends linearly on the fraction of baryons that have collapsed into haloes $f_{\rm coll}$ via an efficiency parameter ζ:

$$Q = \zeta \cdot f_{\rm coll}. \tag{5.8}$$

This expression implicitly assumes that the ionizing efficiency of galaxies is independent of their mass. This is very unlikely to be true of real galaxies (e.g., Trenti et al. 2010; Tacchella et al. 2013; Mason et al. 2015; Behroozi & Silk 2015; Furlanetto et al. 2017), but it allows for a very simple reionization model. In particular, within the Press-Schechter model (Press & Schechter, 1974; Lacey & Cole, 1993),

$$f_{\rm coll}(\delta, R_\alpha, z) = {\rm erfc}\left(\frac{\delta_{\rm crit}(z) - \delta_0}{\sqrt{2(\sigma^2_{\rm min} - \sigma^2_{R_\alpha})}}\right), \tag{5.9}$$

where $\delta_{\rm crit}(z)$ is the linearized density threshold for spherical collapse (approximately 1.69 divided by the growth factor of dark matter structure, Eisenstein & Hu, 1998), δ_0 is the linearized density of the region δ scaled to $z = 0$ (again, via the growth factor), σ_{R_α} is the linear r.m.s fluctuation of the dark matter density field on the scale of R_α, and $\sigma_{\rm min}$ is the same on the scale of the smallest virialized halo allowed to form a galaxy. We take that smallest scale to correspond to a halo virial temperature $T_{vir} = 10^4 K$, when atomic line cooling becomes efficient enough for gas clouds to collapse and fragment for star formation (Loeb & Furlanetto, 2013).

In this simple model in which the ionized fraction increases monotonically with the local collapse fraction, a region is fully ionized if it has a sufficiently high density. Setting $Q = 1$, we obtain a relationship between the efficiency ζ and $f_{\rm coll,vis}$, the collapse fraction required to ionize a region of size R_α: $\zeta = 1/f_{\rm coll,vis}$ (see Fig. 5.2). From equation (5.9), we can also obtain $p(f_{\rm coll}|z, R)$.

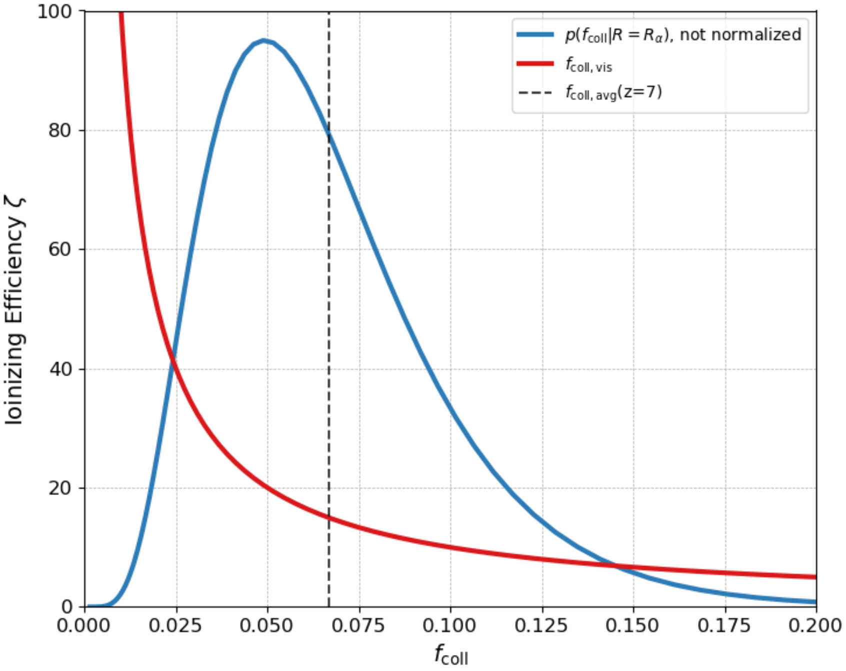

Figure 5.2: Properties of our simple reionization model. The *red* curve shows the ionizing efficiency required to ionize an independent region as a function of its collapse fraction (see eq. 5.8). The *blue* curve shows $p(f_{\rm coll}|R = R_\alpha)$, the (un-normalized) probability distribution of collapse fractions for regions of size $R = R_\alpha$. The *black dashed* curve shows the average collapse fraction at $z = 7$.

Choosing a value for $f_{\mathrm{coll,vis}}$ then defines the fraction of equal-mass regions at each redshift that can contain visible LAEs:

$$\bar{F}_{\mathrm{vis}}(z) = \int_{f_{\mathrm{coll,vis}}}^{1} p(f_{\mathrm{coll}}|z, R_\alpha) df_{\mathrm{coll}}. \tag{5.10}$$

Weighting by the number of LAEs inside a region of f_{coll} gives the fraction of LAEs in that are visible at each redshift:

$$\bar{F}_{\mathrm{surv}}(z) = \frac{1}{N_{\mathrm{exp,avg}}} \int_{f_{\mathrm{coll,vis}}}^{1} N_{\mathrm{exp}}(f_{\mathrm{coll}}) \cdot p(f_{\mathrm{coll}}|z, R_\alpha) df_{\mathrm{coll}}, \tag{5.11}$$

where $N_{\mathrm{exp,avg}}$ is the average number of LAEs in a region of size R_α and $N_{\mathrm{exp}}(f_{\mathrm{coll}})$ is the number of LAEs in a region of size R_α but with collapse fraction f_{coll}. Since each collapse fraction has a unique δ value on this scale, $N_{\mathrm{exp}}(f_{\mathrm{coll}}(\delta)) = N_{\mathrm{exp,avg}}(1 + \delta \cdot b)$ (like eq. 5.1). Similarly, we can weight by the ionization of each region through $Q = f_{\mathrm{coll}}/f_{\mathrm{coll,vis}}$ (taking $Q = 1$ for $f_{\mathrm{coll}} > f_{\mathrm{coll,vis}}$) to obtain the average ionization of the Universe at some redshift:

$$\bar{Q}(z) = \int_{0}^{1} \frac{f_{\mathrm{coll}}}{f_{\mathrm{coll,vis}}} \cdot p(f_{\mathrm{coll}}|z, R_\alpha) df_{\mathrm{coll}}. \tag{5.12}$$

Note that $\bar{Q}$ is *not* simply the globally-averaged collapse fraction multiplied by ζ, because some photons in over-dense regions are wasted thanks to absorption by small-scale features: this is important because we have assumed λ_i to be smaller than the size of the independent regions.

