Thinking outside the box; fast error estimation for next generation galaxy surveys

by

Gregory Schreiner

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Science
in
Physics

Waterloo, Ontario, Canada, 2023

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Since the discovery of the accelerated expansion of the universe in the late 90s, the flat ΛCDM model has reigned as the best explanation for the cosmological phenomena we have observed. In spite of over two decades of study, the identities and properties of both cold dark matter (CDM) and dark energy (Λ) remain a mystery. The goal of modern precision cosmology is to measure cosmological parameters using a multitude of probes in the pursuit of deviations from the ΛCDM framework or insights into the properties of dark matter and dark energy. If the probe being considered is the matter power spectrum measured from a galaxy redshift survey, then the precision of the derived parameters is determined by the power spectrum covariance matrix. The analytic form of the covariance matrix is difficult to estimate, so common practice is to run many simulations of the survey volume to get a brute-force estimate.

Next-generation cosmological surveys are set to collect higher resolution data within a larger survey volume than ever before. The complexity and number of simulations that will be required to estimate the covariance matrix of these surveys is threatening to become too computationally expensive for even the most advanced computer clusters. Thus, there is an urgent need to develop novel techniques for reducing the computation time required to achieve such precise covariance estimates. While many proposed methods seek to reduce the number of simulations required, it is also possible to leverage the volume scaling of the covariance matrix, allowing one to reduce the size of the simulations required.

Super-sample covariance (SSC) is a contribution to the covariance matrix made by modes of the power spectrum that are larger than the volume of a survey or simulation. If this volume scaling of the covariance is to be taken advantage of, then the SSC within the simulations must be accurately modeled. To this end, I review methods of running separate universe (SU) simulations to account for the effects of SSC. While these methods have all been shown to recover the SSC with reasonable accuracy, they have been largely developed and tested in isolation from one another. I present my work in directly comparing the accuracy of these methods in recovering the SSC effect using ensembles of N-body simulations.

Even with SSC accurately modeled, the volume scaling of the covariance does not hold for arbitrarily small volume simulations; at some point, the analytic behaviour of the covariance is expected to break down. I push the volume scaling to its limit by running many thousands of simulations at different volumes and scaling the covariance to match that of a larger volume survey. The SSC term has a nontrivial relation to the simulation volume, preventing it from scaling in the same way as the other components of the covariance. In

light of this, I present a way to include SSC such that the scaled covariance could still be recovered with good accuracy. I find the scaled covariance matches the large volume covariance to within $\sim 4\%$ or better on most scales, with higher k bins being biased low due to missing a small component of the SSC. The scaled covariance at very low k for very small simulations is substantially lower than the large mock covariance at those scales due to very few modes of that scale being present in the small volume simulations. This creates a skewness in the distribution of power at those scales. By computing the number of modes required to avoid this skewed distribution of power, I derive a way to estimate the minimum simulation volume that could be used to accurately model the covariance at a given scale. The accurate modeling of SSC and optimal leveraging of the volume scaling of the covariance matrix are powerful complementary tools with the potential to substantially reduce the computational cost of covariance matrix estimation for future galaxy survey data.

Acknowledgements

I would like to thank my supervisor Prof. Will Percival for giving me the amazing opportunity to study at the Waterloo Centre for Astrophysics. Your encouragement, insights, and dedication to your students were irreplacable ingredients to the success of this project. I would also like to thank Alex Krolewski and Shahab Joudaki for co-supervising me on this project, and for their boundless patience in answering questions I was too embarassed to ask Will. Throughout the constant hurdles and roadblocks I faced on this project, the three of you were always there to offer guidance and support.

Thank you to my Advisory Committee, Prof. Niayesh Afshordi and Prof. James Taylor, for their guidance on the course of my degree. Your constructive feedback from our Committee meetings gave me valuable perspective on where I needed to focus my efforts.

Thank you to my fellow researchers both in the Percival lab group and the WCA as a whole with whom I shared both scientific discussions and entertaining banter. These conversations broadened my horizons and offered well-needed stress relief, respectively. In particular, I would like to thank my office mates Mike Chapman, Alex Woodfinden, Tristan Fraser, Alan Nguyen, Batia Friedman-Shaw, and James Morawetz for letting me bounce ideas off of them and for sympathizing when my code was broken. Special thanks to Marie-Joëlle Gingras for the never-ending supply of bubble gum.

Thank you to Adrian Bayer and Prof. Wayne Hu for useful discussions regarding super-sample covariance in simulations. Your insightful advice helped me get past the last great hurdle of this project.

Thank you to my family and friends, whose emotional support carried me through any hardships I found myself in. I truly appreciate your interest in my project and your patience as I tried to describe what a covariance matrix was using as little math as possible. An extra special thank you goes to my smarter half Chelsea Frank. You are my rock, and every day you challenge me to be the best version of myself I can be.

Finally, thank you to Raithin, Odin, Rukat, and Darva, who have waited so patiently for their dungeon master to return.

Dedication

To my dad. It means the world to know you're proud of me.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

ΛCDM Λ Cold Dark Matter 3, 8, 12, 16, 17, 25, 41, 48, 49, 68

2LPT second-order Lagrangian Perturbation Theory 23

2PCF 2-Point Correlation Function 6, 7, 70

BAO Baryon Acoustic Oscillation 10, 11, 13

BC Beat Coupling 43

CDM Cold Dark Matter 4

CIC Clouds-in-Cells 48, 49, 66

CMB Cosmic Microwave Background 11, 13, 16, 17, 22

COLA COmoving Lagrangian Acceleration 22–24, 50

DESI Dark Energy Spectroscopic Instrument 69

FFT Fast Fourier Transform 22, 24, 25, 49

FLRW Friedmann-Lemaitre-Robertson-Walker 1, 2, 35, 68

HSV Halo Sample Variance 43

L-PICOLA Lightcone-enabled Parallel Implementation of COmoving Lagrangian Acceleration 23, 24, 35, 39, 54

LD Linear Dilation 43

LPT Lagrangian Perturbation Theory 23

PBCs Periodic Boundary Conditions 31, 32

PM Particle Mesh 22–24, 30, 48–50, 66

SC Spherical Collapse 27, 36, 39, 41–48, 68

SSC Super-Sample Covariance 27, 29–32, 35, 36, 39, 41–43, 47–50, 52, 54, 55, 57, 58, 60, 61, 63, 65, 66, 68–71

SU Separate Universe 32, 33, 36, 37, 48, 51

Chapter 1

Introduction

1.1 Standard model of cosmology

The current model of cosmology is dependent on the cosmological principle and Einstein's theory of general relativity. The cosmological principle states that, on sufficiently large scales ($> 100\ Mpc$), the universe is homogeneous and isotropic. The homogeneity of the universe implies that any given (large) patch of the universe contains about the same amount of stuff and obeys the same laws of physics as any other patch. The isotropy of the universe implies that, on large scales, the universe looks the same in all directions. General relativity describes the geometry of the universe as a 4-dimensional spacetime. The presence of matter or energy causes this spacetime to curve, resulting in the force of gravity. This relationship is given by Einstein's field equation [25]

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \tag{1.1}$$

where $G_{\mu\nu}$ is the Einstein tensor describing the curvature of spacetime, $g_{\mu\nu}$ is the metric tensor, Λ is the cosmological constant, G is the gravitational constant, c is the speed of light, and $T_{\mu\nu}$ is the stress-energy tensor describing the density and flux of energy and momentum. Einstein's field equation and the cosmological principle can be combined to derive the general metric of spacetime, known as the Friedmann-Lemaitre-Robertson-Walker (FLRW) metric [31, 43, 60, 78]. The FLRW metric can be written in polar coordinates as

$$ds^2 = -c^2 dt^2 + a(t)^2 \left[dr^2 + S_\kappa(r)^2 d\Omega^2 \right], \tag{1.2}$$

where

$$d\Omega^2 \equiv d\theta^2 + sin^2\theta d\phi^2 \tag{1.3}$$

and

$$S_\kappa(r) = \begin{cases} Rsin(r/R) & \kappa = +1 \\ r & \kappa = 0 \\ Rsinh(r/R) & \kappa = -1 \end{cases} \tag{1.4}$$

Here, $a(t)$ is the scale factor of the universe. The scale factor describes how distances in the universe change with time as space expands and is normalized to $a(t) = 1$ at present day. In equation 1.4, R is the radius of curvature of the universe, and the value of κ represents the direction of the curvature: $\kappa = 1$ represents closed spherical geometry, $\kappa = 0$ represents flat geometry, and $\kappa = -1$ represents open hyperbolic geometry.