Finally, we obtain a relationship between the average ionized fraction of the Universe $\bar{Q}(z)$ and the fraction of LAEs that are visible $\bar{F}_{\mathrm{surv}}(z)$ (averaged across the Universe) for any choice of $f_{\mathrm{coll,vis}}$ (as shown in Fig. 5.3). Higher bias values result in a larger fraction of all LAEs being visible for a given ionization fraction of the Universe. We also plot in Figure 5.3 the total ionizing photon production rate relative to the number of hydrogen atoms, $\zeta \cdot f_{\mathrm{coll,avg}}$. A value above the line $x = y$ implicitly means photons are escaping their host galaxies, but being absorbed before contributing to reionization. The most over-dense R_α regions are overproducing photons but can't help their neighbors reionize. This effect is only significant once the Universe is mostly reionized.

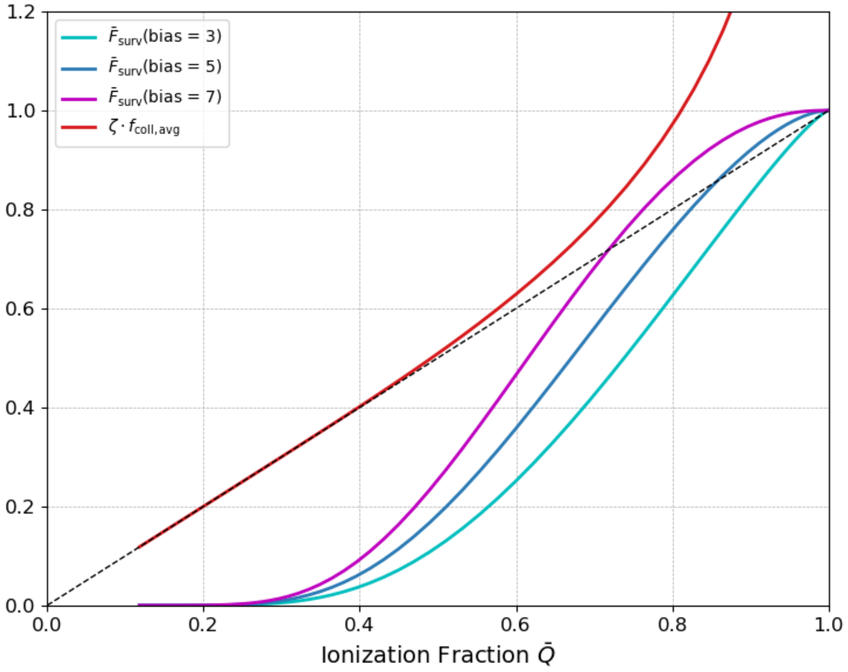

Figure 5.3: The fraction of LAEs that are visible $\bar{F}_{\mathrm{surv}}$ as a function of average ionization fraction $\bar{Q}$ for various bias values. Higher bias values result in a larger fraction of all LAEs being visible for a given ionization fraction of the Universe. The *red* line shows the total ionizing photon production rate $\zeta \cdot f_{\mathrm{coll,avg}}$, for which a value above the line $x = y$ means photons are escaping their host galaxies but being absorbed before contributing to reionization. The most over-dense R_α regions are overproducing photons but can't help their neighbors reionize. This effect is only large when the Universe is mostly reionized.

5.5.2 LAE galaxy associations in a partially ionized Universe

In the previous subsection, we examined the visibility of LAEs in regions of size R_α, the minimum ionized bubble size to host such sources. But in practice, galaxy associations may subtend significantly larger scales in real surveys. The Hu et al. (2021) volume is larger than R_α, meaning it may have some sub-chunks that are ionized, and some that are not. In this section, we consider the distribution of $F_{\rm surv}(\delta, R > R_\alpha)$: the fraction of LAEs that are visible in a region of size $R > R_\alpha$ and density δ.

Such a region has $N = (R/R_\alpha)^3$ sub-chunks of size R_α, each with a density δ_i distributed around δ with standard deviation $\sigma_{\rm w} = \sqrt{\sigma_{R_\alpha}^2 - \sigma_R^2}$, where σ_{R_α} and σ_R are the rms density fluctuation of the dark matter on a scale of R_α and R, respectively. Each chunk then also has corresponding $f_{\rm coll,i}$. Each of those sub-chunks are either ionized ($f_{\rm coll,i} > f_{\rm coll,vis}$, allowing LAEs to be visible) or not (so that their LAEs are invisible). We can then calculate $F_{\rm vis}$, $F_{\rm surv}$, and Q for the region of size R via weighted averages over the sub-chunks rather than integrals (as we did in the last section).

In Figure 5.4, we show that $F_{\rm surv}$ can vary widely between volumes of the same size R and overall density δ (more-so than Q in those regions!), due to the effects of Poisson noise and density fluctuations on the sub-chunk scale R_α. When the value of $f_{\rm coll,vis}$ is large (or equivalently when $\delta_{\rm vis}$ is large and $\bar{Q} \ll 1$), the distribution of $F_{\rm surv}$ can become bimodal, with a large fraction of regions having zero visible LAEs and all of the visible LAEs confined to a few very over-dense regions.

5.5.3 The inference framework

Within the framework described above, we now imagine that a survey has found an association of LAEs, and we use the model in order to measure three interesting quantities: *(a)* the local density of the region (as in section 5.4), *(b)* the local ionized fraction in the region, and *(c)* the *global* ionized fraction. We next describe how we make such inferences in practice. This procedure must be tuned to the specific construction of the survey; here, we imagine the simple case of choosing the most extreme apparent over-density in a survey, similar to

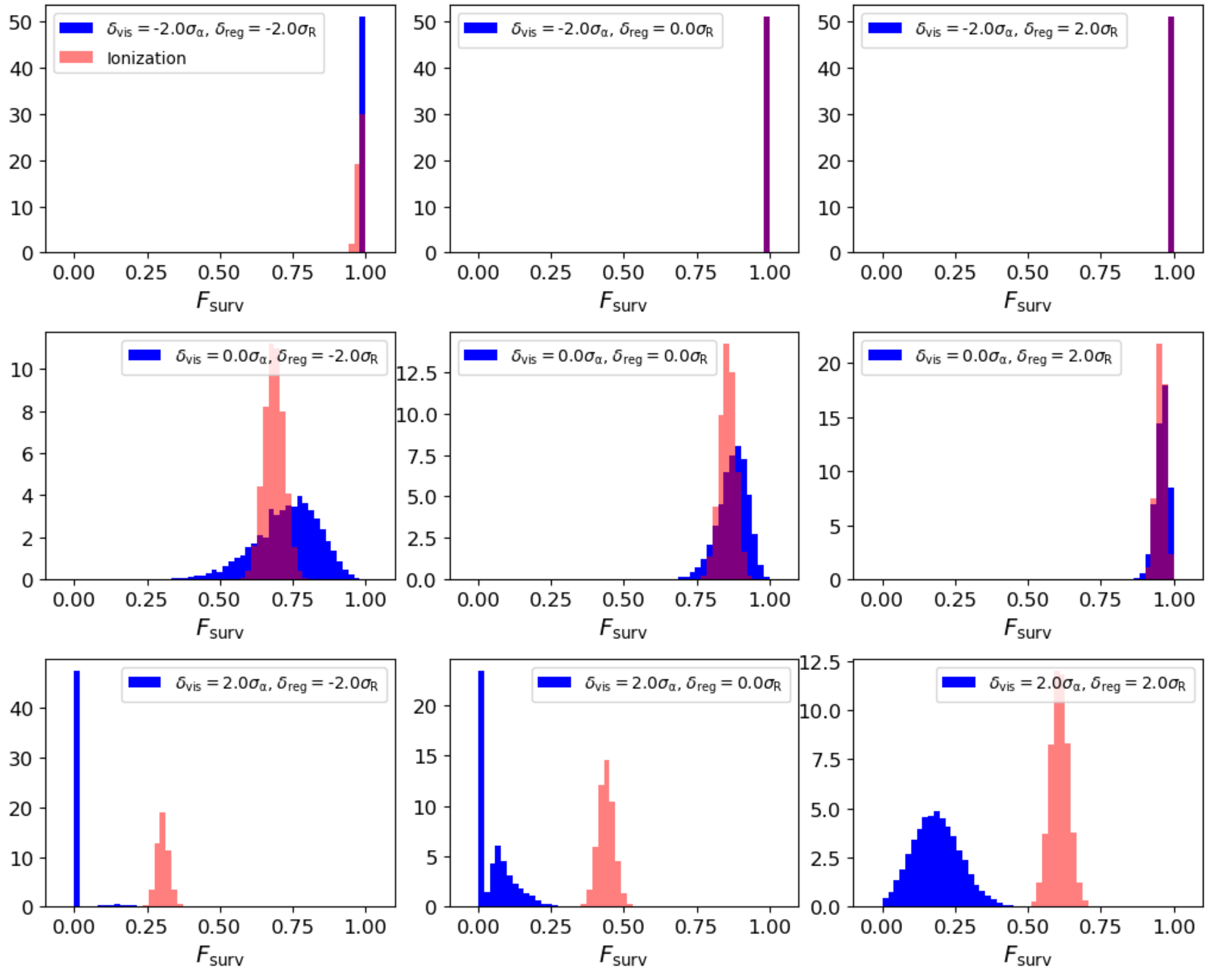

Figure 5.4: The distribution of the LAE survival fraction $F_{\rm surv}$ (blue histograms) and local ionized fraction (red histograms) for regions matching the Hu et al. (2021) volume. The panels vary the density $\delta_{\rm reg}$ of that region (left to right) and the density required for a region of size R_α to fully ionize, $\delta_{\rm vis}$ (top to bottom). $F_{\rm surv}$ can vary widely between regions of the same size and density (more-so than Q in those regions!), due to the effects of Poisson noise. When the value of $\delta_{\rm vis}$ is large ($\bar{Q} \ll 1$), the distribution of $F_{\rm surv}$ can become bimodal, with a large fraction of regions having zero visible LAEs.

the method of Hu et al. (2021).

We generate the likelihood of finding $N_{\rm LAE}$ sources in a region with *(i)* radius $R > R_\alpha$ pMpc among a larger survey with $N_{\rm vols}$ regions of the same size, *(ii)* average visible LAE number $N_{\rm exp,avg,vis}$, *(iii)* density δ, and *(iv)* collapse fraction required for sub-chunks to be ionized $f_{\rm coll,vis}$ by running a Monte-Carlo simulation.

1. First, we choose a value for $f_{\rm coll,vis}$, which defines $\bar{F}_{\rm surv}$ and $\bar{Q}$. As shown in Figure 5.2, this implicitly determines the effective ionizing efficiency ζ. This also defines the *actual* expected number of LAEs through $N_{\rm \ exp,avg} = N_{\rm exp,avg,vis}/\bar{F}_{\rm surv}$.

2. Next, we generate a set of $N_{\rm vols}$ volumes each with a dark matter density δ drawn from a normal function with standard deviation σ_R. Each of those volumes has $N_{\rm sub} = (R/R_\alpha)^3$ sub-chunks with densities $\delta_{\rm sub}$ drawn from a normal[2] at δ with standard deviation $\sigma_{\rm w}$. These densities each have corresponding $f_{\rm coll,sub}$.

3. The number of sources expected per sub-chunk is $N_{\rm exp,sub} = N_{\rm avg,\ sub}(1 + b \cdot \delta_{\rm sub})$, with $N_{\rm avg,\ sub} = N_{\rm \ exp,avg}/N_{\rm sub}$. We then draw from a Poisson distribution for each subchunk with $\lambda = N_{\rm exp,sub}$ to get the number of LAEs per subchunk $N_{\rm draw,sub}$. Finally, we sum all LAEs in subchunks with $f_{\rm coll,sub} > f_{\rm coll,vis}$ to get the total number of *observed* LAEs. Then, out of the $N_{\rm vols}$ volumes, we keep the one with the most observed LAEs, mimicking the procedure of Hu et al. (2021).

4. We repeat (i) - (iii) many times, each time choosing a random value for $N_{\rm exp,avg,vis}$ and b according to their uncertainty. The likelihood $P(N_{\rm LAE}|b, f_{\rm coll,vis})$ is the fraction of volumes chosen in step (iii) that have the correct number of observed LAEs ($N_{\rm obs} = N_{\rm LAE}$) This step implicitly marginalizes over the density δ and uncertainty in $N_{\rm exp,avg,vis}$ and b.

---

[2] When drawing this way, the sub-chunk densities $\delta_{\rm sub}$ will not add up to exactly δ. However, with the large number of sub-chunks considered in this paper, the deviation is small. When $N_{\rm sub}$ is small, δ must be re-calculated from the average of all $\delta_{\rm sub}$ values. However, this process imposes a broadening in the distribution of δ that must be corrected for.

5. We then repeat (i)-(iv) for many values of $f_{\rm coll,vis}$, each having a corresponding $\bar{F}_{\rm surv}$ and $\bar{Q}$ (depending on the random draw of b). This allows us to construct the likelihood as a function of $\bar{Q}$.
6. In order to transform the likelihood into a posterior distribution, we require a prior on the intrinsic LAE density. Many such choices are possible; below, we multiply by a prior that enforces that the number density of LAE hosts (whether visible as LAEs or not) must *increase* monotonically as redshift decreases (see below for an explanation of this prior).
7. Finally, we take all volumes selected above across every value of $f_{\rm coll,vis}$, weight by our prior, and make a histogram of the densities δ. This histogram is the posterior of δ marginalized over $\bar{Q}$. Similarly, we take these volumes, weight by the prior, and construct the posterior around the ionized fraction of the LAE association, marginalized over $\bar{Q}$ and δ.