Using the FLRW metric and Einstein's field equations, equations describing the relationship between the energy content of the universe, its curvature, and its scale factor can be derived. The first of these equations is the Friedmann equation

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{3\pi G}{3c^2}\varepsilon(t) - \frac{\kappa c^2}{R_0^2}\frac{1}{a(t)^2}, \tag{1.5}$$

where $\varepsilon(t)$ is the total energy density of the universe, and R_0 is the radius of curvature of the universe measured at present day. The Friedmann equation tells us how fast the universe is expanding given its energy content. The expansion rate is often written in terms of the Hubble parameter $H \equiv \dot{a}/a$. The second equation of interest is the acceleration equation, which describes how the expansion rate changes over time. It is given by

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3c^2}(\varepsilon + 3P), \tag{1.6}$$

where P is the pressure within the universe. It is derived by combining the Friedmann equation with the fluid equation which describes how the energy density of the universe changes over time

$$\dot{\varepsilon} + 3\frac{\dot{a}}{a}(\varepsilon + P) = 0. \tag{1.7}$$

To solve these equations, the relationship between energy density and pressure must be known for the different components of the universe. This relation is known as the equation of state, and for substances relevant in a cosmological context it takes the general form

$$P = w\varepsilon, \tag{1.8}$$

where w is the equation of state parameter for the substance being considered. The equation of state parameters for common components of the universe are

$$\begin{array}{ll} w = 0 & \text{for nonrelativistic matter} \\ w = 1/3 & \text{for radiation} \\ w = -1 & \text{for a cosmological constant.} \end{array} \quad . \tag{1.9}$$

These can be plugged into the fluid equation to relate the energy density of each component to the scale factor. Nonrelativistic matter scales as $\varepsilon_m \propto a^{-3}$ since the density of matter dilutes as the volume of the universe increases. The number density of photons scales as a^{-3}, but the energy of each photon also scales as a^{-1} as its wavelength is stretched out by the expansion of the universe, resulting in the photon energy density scaling as $\varepsilon_r \propto a^{-4}$. The cosmological constant is gets its name from the fact that its energy density remains constant as the universe increases. For a universe with these three components, its expansion history can be split into the radiation-dominated, matter-dominated, and dark-energy-dominated eras, defined by the component that accounted for the majority of the energy density at that time.

When working with the Friedmann equation, it is useful to define the critical density $\varepsilon_c(t)$

$$\varepsilon_c(t) \equiv \frac{3c^2}{8\pi G} H(t)^2 \tag{1.10}$$

at which the curvature of the universe is exactly zero. From this, the energy density of the universe can be expressed using a dimensionless density parameter

$$\Omega(t) \equiv \frac{\varepsilon(t)}{\varepsilon_c(t)}. \tag{1.11}$$

The dimensionless density parameter Ω can be split into different components corresponding to matter Ω_m, radiation Ω_r, and the cosmological constant Ω_Λ. The Friedmann equation can be rewritten in terms of these density parameters

$$H^2 = H_0^2 \left[\Omega_{r,0} a^{-4} + \Omega_{m,0} a^{-3} + \Omega_{k,0} a^{-2} + \Omega_{\Lambda,0}\right], \tag{1.12}$$

where $\Omega_k = \kappa c^2 / R_0^2$ and subscript 0 denotes values at present day.

1.1.1 The ΛCDM model

The ΛCDM model, also known as the concordence model, is the current model of cosmology that best fits the observed properties of the universe. This model is spatially flat and has its present day energy density made up of three main components: baryonic matter, Cold Dark Matter (CDM), and dark energy in the form of a cosmological constant. Baryonic matter consists of protons, neutrons, electrons, and other similar massive particles. It makes up only a small fraction of the universe's energy budget. Dark matter is a currently unidentified type of matter that interacts via gravity [82, 53, 24, 61], but has not been detected to interact via any other mechanisms. As such, all measurements of the position

and abundance of dark matter are made indirectly by observing its gravitational effects on photons and baryonic matter. Indirect observations of dark matter have revealed that the overwhelming majority of it is moving at nonrelativistic speeds ($v << c$), so it is referred to as Cold Dark Matter (CDM). The tightest constraints on the amounts of matter and dark energy in the universe come from the Planck collaboration [4], finding values of $\Omega_m = 0.315 \pm 0.007$ and $\Omega_\Lambda = 0.685 \pm 0.007$.

1.2 Growth of structure

While the universe on large scales obeys the cosmological principle, the matter within is not distributed perfectly evenly. Initial small perturbations in the density field $\rho(\mathbf{x})$ grow due to the uneven force of gravity acting on nearby matter, causing overdense regions to become denser as time passes and underdense regions to become emptier. When working with the density field of the universe, it is useful to define the overdensity field

$$\delta(\mathbf{x}, t) = \frac{\rho(\mathbf{x}, t) - \overline{\rho}}{\overline{\rho}}, \tag{1.13}$$

where $\overline{\rho}$ is the mean density of the universe, such that $\delta > 0$ corresponds to regions of space with higher than average density and $\delta < 0$ corresponds to lower than average density. It is also useful to define its Fourier transform

$$\tilde{\delta}(k) = \frac{1}{\sqrt{V}} \int \delta(\mathbf{x}, t) e^{-i\mathbf{k}\cdot\mathbf{x}} d^3\mathbf{x}, \tag{1.14}$$

where V is the volume being integrated. It is useful to work with $\tilde{\delta}$ because, if $\tilde{\delta}$ is a Gaussian random field (as is true in the early universe), the different $\mathbf{k}$ modes will each evolve independently of one another. It should be noted that high $\mathbf{k}$ modes correspond to small wavelength perturbations and thus small-scale structure, and low $\mathbf{k}$ modes correspond to large scales.

1.2.1 Linear evolution

Under the assumption that these density perturbations are small ($\delta \ll 1$), an equation can be derived for how these perturbations change over time. Combining The continuity, Euler, and Poisson equations and assuming constant entropy across the entire field gives

$$\frac{d^2\tilde{\delta}}{dt^2} + 2\frac{\dot{a}}{a}\frac{d\tilde{\delta}}{dt} = \tilde{\delta}\left(4\pi G\overline{\rho} - \frac{c_s^2 k^2}{a^2}\right), \tag{1.15}$$

where c_s is the speed of sound. There is a characteristic length known as the Jeans length

$$\lambda_J = c_s \sqrt{\frac{\pi}{G\rho}}, \tag{1.16}$$

below which the pressure within a fluid is able to effectively oppose gravity and above which the perturbations are free to grow. During the radiation-dominated era, the speed of sound in the universe was $c_s = c/\sqrt{3}$. This caused the Jeans length during that era to be quite large, suppressing the growth of structure. Once the matter-dominated era began and especially after recombination occurred, c_s decreased substantially and matter perturbations were able to grow in earnest. Considering the density field in the matter-dominated era and on scales larger than λ_J, equation 1.15 simplifies to

$$\frac{d^2\tilde{\delta}}{dt^2} + 2\frac{\dot{a}}{a}\frac{d\tilde{\delta}}{dt} - 4\pi G\overline{\rho}\tilde{\delta} = 0, \tag{1.17}$$

which can be solved numerically for a given cosmology.