5.5.4 The choice of priors

The most important prior in our framework is that on the underlying density of a region, which is, to linear order, simply a gaussian centered at zero with standard deviation σ_R. This prior is determined by integrating the power-spectrum of dark matter fluctuations over the desired physical scale (see e.g., Newman & Davis, 2002; Stark et al., 2007; Muñoz et al., 2010; Robertson, 2010; Trapp & Furlanetto, 2020) to obtain σ_R. This prior is highly constraining; in sections 5.4 and 5.6, the posterior of the density is prior-dominated. Fortunately, this prior is also very well-specified by a variety of cosmological probes.

A second prior, limiting the underlying galaxy density, also turns out to be important. In practice, one reasonable solution for a survey with a single, large LAE association and no other visible LAEs would be a highly neutral Universe with a single large ionized bubble. However, this would require that the true number density of LAE hosts be much larger than the observed density, because most of the Lyman-α lines are attenuated by the IGM. It is thus helpful to include a prior on the underlying density of the host galaxies. There are a variety of ways one can do this; we make a conservative choice here.

In particular, we include a prior requiring that there are not intrinsically more LAEs – for example – at $z = 6.93$ than at $z = 5.7$. In reality, there are likely many more LAEs at $z = 5.7$ than $z = 6.93$, because the underlying luminosity function of galaxies is evolving rapidly, so this serves as a conservative bound. Wold et al. (2022) finds the LAE luminosity density ratio between $z = 6.9$ and 5.7 to be $\rho_{z=6.9}/\rho_{z=5.7} = 0.63^{+0.13}_{-0.15}$. As long as the number density of LAEs stayed constant or increased from $z = 6.9$ to 5.7, this measurement can be converted to a lower-bound on the value of $\bar{F}_{\rm surv}(z = 6.9) > 0.63_{-0.15}$. We then convert to a lower bound on $\bar{Q}$ using our mapping between $\bar{F}_{\rm surv}$ and $\bar{Q}$ (see Fig. 5.3). Again, this is a conservative approach, because there are many fewer galaxies overall at $z \sim 7$ than at $z \sim 6$. One could incorporate a full model for the LAE population as a more sophisticated prior, or else simultaneously constrain both the overall galaxy evolution and the LAE distribution.

Finally, we use the bias value and its uncertainty from Hu et al. (2021) (in turn taken from Ouchi et al., 2018) of $b = 7.31 \pm 1.01$ (after converting to a linear bias). A theoretical model of cosmic variance predicts a value of the linear bias between 2-4 for halo masses between 10^9 and $10^{11} M_\odot$ at this redshift and scale (Trapp & Furlanetto, 2020), significantly lower than the observationally obtained value. However, we have shown in section 5.4 and will show in section 5.6 that switching to a lower bias value has little effect on the posteriors, likely owing to the fact that the prior on the linear density is dominating.

5.6 Applying the full inference framework

As an example of the inference framework, we now apply our procedure to the LAE association observed by Hu et al. (2021). In this case, $R = 3.7$ pMpc giving $N_{\rm sub} = 51$, $N_{\rm vols} = 20$, $N_{\rm exp,avg,vis} = 2.9 \pm 1.1$, and $b = 7.31 \pm 1.02$.

5.6.1 What can we learn about reionization?

Figure 5.5 shows the likelihood and posterior of the globally-averaged ionized fraction of the Universe at $z = 6.93$. The median and inner 68.27% probability of the posterior is $\bar{Q} = 0.70^{+0.14}_{-0.09}$. At $\bar{Q} \lesssim 0.7$, the posterior is prior-dominated. The prior represents a case where the number density of LAEs *stayed the same* between z =6.9 and 5.7. In reality, the number density likely increased dramatically, which would push the posterior to even higher values of $\bar{Q}$.

We note that our analysis disfavors both small ionized fractions (largely due to the prior on the LAE abundance) and a nearly ionized Universe. The latter is perhaps the most interesting aspect, as the constraint comes from finding a single region with so many LAEs and is not driven by any of our priors. Thus our simple model agrees with other measurements that suggest reionization is incomplete (but relatively advanced) at $z \sim 7$ (Greig et al., 2017; Inoue et al., 2018; Mason et al., 2018; Davies et al., 2018).

5.6.2 Is the association inside a large ionized bubble?

Figure 5.6 shows the histogram of the ionization states Q from each MC simulation that resulted in finding 14 sources, weighted by the prior on $\bar{Q}$, and the same for $F_{\rm surv}$, the fraction of LAEs in the Hu et al. (2021) volume that are visible. The 95.45% confidence lower limits for these quantities are $Q > 0.75$ and $F_{\rm surv} > 0.61$. Thus the qualitative association of this LAE association with a large ionized region is largely validated: the modal outcome is in fact a completely ionized region, and the volume is highly ionized with high confidence. Moreover, the fraction of LAE hosts that are visible as line emitters is also quite large.

5.6.3 Is the association a protocluster?

Figure 5.7 shows the posterior of the density, which yields a 68% confidence interval of $\delta = 0.13 \pm 0.05$, only slightly lower than the results from the fully-ionized case (see Fig. 5.1). Given the r.m.s. density fluctuation in the linear density field for the Hu et al. (2021) volume