All density perturbations satisfying these assumptions will evolve identically to one another. In this case, one can define the linear growth factor,

$$D(a) = D(1)\frac{\delta(a)}{\delta(1)}, \tag{1.18}$$

characterizing how much the amplitudes of these perturbations have grown from scale factor a to present day. The linear growth factor can be calculated for any given cosmology as

$$D(a) = \frac{5\Omega_m}{2}\frac{H(a)}{H_0}\int_0^a \frac{da'}{(a'H(a')/H_0)^3}, \tag{1.19}$$

where the form of $H(a)$ is determined by the Friedmann equation. Here, the growth factor is normalized such that, during the matter-dominated era, $D(a) = a$. Using equation 1.18, the linear evolution of δ can be calculated for any scale factor of interest. The linear evolution of the power spectrum can also be calculated as

$$P(k, a) = \left(\frac{D(a)}{D(1)}\right)^2 P(k, 1). \tag{1.20}$$

1.2.2 Nonlinear growth and spherical collapse

While linear theory serves well for small density perturbations, the approximations used to build it break down once the perturbations grow to $\delta \sim 1$. The simplest nonlinear extension

to linear theory is the spherical top-hat collapse model [34, 56]. The method considers the density perturbation to be a homogeneous sphere that begins with slightly higher overdensity than the background universe. As space expands, the radius of the sphere grows to a maximum value, after which point the gravity of the region is strong enough to overcome the Hubble flow and the region begins collapsing. Under this approximation, this collapse would theoretically continue until the radius of the region was zero, but in reality this spherical region instead becomes a virialized structure called a halo. From this formalism, a critical density for collapse $\delta_c = 1.686$ can be defined. This critical density corresponds to the linearly extrapolated present day density of a region required for the region to finish collapsing by today. This δ_c is useful in identifying the scales at which linear theory is expected to hold well and which scales undergo nonlinear evolution. Length scales greater than 40 $h^{-1}Mpc$ ($k \lesssim 0.15 hMpc^{-1}$) are considered linear, scales between 40 and 20 $h^{-1}Mpc$ ($0.15 < k < 0.3 hMpc^{-1}$) are quasilinear, and scales below 20 $h^{-1}Mpc$ ($k > 0.3\ hMpc^{-1}$) are considered nonlinear.

1.3 Quantifying matter clustering

The large scale structure of the universe is commonly studied through the lens of clustering statistics. In Sec. 1.3.1, I describe the most common 2-point statistics used to quantify matter clustering. In Sec. 1.3.2, I briefly cover higher order functions used to fully describe a non-Gaussian matter distribution. Sec. 1.4.7 introduces the power spectrum covariance matrix and its significance in evaluating the likelihood of a cosmological model.

1.3.1 2-point statistics

The most straightforward way to quantify the clustering of matter is using the 2-Point Correlation Function (2PCF) of the tracers of the overdensity field

$$\xi(\mathbf{x}') = \langle \delta(\mathbf{x} + \mathbf{x}')\delta(\mathbf{x}) \rangle \,. \tag{1.21}$$

The 2PCF captures how likely one is to find a tracer at position $\mathbf{x}+\mathbf{x}'$ given there is one at position $\mathbf{x}$ compared to the expectation from a Poisson distribution. In the case of galaxy surveys, this is usually calculated by pair counting where the separation of each pair of tracers is counted and binned in terms of separation distance.

The same clustering information captured in the 2PCF can also be quantified by its Fourier transform, the matter power spectrum

$$P(\mathbf{k}) = \left\langle \tilde{\delta}(\mathbf{k})\tilde{\delta}^*(\mathbf{k}) \right\rangle . \tag{1.22}$$

Working with the power spectrum has the advantage of the understanding of large-scale modes that follow linear evolution. These large-scale modes will evolve independently of one another, and their behaviour can be well predicted analytically.

The information in the power spectrum is often convenient to compress in terms of multipole moments

$$P_l(k) = (2l+1) \int_0^1 P(k,\mu)\mathcal{L}_l(\mu)d\mu, \tag{1.23}$$

where $\mathcal{L}_l$ are the Legendre polynomials and μ is the cosine of the angle to the line of sight. In the case of a homogeneous and isotropic universe, all moments beyond the monopole ($l = 0$) are expected to be zero. In real galaxy surveys, redshift space distortions can cause anisotropies along the line of sight, making the higher order multipoles useful for detecting them. The data used in this work are from simulations without redshift space distortions and so will focus on the monopole moment of the power spectrum $P(k)$.

1.3.2 Beyond the 2-point functions

In the case of a Gaussian random field, the field can be completely described by its 2-point statistic, and all higher order statistics are able to be derived from the 2-point function. The density field is approximately Gaussian at early times before significant gravitational evolution has occurred. It is also Gaussian on large scales where linear theory holds, even at late times. Nonlinear evolution of small scale structure at late times, however, induces non-Gaussianity in the field. In this case, higher order clustering statistics can no longer be calculated directly from the 2-point statistics and must be measured. The most common of these higher order statistics are the 3-point and 4-point correlation functions, whose Fourier transforms are the matter Bispectrum

$$B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) = \left\langle \tilde{\delta}(\mathbf{k}_1)\tilde{\delta}(\mathbf{k}_2)\tilde{\delta}(\mathbf{k}_3) \right\rangle \tag{1.24}$$

and Trispectrum

$$T(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4) = \left\langle \tilde{\delta}(\mathbf{k}_1)\tilde{\delta}(\mathbf{k}_2)\tilde{\delta}(\mathbf{k}_3)\tilde{\delta}(\mathbf{k}_4) \right\rangle , \tag{1.25}$$

respectively. In principle the n-point correlation function can be calculated for arbitrary n, but increasing computational cost and vanishing statistical significance as n increases makes doing so unnecessary for most cosmological applications.

1.4 Observations of ΛCDM

The ΛCDM model is remarkable in its ability to match a multitude of independent observational probes with a small number of parameters. Probes such as the distance ladder, the cosmic microwave background, redshift space distortions, baryon acoustic oscillations, and the shape of the power spectrum all provide unique constraints on the model parameters. In this section I summarize some of the key measurements and observations that lead to the ΛCDM model becoming the most favoured.

1.4.1 The distance-redshift relation

A photon emitted by a galaxy that is moving radially relative to an observer will have its wavelength shifted via the Doppler effect

$$\lambda_{obs} = \frac{c + v_r}{c}\lambda_{emit}, \tag{1.26}$$

where v_r is the galaxy's radial velocity relative to the observer and λ_{emit} is the wavelength of the photon when it was emitted. This shifting of wavelengths is commonly expressed as a redshift z, where

$$1 + z = \frac{\lambda_{obs}}{\lambda_{emit}}. \tag{1.27}$$

The radial velocity v_r in equation 1.26 can be split into two components: the Hubble flow caused by the expansion of the universe, and the peculiar velocity which is the galaxy's deviation from the Hubble flow. Since the velocity from the Hubble flow is linearly related to the distance to the galaxy, v_r can be rewritten as

$$v_r = H_0 r + v_{pec,r}, \tag{1.28}$$

where r is the line of sight distance from the observer to the galaxy. Assuming the peculiar velocity is negligible compared to the Hubble flow, the redshift z of the photon can be related to the scale factor at the time of emission by

$$a = \frac{1}{1 + z} \tag{1.29}$$

as well as the recession velocity and by extension the distance to the source

$$r = \frac{zc}{H_0}. \tag{1.30}$$

Typical peculiar velocities are on the order of hundreds of km/s, so given $H_0 \sim 70\ kms^{-1}Mpc^{-1}$, the approximation of negligible peculiar velocity holds for objects more distant than tens of Mpc.

1.4.2 Defining distances in cosmology

It is common in cosmology to use comoving coordinates for distance measurements, which are a coordinate system that expands along with the Hubble flow. Two points in space will remain the same comoving distance apart as the scale factor increases with time. The comoving distance between two points d_c is related to the proper distance

$$d_p = a \cdot d_c, \tag{1.31}$$

where d_p is the proper distance, which is the distance between the two points as would be measured by a ruler. The comoving distance between two points depends on the expansion history of the universe, introducing a factor of H_0 into its calculation. While multiple probes have determined the value of H_0 at high precision, there is significant disagreement between the values of H_0 measured by these different probes (see sec. 1.5). The choice of H_0 would then greatly affect the comoving distance calculated. To avoid this, it is common to report distances in units of $h^{-1}Mpc$, where $h = H_0/(100\ kms^{-1}Mpc^{-1})$. This way, distances can be reported with no assumed value of h, allowing the distance in Mpc to be calculated for any desired H_0. In this thesis, distances will typically be reported in comoving $h^{-1}Mpc$ units, and any departure from this will be explicitly stated.

When measuring distance to an object at redshift z along the line of sight, determining the comoving distance requires an integration over the expansion history between the time of emission and observation

$$d_c \equiv d_H \int_0^z \frac{dz'}{E(z')}, \tag{1.32}$$

where $d_H = c/H_0$ is the Hubble distance, and $E(z) = H(z)/H_0$ is the dimensionless Hubble parameter. The distance between two objects at the same redshift but separated by some angle on the sky $d\theta$ is $d_M d\theta$, where d_M is the transverse comoving distance

$$d_M = \begin{cases} d_H \frac{a}{\sqrt{\Omega_k}} \sinh\left[\sqrt{\Omega_k}\frac{d_c}{d_H}\right] & \Omega_k > 0 \\ d_c & \Omega_k = 0 \\ d_H \frac{1}{\sqrt{|\Omega_k|}} \sin\left[\sqrt{|\Omega_k|}\frac{d_c}{d_H}\right] & \Omega_k < 0 \end{cases} \tag{1.33}$$

This can then be used to find the angular diameter distance d_a of the object, which is the ratio between its transverse size to its angular size

$$d_a = \frac{d_M}{1+z}. \tag{1.34}$$

The angular diameter distance is of particular interest for observing standard rulers such as the Baryon Acoustic Oscillation (BAO) scale (see sec. 1.4).

The final distance of interest is the luminosity distance

$$d_L = \sqrt{\frac{L}{2\pi S}} = (1+z)d_M, \tag{1.35}$$

where L is the intrinsic luminosity of an object and S is the bolometric flux of that object measured by an observer. This distance measure is useful when determining distances to objects of known luminosity such as type Ia supernovae.

1.4.3 The distance ladder

Making use of the distance-redshift relation requires calibration by an accurate measurement of H_0. This can be done by using a distance ladder. A distance ladder works by first establishing a distance-luminosity relation to some nearby astrophysical objects through geometric means, then using that relation to determine the distance-luminosity relation of more luminous objects that are farther away. The first rung of the ladder is commonly established by using parallax to measure the distance to nearby Cepheid variable stars. The total luminosity of a Cepheid is related to the period at which its luminosity pulsates, so the distance to such stars can be calculated from a measurement of the photon flux observed from them. This is used to measure the distance to more distant Cepheids that are too far to have their distances measured using parallax. The final rung of the ladder is type Ia supernovae, caused when white dwarf stars accrete enough matter from a companion star to overcome the electron degeneracy pressure keeping them from collapsing. These supernovae are extraordinarily bright and have relatively similar luminosities to one another, making them ideal for measuring the distance-redshift relation out to redshift $z \sim 1$. Measurements of type Ia supernovae provided the first evidence for the accelerated expansion of the universe, which is now believed to be driven by dark energy [58, 57]. Currently, the most precise measurement of H_0 using the distance ladder comes from the SH0ES collaboration, who report a value of $H_0 = 73.04 \pm 1.04\ kms^{-1}Mpc^{-1}$ [59].

1.4.4 The cosmic microwave background

While the distance ladder can provide tight constraints on the expansion history of the universe, precise measurements of cosmological parameters can also be made from the cosmic microwave background in the early universe. When the universe was only a few

hundred thousand years old, it consisted of a hot, densely-packed sea of charged particles and photons. The photons in this hot ionized plasma were coupled strongly to the charged particles, meaning the average photon did not travel far between interactions with matter. These constant photon interactions also exerted a pressure on the plasma that opposed gravity, preventing the growth of structure. The universe was effectively opaque until it expanded and cooled enough for the charged particles to combine together into neutral elements during an epoch known as reionization. The photons interacted much less strongly with the neutral matter particles and began free-streaming. These photons are still observable today as the Cosmic Microwave Background (CMB) being emitted from all directions. These photons follow a blackbody spectrum with temperature $T_{CMB} = 2.72548 \pm 0.00057K$ [30].

The CMB provides a snapshot of the distribution of matter in the universe at the time of recombination, meaning plenty of cosmological information can be inferred from its measurements. The CMB on large scales is homogeneous and isotropic, providing a test of the cosmological principle [4]. There are small anisotropies in the CMB caused by small perturbations in the density field at the time of recombination. These tiny perturbations (called primary anisotropies) became the seeds of large scale structures such as clusters and voids. The CMB photons also interact with post-recombination matter via gravitational lensing [12], the Sachs-Wolfe effect [62], and the Sunyaev-Zel'dovich effect [70, 68, 69].

Since the primary CMB anisotropies are correlated with the early matter anisotropies, the CMB temperature power spectrum could be used to measure the matter power spectrum at recombination (see fig. 1.1). These initial perturbations are driven by Baryon Acoustic Oscillation (BAO)s. Before recombination, the balance between gravity and radiation pressure created acoustic standing waves in the density field. At the moment of recombination, these BAOs were frozen in place and the existing anisotropies were able to grow in the absence of the photon pressure. The size of the largest BAOs is determined by how far the acoustic oscillations were able to travel over the lifetime of the universe before recombination. These large scale BAOs create a distinctive peak in the correlation functions of both the CMB and the late-time density field. The position of this BAO peak and successive smaller BAO peaks can be used to provide strong constraints on cosmological parameters at different points during the universe's expansion history.

1.4.5 Spectroscopic galaxy surveys

Spectroscopic galaxy surveys provide information from the universe somewhere in between the late and early times covered by the distance ladder and CMB, respectively. These

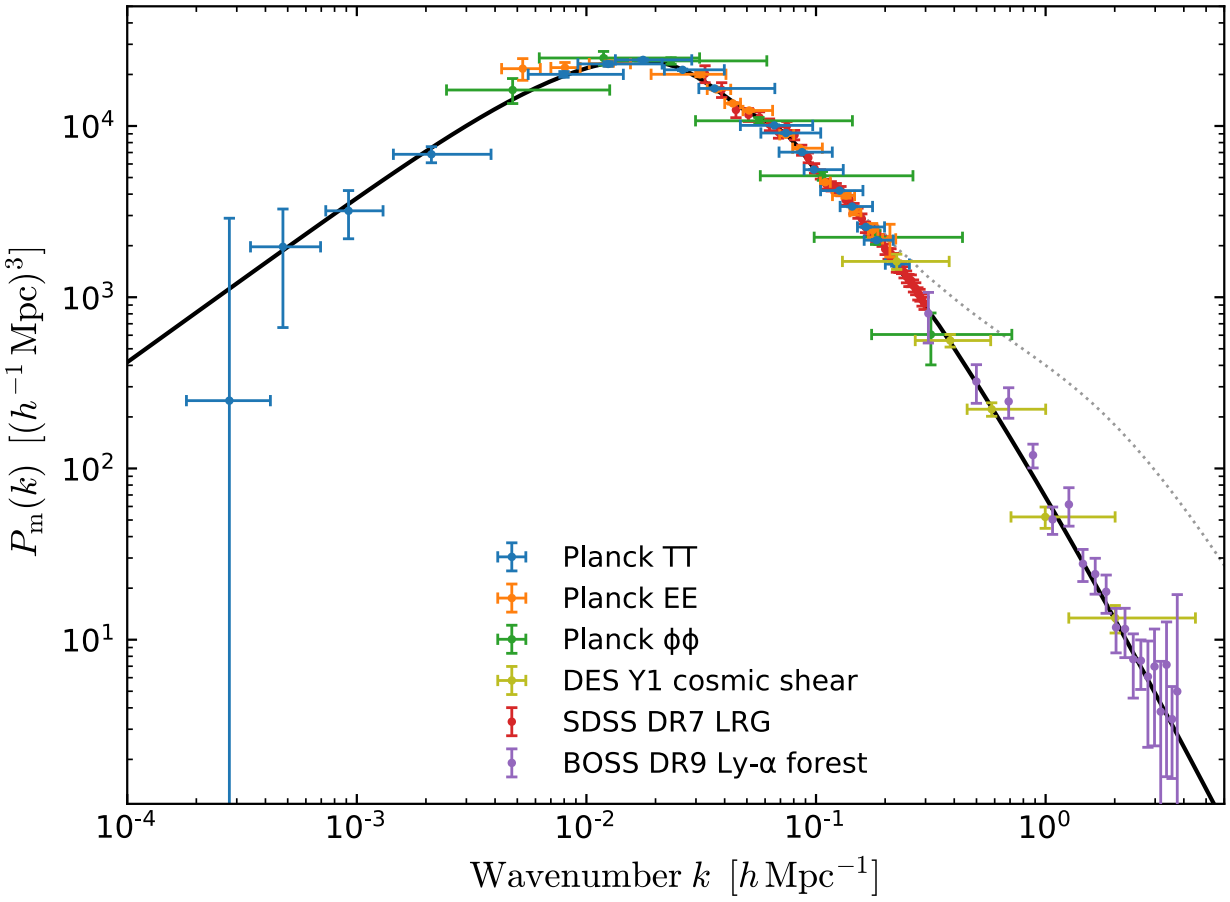

Figure 1.1: Linear-theory matter power spectrum at $z = 0$ inferred from different cosmological probes. The solid black line shows the predicted linear power spectrum from the ΛCDM model, while the dotted line shows the nonlinear power spectrum. The ΛCDM model is remarkable in its ability to explain results from such a wide variety of independent probes. Figure sourced from [3].

surveys record the angular positions and emission spectra of a large number of galaxies. By observing the redshift of certain emission lines from the galaxies, the distance-redshift relation can be used to construct a 3D map of the patch of the universe covered by the survey. From this map, features in the clustering statistics can be used to extract information about cosmological parameters.