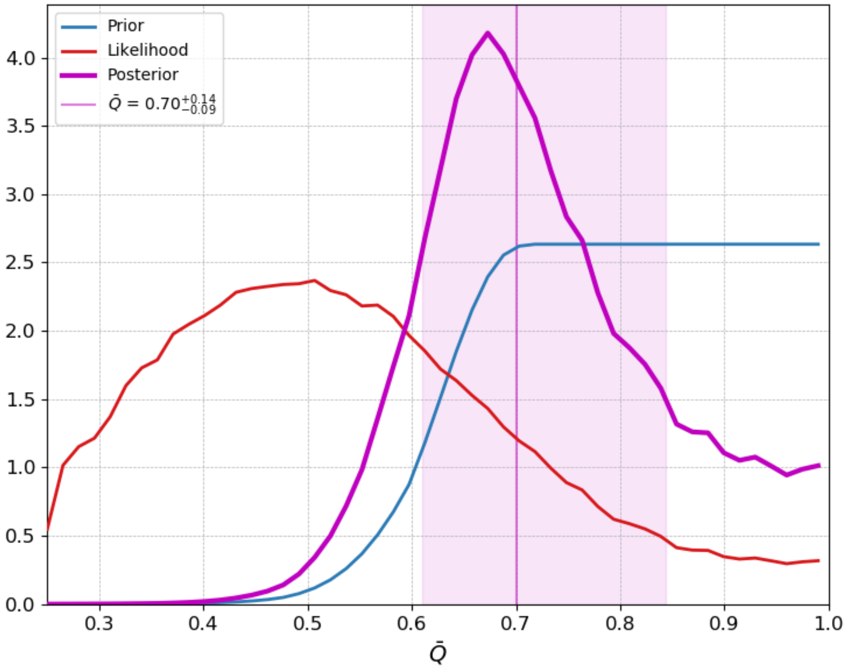

Figure 5.5: Constraints on the global ionized fraction $\bar{Q}$ from the Hu et al. (2021) observation. The *magenta* curve shows the posterior on the ionization of the Universe and 68.27% confidence interval $\bar{Q} = 0.70^{+0.14}_{-0.09}$. The *red* curve shows the likelihood of finding 14 sources given the average ionization of the Universe $\bar{Q}$ at $z = 6.93$ and other model parameters, while the *blue* curve shows the prior on the ionized fraction of the Universe, assuming the intrinsic number of LAEs increased from $z = 6.93$ to 5.7. At $\bar{Q} \lesssim 0.7$, the posterior is prior-dominated.

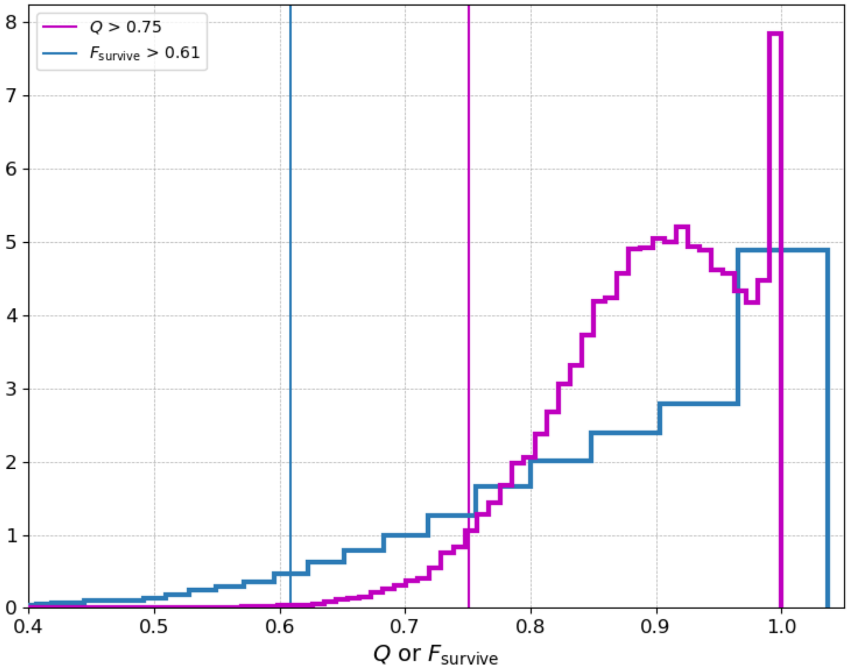

Figure 5.6: The Hu et al. (2021) volume is highly ionized with most bright LAEs visible as line emitters, as suspected by those authors. The *magenta* histogram shows the ionization states Q from each MC simulation that resulted in finding 14 sources, weighted by the prior on $\bar{Q}$. The *blue* curves the same but for F_{surv}, the fraction of LAEs that are visible. The vertical bars indicate the 95.45% lower limits of these quantities, $Q > 0.75$, $F_{\mathrm{surv}} > 0.61$.

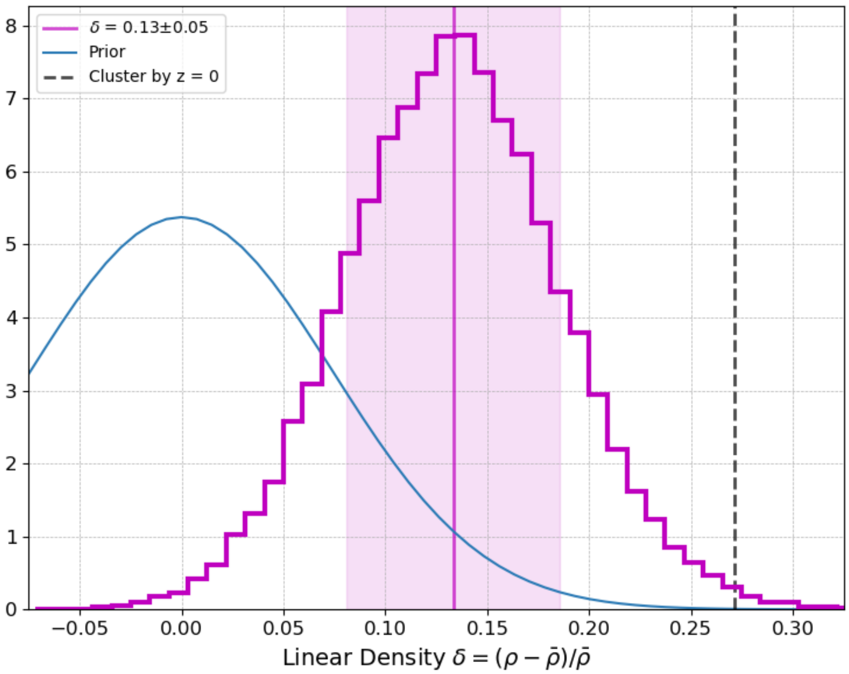

Figure 5.7: Constraints on the overall density of the Hu et al. (2021) volume. The *magenta* curve shows the histogram of all densities that resulted in finding 14 sources in the MC simulation described in section 5.5.3, weighted by the prior on $\bar{Q}$. This gives a 68% confidence interval of $\delta = 0.13 \pm 0.5$, only slightly lower than the results from the fully-ionized case (see Fig. 5.1). The entire region is not likely to have collapsed into a single cluster by $z = 0$, which would require a linear density of $\delta_{\mathrm{pc}} = 0.27$ (*black dashed* line). The prior on the density is shown in *blue*.

σ_R =0.074, this corresponds to a 1.8 ± 0.7 σ_R over-density. Despite its extreme apparent over-density, the entire region is not likely to have collapsed into a single cluster by $z = 0$, which would require a linear density of $\delta_{\mathrm{pc}} = 0.27$. However, it is possible that one or more sub-regions within this larger volume are at sufficiently high density to collapse into clusters by $z = 0$.