One such feature is the BAO scale. Just as the BAO leaves an imprint on the temperature of the CMB, it also leaves that same imprint on the clustering of galaxies [26]. The scale of the BAO (also known as the sound horizon) acts as a standard ruler with a known length, determined as

$$r_s = \int_z^\infty \frac{dt}{H(z)} c_s(t), \tag{1.36}$$

where c_s is the speed of sound in the photon-baryon fluid before recombination. The BAO peak allows one to extract information about the expansion history independent of the distance ladder and the CMB, providing complementary constraints for various cosmological parameters.

Galaxy surveys also contain information about the growth of structure encoded within their peculiar velocities. When converting measured galaxy redshifts to radial distances assuming the Hubble flow is the sole cause of the recession velocity, the peculiar velocities of these galaxies cause their apparent positions to be shifted along the line of sight relative to their true positions. These shifts are known as redshift space distortions. The peculiar velocities of galaxies depends on the gravitational effect from their local surroundings, so measuring the degree to which galaxy positions are distorted reveals information about the matter present in the surrounding area.

Further information can be gleamed from the shape of the matter power spectrum. The initial post-inflation power spectrum was a power law of the form $P(k) \propto k^{n_s}$. When the universe was dominated by radiation, modes of the power spectrum did not grow. However, modes larger than the horizon scale (the largest distance at which two points in the universe could be in causal contact with one another) were able to grow similarly to how they would during the matter dominated era. This suppressed the small-scale power spectrum, with smaller scales being suppressed more due to entering the horizon earlier. Once the universe became matter dominated, all modes of the power spectrum became free to grow at the same rate, effectively freezing the shape of the power spectrum. The turnover scale of the power spectrum can then be used to constrain the ratio of matter to radiation.

1.4.6 Galaxy lensing surveys

Another type of survey that can provide complementary cosmological information is a galaxy lensing survey. In accordance with General Relativity, matter distorts spacetime around it. Any light passing through this distorted spacetime will have its path bent analogous to how light refracts when passing through a lens. In the case of light from a distant galaxy being lensed by a mass between it and the observer, this causes the galaxy to appear magnified and distorted compared to its true shape. When the galaxy's light passes in close proximity to a very massive lens, its shape is significantly distorted, often forming an arc or a ring shape around the lens. This phenomenon is known as strong gravitational lensing. If the source galaxy and the lens are not closely aligned, then the distortion is much more subtle and a statistical analysis of many background galaxies is required to infer the properties of the lens. This is known as weak gravitational lensing.

The basis for weak gravitational lensing analysis is the assumption that there is no preferred orientation for galaxies in the universe. Lensing surveys measure the position and alignment (also called shear) of distant galaxies in the sky. Any detection of a preferred orientation in some region would be caused by the presence of a foreground lens, allowing the matter distribution within the foreground to be inferred. There are three two-point statistics commonly used in lensing analyses: the galaxy-galaxy correlation, capturing the correlation between distant galaxy positions; the shear-shear correlation, capturing the correlation between the galaxy distortions; and the galaxy-shear cross-correlation, relating the positions of the observed galaxies to how much they are distorted.

The results are primarily sensitive to the cosmological parameters Ω_m and σ_8. While Ω_m parameterizes the total amount of matter in the universe, σ_8 is a measure of how tightly clustered that matter is; if one was to measure the overdensity of the universe within a sphere of radius 8 $h^{-1}Mpc$, one would expect to measure a value of $0 \pm \sigma_8$. Weak lensing provides a degenerate measurement of these two parameters which is commonly reported as a new parameter S_8:

$$S_8 \equiv \sigma_8 \sqrt{\frac{\Omega_m}{0.3}}. \tag{1.37}$$

The current best constraints on S_8 come from the Dark Energy Survey year 3 results, which reports a value of $S_8 = 0.776 \pm 0.017$ [1].

1.4.7 The power spectrum covariance matrix

When testing a cosmological model against survey data, it is necessary to compute a covariance matrix to test the likelihood of that model describing the data. The likelihood

of measuring an observed matter power spectrum $P_{measured}$ given some predicted matter power spectrum from a model P_{model} is given by

$$\mathcal{L} = \left(P_{measured}(k) - P_{model}(k)\right) C^{-1} \left(P_{measured}(k) - P_{model}(k)\right), \tag{1.38}$$

where C^{-1} is the inverse of the covariance matrix of the matter power spectrum

$$C(k_i, k_j) = \langle P(k_i) P(k_j) \rangle - \langle P(k_i) \rangle \langle P(k_j) \rangle . \tag{1.39}$$

The covariance matrix is a 4-point clustering statistic. In the case of δ being a Gaussian random field, this means that the covariance matrix can be calculated directly from the power spectrum

$$C_{Gauss}(k_i, k_j) \equiv \frac{1}{V} \frac{(2\pi)^3}{V_{k_i}} 2P(k_i)^2 \delta^K_{ij}, \tag{1.40}$$

where V is the survey volume, V_{k_i} is the volume of the i^{th} spherical shell in k space, and δ^K_{ij} is the Kronecker delta function. In this case the covariance C_{Gauss} would be diagonal, reflecting the fact that the modes of a Gaussian random field are independent of each other. Once nonlinear evolution has occurred and modes of the power spectrum become correlated, the off diagonal terms of the covariance become nonzero. In this case, it has been established that the covariance matrix picks up a contribution from the matter trispectrum [28, 49, 64, 71]

$$C(k_i, k_j) = C_{Gauss}(k_i, k_j) + \frac{1}{V} \overline{T}(k_i, k_j), \tag{1.41}$$

were $\overline{T}$ is the bin-averaged trispectrum

$$\overline{T}(k_i, k_j) = \int_{|\mathbf{k}| \in k_i} \frac{d^3\mathbf{k}}{V_{k_i}} \int_{|\mathbf{k}'| \in k_j} \frac{d^3\mathbf{k}'}{V_{k_j}} T(\mathbf{k}, -\mathbf{k}, \mathbf{k}', -\mathbf{k}'). \tag{1.42}$$

When analyzing a galaxy survey, one wants to minimize the covariance matrix in order to get the tightest constraints possible on the cosmological parameters. This entails trying to minimize the contributions of both cosmic variance and shot noise. Cosmic variance comes from the fact that only a finite number of modes can be measured from a survey, with larger volume surveys having a higher number of modes. If there are too few modes at a scale of interest, their sampling will be noisy and not accurately represent the variance of those modes in the universe as a whole. The shot noise component arises from the fact that the sample of observed galaxies are Poisson sampled from the underlying density field, adding a $P_{shotnoise} = \frac{1}{n}$ term to the power, where n is the number density of galaxies. For a finite number of galaxies in a survey, increasing the survey volume will decrease cosmic

variance and increase the shot noise. One can minimize the variance arising from these two sources by minimizing

$$\frac{\sigma_P}{P} = \frac{nP+1}{nP}, \tag{1.43}$$

where P is the estimated magnitude of the power spectrum at a scale of interest for the survey [28]. Equation 1.43 can be minimized by assigning weights to the galaxies in the survey that depends on the number density in its neighbourhood. These weights, denoted as w_{FKP}, are

$$w_{FKP} = \frac{1}{1+n(z)P}, \tag{1.44}$$

where $n(z)$ is allowed to take on different values depending on the redshift distribution of the data.

1.5 Tensions in the ΛCDM model

While the ΛCDM model performs extraordinarily well in matching current observations, there are tensions in the measured values of some of the cosmological parameters measured by different probes.