5.6.4 Sensitivity to model parameters

We explore the sensitivity of our results to model choices by re-running our inference in the following cases:

1. No marginalization: when re-doing our inference without marginalizing over the uncertainty in $N_{\rm exp,avg,vis}$ nor b, the posterior on the ionization moves to $\bar{Q} = 0.69^{+0.13}_{-0.09}$. The density posterior becomes $\delta = 0.14 \pm 0.05$.
2. Reduced bias: when changing the bias values from $b = 7.31$ to 5 (with same relative uncertainty), the posterior on the ionization moves to $\bar{Q} = 0.70^{+0.12}_{-0.09}$. The density posterior becomes $\delta = 0.14 \pm 0.05$.
3. $R_\alpha = 0.75$ Mpc: when reducing the parameter R_α, the posterior on the ionization moves to $\bar{Q} = 0.67^{+0.17}_{-0.11}$. The density posterior does not change at all, staying at $\delta = 0.14 \pm 0.05$.
4. $R_\alpha = 1.5$ Mpc: when increasing the parameter R_α, the posterior on the ionization moves to $\bar{Q} = 0.76^{+0.10}_{-0.07}$. The density posterior does not change at all, staying at $\delta = 0.14 \pm 0.05$.

$N_{\rm exp,avg,vis}$ and the bias value b do not appear to have a strong effect on the inference of the Hu et al. (2021) region's density nor the average ionization $\bar{Q}$. The choice of R_α has a larger effect on the results, though not dramatically.

In equation (5.8), we assume all galaxies contribute equally to reionization. A more realistic model of reionization would change the mapping between $\bar{Q}$ and $\bar{F}_{\rm surv}$ (see Fig. 5.3). However, this change might only have a small effect on our results, given that they do not vary strongly with the choice of bias, which modifies the mapping between $\bar{Q}$ and $\bar{F}_{\rm surv}$ significantly. A more realistic model of reionization would likely have more of an effect on our results through $F_{\rm surv}$, the distribution of which is important in our model and is evaluated on a small scale that is thus more subject to potentially complex behaviors of reionization.

In our MC simulation, we included a treatment of the "look-elsewhere" effect. That is, the Hu et al. (2021) volume was chosen out of a much larger volume because it had the most sources. This could potentially bias the inference of the region's density, ionization, etc. (especially because this volume was selected "by eye" rather than in a blind tiling of the survey volume). If we remove the treatment by choosing $N_{\rm vols} = 1$ in *(ii)* of section 5.5.3

instead of $N_{\rm vols} = 20$, the peak in the likelihood for $\bar{Q}$ is shifted significantly from $\sim 0.5 \rightarrow 0.3$. However, because our prior rules out those values, the resulting effect on the posterior is very small, shifting it to lower $\bar{Q}$ by only ~ 0.01. Thus the most crucial aspect of our model is likely the prior on the underlying galaxy density, which is essential to interpreting the $F_{\rm surv}$ values.

5.7 Conclusions

There is a long history of close study of “extreme" objects in astrophysics, as they can offer sharp tests of our physical paradigms. At high redshifts, associations of bright galaxies are often labeled as “protoclusters,” although the mapping to present-day clusters has only been qualitative. Because the neutral IGM modulates Lyman-α absorption, associations of LAEs have similarly been suspected to identify large ionized regions. In this paper, we have introduced a framework to make these identifications in a statistically rigorous manner. We first construct an analytic form of the posterior for the underlying density of a region given a finite number of observed galaxies, ignoring any modulation from reionization. This framework combines the effects of cosmic variance and Poisson noise, making use of a strong prior on the density of a region.

Inferences about the ionization field require a model of the reionization process. Here we have used a simple model motivated by the recent measurement of a surprisingly short mean free path for ionizing photons during this era (Becker et al., 2021). Using this model in a Monte Carlo simulation, we found that sufficiently large LAE associations can not only identify ionized volumes with high reliability but also constrain reionization on a global scale.

Assuming a fully-ionized Universe at $z = 6.93$, we calculate the linear dark-matter density of the Hu et al. (2021) volume to be $\delta = 0.16 \pm 0.05$. When considering a partially-ionized Universe via a simple model of re-ionization, we calculate the density to be slightly lower, $\delta = 0.13 \pm 0.05$. These densities are $> 2\sigma$ below the required linear density of $\delta_{\rm pc} = 0.27$ for this region to collapse into a single virialized object by $z = 0$. The Hu et al. (2021) volume gives a constraint on the ionized fraction of the Universe at $z = 6.93$, $\bar{Q} = 0.70^{+0.14}_{-0.09}$. This result is strongly prior-dominated at low values of $\bar{Q}$, and should only be interpreted in the context of our very simple model of re-ionization. We constrain the ionization of the Hu et al. (2021) volume itself to be $Q > 0.75$ at 95.45% confidence.

Our inferences about reionization are subject to systematic uncertainties about the underlying reionization model, but we already find that even a *single* well-defined LAE association offers competitive constraints on the global ionized fraction at $z \sim 7$ (Greig et al., 2017; Inoue et al., 2018; Mason et al., 2018; Davies et al., 2018); our results are consistent with other methods and have similar uncertainties. Our approach is more similar to a "counts-in-cells" method that leverages the non-gaussianity that reionization induces in the LAE distribution.

An advantage of our framework is that it identifies ionized regions in well-specified locations on the sky – providing targets for detailed studies of the effect of these ionized regions on the galaxy populations.

The constraints from the simple exercise in this paper suggest that the counts-in-cells approach may be very powerful. We have focused on a single association using a simple model. Future improvements to this framework could include: *(i)* a more complete reionization/Lyman-α absorption model, *(ii)* incorporating the distribution of source luminosities with a LAE luminosity function, *(iii)* considering all observed regions simultaneously rather than just a single association (or in other words implementing a full counts-in-cells framework), *(iv)* simultaneously incorporating information from photometric galaxy selection (e.g., Yoshioka et al. 2022) and LAE surveys, and *(v)* considering the expansion/contraction of a region depending on its density, which increases the relative odds of finding an under-dense region (Muñoz et al., 2010; Trapp & Furlanetto, 2020).

Future observations with JWST and other telescopes will discover many LAEs. This hugely increased sample – combined with a more accurate model of re-ionization – is a promising avenue for calculating the ionized fraction of the Universe throughout re-ionization.

CHAPTER 6

Conclusion

The models and frameworks developed in this dissertation were calibrated and tested using currently available data, but are intended mainly as a tool for future observations, especially those made by JWST and the Roman Space Telescope. Together, these instruments will delve deeper, search wider, and measure finer than any combination of current observatories. Of particular interest are the hundreds of large-scale density measurements that JWST will provide at a variety of redshifts (see Chapter 3). These densities are caused by the small inhomogeneities in the soup before the cosmic dawn and are the larger cousins of the seeds of galaxies themselves. We expect these densities to play an important role in reionization, both due to their direct effect on its progression (Becker et al., 2018; Davies et al., 2018; Christenson et al., 2021) and due to their effects on large-scale radiation backgrounds that can suppress star-formation at these early times (Thoul & Weinberg, 1996; Iliev et al., 2007; Noh & McQuinn, 2014). Further, understanding the assembly history of unusual objects like galaxy clusters is facilitated by measuring these large-scale densities (e.g., Chiang et al., 2017). We forecast the effectiveness of future JWST surveys in measuring these densities in Chapter 3, finding they will do so to the maximum precision set by Poisson noise. Simultaneously, they will measure the luminosity function of galaxies to extreme precision (if all surveys can be effectively combined).

In Chapter 5, we develop a statistically robust method for measuring the over/under-density of a region with some observed number of LAEs (or normal galaxies). When combined with a model of reionization, this model can also produce posteriors for the ionization of such a region and the average ionization fraction of the Universe. We apply this method to one such association of LAEs, finding that while still prior-dominated in this case, our single inference can still meaningfully constrain the ionization of the Universe (in the context of our simple model of reionization) as well as that region's density. JWST will find many more such regions and with higher precision in nearly all aspects, the inferences from which can be combined for a more accurate measure of the average ionization of the Universe.

Bibliography

Akeson, R., Armus, L., Bachelet, E., et al. 2019, arXiv e-prints, arXiv:1902.05569

Atek, H., Richard, J., Kneib, J.-P., & Schaerer, D. 2018, MNRAS, 479, 5184

Barkana, R., & Loeb, A. 2004, ApJ, 609, 474

Becker, G. D., D'Aloisio, A., Christenson, H. M., et al. 2021, MNRAS, 508, 1853

Becker, G. D., Davies, F. B., Furlanetto, S. R., et al. 2018, ApJ, 863, 92

Behroozi, P., Wechsler, R. H., Hearin, A. P., & Conroy, C. 2019, MNRAS, 488, 3143

Behroozi, P. S., & Silk, J. 2015, ApJ, 799, 32

Bhowmick, A. K., Somerville, R. S., Di Matteo, T., et al. 2020, MNRAS, 496, 754

Bond, J. R., Cole, S., Efstathiou, G., & Kaiser, N. 1991, ApJ, 379, 440

Bouwens, R. J., Stefanon, M., Oesch, P. A., et al. 2019, ApJ, 880, 25

Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015, ApJ, 803, 34

Bouwens, R. J., Oesch, P. A., Labbé, I., et al. 2016, ApJ, 830, 67

Bouwens, R. J., Oesch, P. A., Stefanon, M., et al. 2021, AJ, 162, 47

Bowler, R. A. A., Jarvis, M. J., Dunlop, J. S., et al. 2020, MNRAS, 493, 2059

Bowler, R. A. A., Dunlop, J. S., McLure, R. J., et al. 2014, MNRAS, 440, 2810

—. 2015, MNRAS, 452, 1817

Calvi, R., Rodríguez Espinosa, J. M., Mas-Hesse, J. M., et al. 2019, MNRAS, 489, 3294

Chiang, Y.-K., Overzier, R. A., Gebhardt, K., & Henriques, B. 2017, ApJ, 844, L23

Christenson, H. M., Becker, G. D., Furlanetto, S. R., et al. 2021, ApJ, 923, 87

Davies, F. B., Becker, G. D., & Furlanetto, S. R. 2018, ApJ, 860, 155

Davies, F. B., Bosman, S. E. I., Furlanetto, S. R., Becker, G. D., & D'Aloisio, A. 2021, ApJ, 918, L35

Davies, F. B., & Furlanetto, S. R. 2022, MNRAS, arXiv:2103.09821

Dekel, A., & Mandelker, N. 2014, MNRAS, 444, 2071

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, ApJ, 142, 414

Dore, O., Hirata, C., Wang, Y., et al. 2019, BAAS, 51, 341

Efstathiou, G., Ellis, R. S., & Peterson, B. A. 1988, MNRAS, 232, 431

Eisenstein, D. J., & Hu, W. 1998, ApJ, 496, 605

Endsley, R., Stark, D. P., Charlot, S., et al. 2021, MNRAS, 502, 6044

Fakhouri, O., Ma, C.-P., & Boylan-Kolchin, M. 2010, MNRAS, 406, 2267

Faucher-Giguère, C.-A., Kereš, D., & Ma, C.-P. 2011, MNRAS, 417, 2982

Faucher-Giguère, C.-A., Quataert, E., & Hopkins, P. F. 2013, MNRAS, 433, 1970

Finkelstein, S. L., Ryan, Russell E., J., Papovich, C., et al. 2015, ApJ, 810, 71

Finkelstein, S. L., Bagley, M., Song, M., et al. 2022, ApJ, 928, 52

Furlanetto, S. R., Hernquist, L., & Zaldarriaga, M. 2004, MNRAS, 354, 695

Furlanetto, S. R., Mirocha, J., Mebane, R. H., & Sun, G. 2017, MNRAS, 472, 1576

Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. 2006, MNRAS, 365, 1012

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

Goerdt, T., Ceverino, D., Dekel, A., & Teyssier, R. 2015, MNRAS, 454, 637

Greig, B., Mesinger, A., Haiman, Z., & Simcoe, R. A. 2017, MNRAS, 466, 4239

Gunn, J. E., & Peterson, B. A. 1965, ApJ, 142, 1633

Haiman, Z. 2002, ApJ, 576, L1

Hayward, C. C., & Hopkins, P. F. 2017, MNRAS, 465, 1682

Hu, W., Wang, J., Infante, L., et al. 2021, Nature Astronomy, 5, 485

Hunter, J. D. 2007, Computing in Science & Engineering, 9, 90

Iliev, I. T., Mellema, G., Shapiro, P. R., & Pen, U.-L. 2007, MNRAS, 376, 534

Inoue, A. K., Hasegawa, K., Ishiyama, T., et al. 2018, PASJ, 70, 55

Ishigaki, M., Kawamata, R., Ouchi, M., et al. 2015, ApJ, 799, 12

Itoh, R., Ouchi, M., Zhang, H., et al. 2018, ApJ, 867, 46

Jung, I., Finkelstein, S. L., Dickinson, M., et al. 2020, ApJ, 904, 144

Kashikawa, N., Shimasaku, K., Matsuda, Y., et al. 2011, ApJ, 734, 119

Kauffmann, O. B., Le Fèvre, O., Ilbert, O., et al. 2020, A&A, 640, A67

Kitzbichler, M. G., & White, S. D. M. 2007, MNRAS, 376, 2

Krumholz, M. R., Burkhart, B., Forbes, J. C., & Crocker, R. M. 2018, MNRAS, 477, 2716

Lacey, C., & Cole, S. 1993, MNRAS, 262, 627

Lin, Y., Oh, S. P., Furlanetto, S. R., & Sutter, P. M. 2016, MNRAS, 461, 3361

Livermore, R. C., Finkelstein, S. L., & Lotz, J. M. 2017, ApJ, 835, 113

Loeb, A., & Furlanetto, S. R. 2013, The First Galaxies in the Universe (Princeton University Press)