1.5.1 H_0 tension

The most significant tension in the ΛCDM model is between the value of H_0 determined by early universe measurements (e.g. CMB) and late time measurements of expansion rate (e.g. type Ia supernovae). The most recent measurements of the CMB from the Planck collaboration find a value of $H_0 = 67.27 \pm 0.60\ km\ s^{-1}Mpc^{-1}$ [4], while measurements from type Ia supernovae by the SH0ES collaboration prefer a higher value of $H_0 = 73.04 \pm 1.04\ km\ s^{-1}Mpc^{-1}$ [59]. These values are in $\sim 5\sigma$ tension with one another, with other early and late time probes of H_0 tending to agree with Planck or SH0ES, respectively (figure 1.2). Extensive searches for unknown systematic errors have been conducted but have so far been unable to relieve the tension. The tension could possibly be relieved by adding new physics to the ΛCDM model, such as modified gravity, early dark energy, or violation of the cosmological principle. A detailed review of proposed modifications that would alleviate the H_0 tension are provided in [23], though none of the proposals are currently favoured over ΛCDM for current observations.

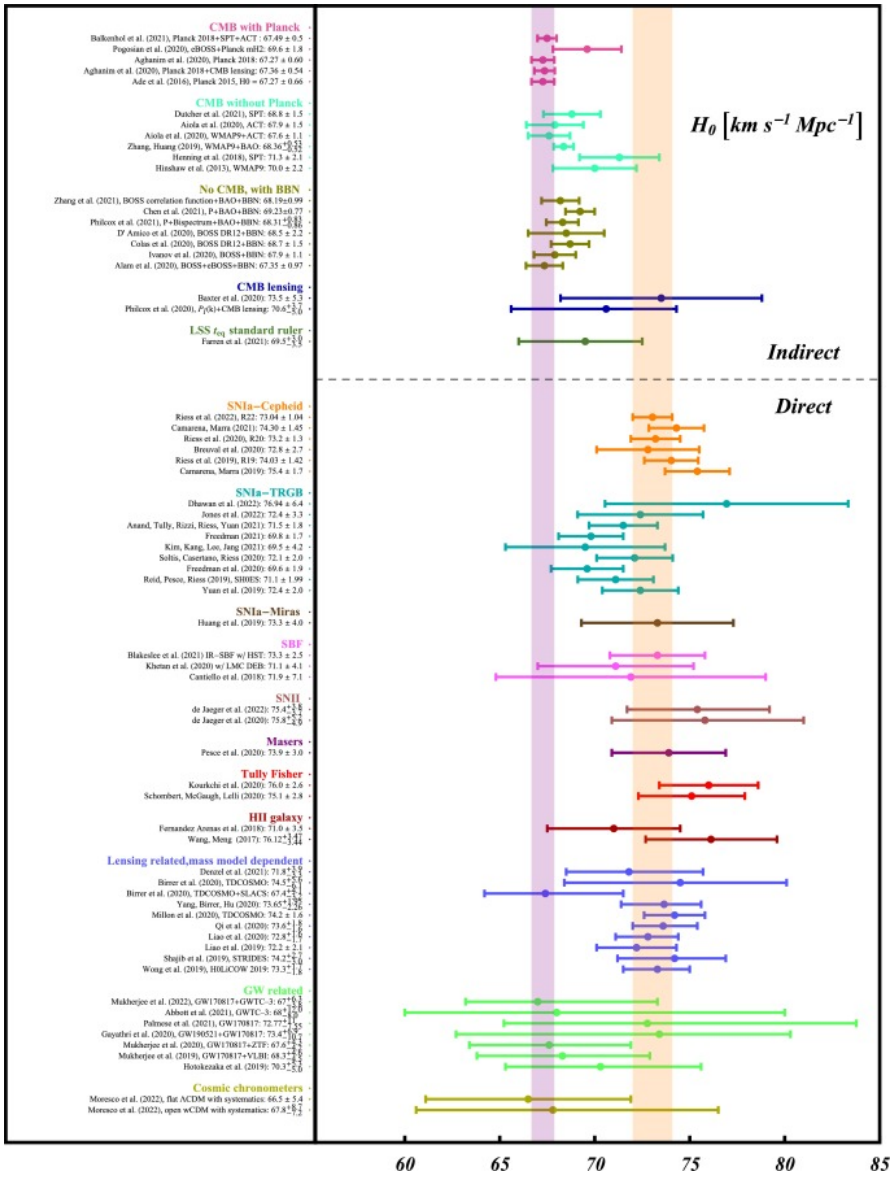

Figure 1.2: 68% constraints on H_0 measurements from different cosmological probes. The measurements are divided into direct probes of H_0 via the expansion rate and indirect probes that assume a particular cosmological model. The purple band shows constraints from Planck 2018 [4], and the orange band shows constraints from SH0ES [59]. Figure sourced from [27].

1.5.2 S_8 tension

There also exists a tension in the value of $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$ which captures information about matter clustering. This value can be obtained from early universe measurements using the CMB or late universe measurements using probes like weak lensing surveys (figure 1.3). While the S_8 tension measured from these two probes is only at the $\sim 2-3\sigma$ significance level, it is troublesome because many of the proposed solutions to either the H_0 tension or S_8 tension exacerbates the other.

1.6 N-body simulations

A cosmological N-body simulation is a simulation of a system of particles within a patch of the universe that interact via gravity. The proper volume of the simulation also expands (or contracts) with time in accordance with the Friedmann equation. In simulations designed to study large scale structure, each particle within the simulation typically represents a clump of dark matter. Since most of the matter in the universe is cold dark matter, tracking how the dark matter density evolves models the formation and evolution of large scale structure.

1.6.1 Basics

Each N-body simulation begins at some initial time that corresponds to an initial redshift z_{init}. The conditions of the simulation at this time are provided by the user in the form of an initial linear power spectrum and a random seed that sets the phases of the power spectrum modes in the simulation. These are used to generate an initial catalog of particle positions corresponding to the power spectrum amplitudes and phases provided. The evolution of the N-body simulations can then be broken up into discrete time steps. Given the particle masses and position at time t, the acceleration of each particle due to the net force of gravity acting on it can be calculated. The accelerations are then used to update the particles' velocities at the next time step $t + \Delta t$, and these velocities are in turn used to update the particles' positions. This process is repeated for a number of time steps until some desired output time (or, equivalently, and output redshift) is reached. Figure 1.4 shows the overdensity field from a simulation at different output redshifts. As time passes, regions that began as small overdensities evolve into haloes as they accumulate surrounding matter, and regions that began underdense empty out as their matter is drawn to denser regions of the simulation. The power spectrum from an N-body simulation contains the

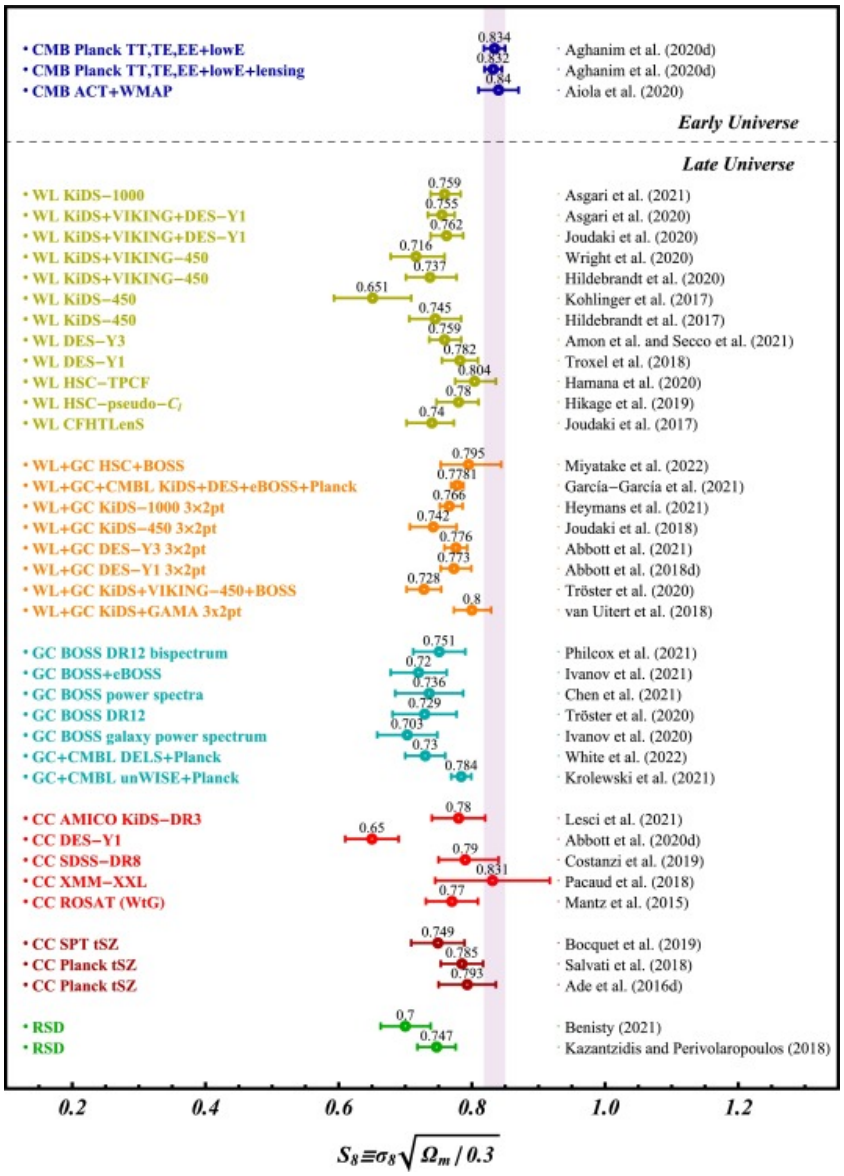

Figure 1.3: 68% constraints on S_8 measurements from different cosmological probes. The measurements are divided into early and late universe probes. The purple band shows constraints from Planck 2018 [4]. Figure sourced from [27].