Lovell, C. C., Vijayan, A. P., Thomas, P. A., et al. 2021, MNRAS, 500, 2127

Madau, P., & Dickinson, M. 2014, ARA&A, 52, 415

Madau, P., & Rees, M. J. 2000, ApJ, 542, L69

Malhotra, S., & Rhoads, J. E. 2004, ApJ, 617, L5

Mason, C. A., Trenti, M., & Treu, T. 2015, ApJ, 813, 21

Mason, C. A., Treu, T., Dijkstra, M., et al. 2018, ApJ, 856, 2

McBride, J., Fakhouri, O., & Ma, C.-P. 2009, MNRAS, 398, 1858

McKinney, W. 2010, in Proceedings of the 9th Python in Science Conference (SciPy)

McLeod, D. J., McLure, R. J., & Dunlop, J. S. 2016, MNRAS, 459, 3812

McLeod, D. J., McLure, R. J., Dunlop, J. S., et al. 2015, MNRAS, 450, 3032

McLure, R. J., Dunlop, J. S., Bowler, R. A. A., et al. 2013, MNRAS, 432, 2696

McQuinn, M., Lidz, A., Zahn, O., et al. 2007, MNRAS, 377, 1043

Mesinger, A., Furlanetto, S., & Cen, R. 2011, MNRAS, 411, 955

Mesinger, A., & Furlanetto, S. R. 2008, MNRAS, 386, 1990

Miralda-Escudé, J. 1998, ApJ, 501, 15

Mirocha, J., Furlanetto, S. R., & Sun, G. 2017, MNRAS, 464, 1365

Mirocha, J., Mason, C., & Stark, D. P. 2020, MNRAS, 498, 2645

Mo, H. J., & White, S. D. M. 1996, MNRAS, 282, 347

Moster, B. P., Somerville, R. S., Newman, J. A., & Rix, H.-W. 2011, ApJ, 731, 113

Muñoz, J. A., Trac, H., & Loeb, A. 2010, MNRAS, 405, 2001

Newman, J. A., & Davis, M. 2002, ApJ, 564, 567

Noh, Y., & McQuinn, M. 2014, MNRAS, 444, 503

Oesch, P. A., Bouwens, R. J., Illingworth, G. D., et al. 2015, ApJ, 808, 104

Oesch, P. A., Bouwens, R. J., Illingworth, G. D., Labbé, I., & Stefanon, M. 2018, ApJ, 855, 105

Oesch, P. A., Bouwens, R. J., Illingworth, G. D., et al. 2013, ApJ, 773, 75

Ono, Y., Ouchi, M., Mobasher, B., et al. 2012, ApJ, 744, 83

Ouchi, M., Ono, Y., & Shibuya, T. 2020, ARA&A, 58, 617

Ouchi, M., Shimasaku, K., Akiyama, M., et al. 2005, ApJ, 620, L1

Ouchi, M., Harikane, Y., Shibuya, T., et al. 2018, PASJ, 70, S13

Partridge, R. B., & Peebles, P. J. E. 1967, ApJ, 147, 868

Patiri, S. G., Betancort-Rijo, J., & Prada, F. 2006, MNRAS, 368, 1132

Pentericci, L., Fontana, A., Vanzella, E., et al. 2011, ApJ, 743, 132

Penzias, A. A., & Wilson, R. W. 1965, ApJ, 142, 419

Perez, F., & Granger, B. E. 2007, Computing in Science & Engineering, 9, 21

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, A&A, 594, A13

Press, W. H., & Schechter, P. 1974, ApJ, 187, 425

Roberts-Borsani, G., Morishita, T., Treu, T., Leethochawalit, N., & Trenti, M. 2022, ApJ, 927, 236

Robertson, B. E. 2010, ApJ, 713, 1266

Robertson, B. E., Ellis, R. S., Dunlop, J. S., et al. 2014, ApJ, 796, L27

Rubiño-Martín, J. A., Betancort-Rijo, J., & Patiri, S. G. 2008, MNRAS, 386, 2181

Sandage, A., Tammann, G. A., & Yahil, A. 1979, ApJ, 232, 352

Santos, M. R. 2004, MNRAS, 349, 1137

Schenker, M. A., Robertson, B. E., Ellis, R. S., et al. 2013, ApJ, 768, 196

Sheth, R. K. 1998, MNRAS, 300, 1057

Somerville, R. S., Lee, K., Ferguson, H. C., et al. 2004, ApJ, 600, L171

Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv e-prints, arXiv:1503.03757

Stark, D. P., Ellis, R. S., & Ouchi, M. 2011, ApJ, 728, L2

Stark, D. P., Loeb, A., & Ellis, R. S. 2007, ApJ, 668, 627

Sun, G., & Furlanetto, S. R. 2016, MNRAS, 460, 417

Tacchella, S., Trenti, M., & Carollo, C. M. 2013, ApJ, 768, L37

Thoul, A. A., & Weinberg, D. H. 1996, ApJ, 465, 608

Tilvi, V., Malhotra, S., Rhoads, J. E., et al. 2020, ApJ, 891, L10

Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709

Trac, H., Cen, R., & Mansfield, P. 2015, ApJ, 813, 54

Tramonte, D., Rubiño-Martín, J. A., Betancort-Rijo, J., & Dalla Vecchia, C. 2017, MNRAS, 467, 3424

Trapp, A. C., & Furlanetto, S. R. 2020, MNRAS, 499, 2401

Trapp, A. C., Furlanetto, S. R., & Yang, J. 2022, MNRAS, 510, 4844

Trenti, M., & Stiavelli, M. 2008, ApJ, 676, 767

Trenti, M., Stiavelli, M., Bouwens, R. J., et al. 2010, ApJ, 714, L202

Trenti, M., Bradley, L. D., Stiavelli, M., et al. 2012, ApJ, 746, 55

Ucci, G., Dayal, P., Hutter, A., et al. 2021, MNRAS, 506, 202

Vale, A., & Ostriker, J. P. 2004, MNRAS, 353, 189 van den Bosch, F. C., Jiang, F., Hearin, A., et al. 2014, MNRAS, 445, 1713

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, Computing in Science & Engineering, 13, 22

Virtanen, P., , Gommers, R., et al. 2020, Nature Methods, 17, 261

Vogelsberger, M., Nelson, D., Pillepich, A., et al. 2020, MNRAS, 492, 5167

Wold, I. G. B., Malhotra, S., Rhoads, J., et al. 2022, ApJ, 927, 36

Wyithe, J. S. B., & Loeb, A. 2005, ApJ, 625, 1

Yoshioka, T., Kashikawa, N., Inoue, A. K., et al. 2022, ApJ, 927, 32

Zitrin, A., Labbé, I., Belli, S., et al. 2015, ApJ, 810, L12