effects of nonlinear growth of structure, resulting in an increase in power over the prediction from linear theory at scales of $k > 0.3\ hMpc^{-1}$ (see figure 1.5).

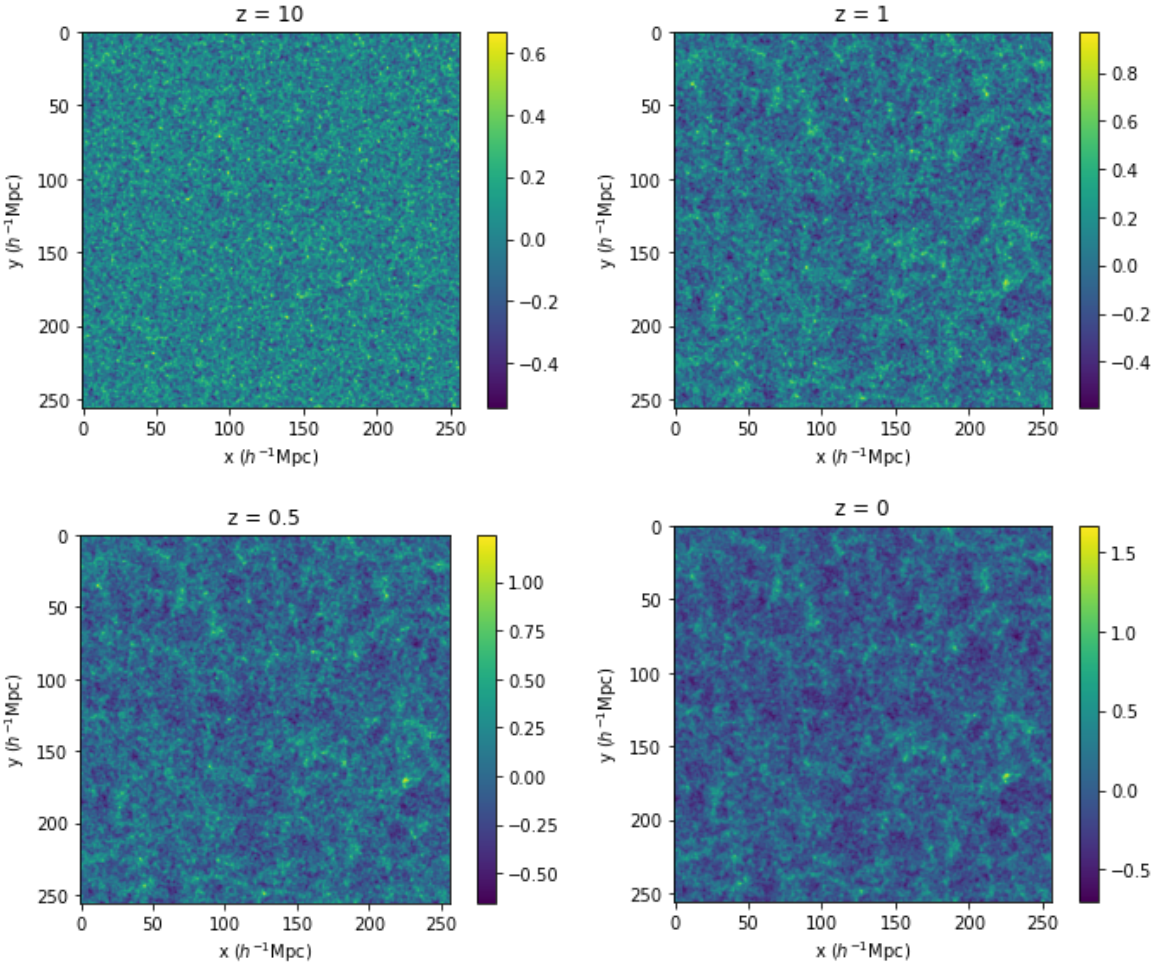

Figure 1.4: Evolution of the overdensity field within an N-body simulation at $z = 10$ (top left), $z = 1$ (top right), $z = 0.5$ (bottom left), and $z = 0$ (bottom right). The overdensity field is computed on a $(256)^3$ grid and projected onto the $x - y$ plane.

The problem with directly computing the exact force acting on each particle is that it is extremely computationally expensive. With the computational complexity scaling as $\mathcal{O}(N^2)$, simulations containing $N \sim 10^7$ particles are already bordering on prohibitively time-consuming on modern hardware [76]. Thus, it is useful to introduce approximations that can vastly decrease the computational time in exchange for slight decreases in overall

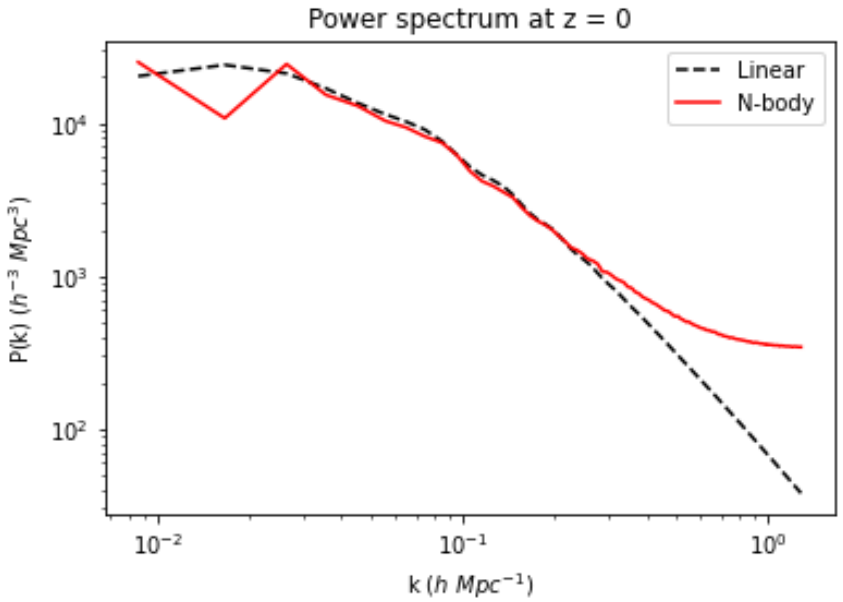

Figure 1.5: Present-day power spectrum within a $L = 625\ h^{-1}Mpc$ cube as predicted by linear theory and measured from an N-body simulation. The contribution to the power from non-linear modes becomes dominant for $k > 0.3\ hMpc^{-1}$

.

accuracy of the simulation. One such approximation is the Particle Mesh (PM) method.

In the PM method, rather than computing the force due to gravity for each pair of particles, the gravitational potential of the entire simulation is computed on a 3d grid. This is done by first computing the overdensity field on the grid and then solving the Poisson equation using a Fast Fourier Transform (FFT). This gravitational potential mesh can then be used to compute the force acting on each particle. Short range interactions between particles in this method are effectively softened since two particles within the same cell as each other don't contribute to each other's net force. The complexity of a PM simulation scales as $\mathcal{O}(N_{mesh} log(N_{mesh}))$, where N_{mesh} is the number of cells in the mesh grid. However, this mesh grid treatment of the gravitational potential is a poor approximation on scales of up to several grid spacings, resulting in reduced accuracy compared to full N-body methods [76].

Another possible approximation is the tree method, where the simulation volume is divided into cubic cells. The forces caused by particles within each cell are computed exactly, while the forces from particles in distant cells are treated as a single large particle located at that cell's centre of mass. The cells in denser regions of the simulation are made smaller to better capture the small-scale clustering. For optimal cell sizes, the complexity of tree method simulations scales as $\mathcal{O}(N\ log(N))$.

1.6.2 COLA method

The COmoving Lagrangian Acceleration (COLA) method is an N-body integration scheme that can significantly reduce the accuracy trade-off incurred by using the PM force scheme [74]. Consider the equation of motion for CMB particles in an expanding universe written in a Lagrangian coordinate system [37]

$$\frac{d^2\mathbf{\Psi}}{d\tau^2} + \mathcal{H}(\tau)\frac{d\mathbf{\Psi}}{d\tau} + \nabla\Phi = 0, \tag{1.45}$$

where Φ is the gravitational potential, τ is the conformal time, and $\mathcal{H} \equiv \frac{dlna}{d\tau}$ is the conformal Hubble parameter. In a Lagrangian frame of reference, the coordinate system moves along with each of the particles in the system, meaning a particle's Lagrangian position $\mathbf{q}$ does not change with time. The particle's kinematics are instead captured in the displacement vector $\mathbf{\Psi}$ which captures how the Lagrangian frame of reference is evolving with time

$$\mathbf{x}(\tau) = \mathbf{q} + \mathbf{\Psi}(\mathbf{q}, \tau). \tag{1.46}$$

Thus, solving for $\mathbf{\Psi}$ yields the time evolution of the particle positions. Lagrangian Perturbation Theory (LPT) seeks to find a perturbative solution for $\mathbf{\Psi}$ [14]

$$\mathbf{\Psi}(\mathbf{q},\tau) = \mathbf{\Psi}^{(1)}(\mathbf{q},\tau) + \mathbf{\Psi}^{(2)}(\mathbf{q},\tau) + \mathbf{\Psi}^{(3)}(\mathbf{q},\tau) + \cdots . \tag{1.47}$$

It is mathematically convenient to define Lagrangian potentials $\nabla_q \phi^{(i)} = \mathbf{\Psi}^{(i)}$. The second order solution of equation 1.46 then takes the form

$$\mathbf{x}(\tau) = \mathbf{q} + D_1(\tau)\nabla_q \phi^{(1)} + D_2(\tau)\nabla_q \phi^{(2)}, \tag{1.48}$$

where $D_1(\tau)$ and $D_2(\tau)$ are the linear and second order growth factors, respectively. This second-order Lagrangian Perturbation Theory (2LPT) solution performs well in solving for large scale structure evolution in N-body simulations, but its accuracy falls off significantly outside of the linear regime. Since at early times ($a \lesssim 0.1$) the overdensity fluctuations in the universe are very small, 2LPT is well suited to determining the early evolution of particles in an N-body simulation. Once initial particle positions are generated at high redshift (typically $z \simeq 1000$), 2LPT is used to evolve them to an initial redshift z_{init} (typically $z_{init} \sim 9$), after which the N-body time-stepping code takes over to compute the non-linear evolution.

The COLA method takes advantage of the fact that the 2LPT solution is exactly known and uses PM methods to estimate the higher order terms of equation 1.46. In the COLA method, the observer is considered to be in a frame of reference comoving with 2LPT coordinates. Equation 1.45 can then be rewritten as

$$\frac{d^2\Psi_{res}}{d\tau^2} + \mathcal{H}(\tau)\frac{d\Psi_{res}}{d\tau} + \nabla\Phi + [2LPT\ solution] = 0, \tag{1.49}$$

where the $[2LPT\ solution]$ term is calculated using 2LPT. Since this frame of reference is accelerating, the 2LPT solution term can be interpreted as a fictitious force acting on the particles. This term is solved analytically, leaving the N-body code to calculate an estimate of the 3rd order and higher in overdensity.

1.6.3 Using L-PICOLA simulation code

In this thesis, the N-body simulations used were generated using the Lightcone-enabled Parallel Implementation of COmoving Lagrangian Acceleration (L-PICOLA) simulation code [37]. L-PICOLA is a parallelized implementation of the COLA method. When running the simulation in parallel across multiple CPUs, the simulation volume is divided

up into slabs along its x axis. Each CPU is assigned one slab for which it evolves its particles using the PM method. At the end of each timestep, particles that have moved outside of their CPU's slab are moved to the correct slab before the next timestep begins. Due to this parallelization, L-PICOLA can run simulations with large numbers of particles quickly while still accurately modeling nonlinear clustering.

To run a simulation, L-PICOLA requires three input files from the user: the run parameters, the linear power spectrum, and the output redshifts. The run parameters file contains the values of the cosmological parameters to be used, as well as the seed for the initial density perturbations, information on the number of CPUs being used, the memory allocated to each slab, and the file names of the input and output files. The linear power spectrum file contains values of the redshift 0 linear power spectrum $P_{lin}(k)$ for a wide range of k. At the start of the simulation, L-PICOLA scales this P_{lin} back to the initial redshift using the growth factor (equation 2.13) and generates initial particle positions that match this power spectrum. The output redshifts file contains a list of desired redshifts at which the simulation is to save an output, as well as the number of timesteps to take between each output. At each output redshift, each CPU saves a catalog containing the x, y, and z components of its slab's particle positions and velocities, with position being given in comoving $h^{-1}Mpc$ units and velocity given in km/s.

1.6.4 Processing pipeline: From particle catalog to power spectrum

Computing the power spectrum from an output simulation catalog was done using the `nbodykit` toolkit [36]. The basic processing pipeline I used is as follows:

1. Load the simulation output files into a single catalog.
2. Paint the catalog onto a 3d mesh grid with resolution equal to the PM resolution used by L-PICOLA.
3. Compute the 1d power spectrum from the FFT of the density field.

The output from this process is an object containing the power in each k bin, the central k value of each bin, and the number of modes in each bin. For all the power spectra computed in this work, the k bins were linearly spaced with width $\Delta k = 10\pi/625\ hMpc^{-1}$ in proper distance units.

When computing $P(k)$ from a mesh like this, one runs into significant aliasing issues on scales approaching the Nyquist frequency of the grid

$$k_{nyq} = \pi N_{mesh}/L_{box}, \tag{1.50}$$

where N_{mesh} is the number of mesh grid cells per side and L_{box} is the side length of the simulation. Aliasing is caused by small scale modes with $k > k_{nyq}$ that cannot be resolved by the mesh being mistakenly identified as modes with $k < k_{nyq}$, contributing to spurious measurement of the power spectrum in k bins near the Nyquist limit. This effect can be counteracted by using interlacing and interpolation methods [65, 40]. Interlacing involves computing the FFT on a second mesh grid shifted π/L_{box} relative to the first one. While this doubles the required computing time, the additional information results in a significant reduction in the aliasing of the power spectrum. The different interpolation methods change how the particles are assigned to a mesh grid; instead of each particle being assigned solely to the nearest grid cell, the particle is convolved with a resampler function that spreads the particle's contribution to nearby cells in a way that better captures the particle's position relative to the neighboring cells. These resamplers have the form

$$W(\mathbf{k}) = \left[\frac{\sin(\pi k_x/2k_{Nyq})\sin(\pi k_y/2k_{Nyq})\sin(\pi k_z/2k_{Nyq})}{(\pi k_x/2k_{Nyq})(\pi k_y/2k_{Nyq})(\pi k_z/2k_{Nyq})}\right]^p, \tag{1.51}$$

where k_x, k_y, k_z are the x, y, and z components of the $\mathbf{k}$ vector, and p is a positive integer determining the order of the resampler. The resampler must later be deconvolved from the power spectrum, but the result is also a marked decrease in aliasing for the cost of a fractionally higher computation time. Figure 1.6 demonstrates the effectiveness of these techniques in counteracting the effect of aliasing.

1.7 Thesis outline

This chapter has introduced the flat ΛCDM concordance model, its key observables, the relevant clustering statistics for this work, and the process by which N-body simulations can be used to study the growth of large-scale structure. Chapter 2 introduces the effect of super-sample modes on the covariance matrix estimator and provides a comparison of different methods of including this effect in N-body simulations. Chapter 3 evaluates the method of estimating the covariance matrix of a large volume survey using an ensemble of small volume simulations. Finally, Chapter 4 summarizes these results and discusses their potential use for analysis of next-generation surveys.

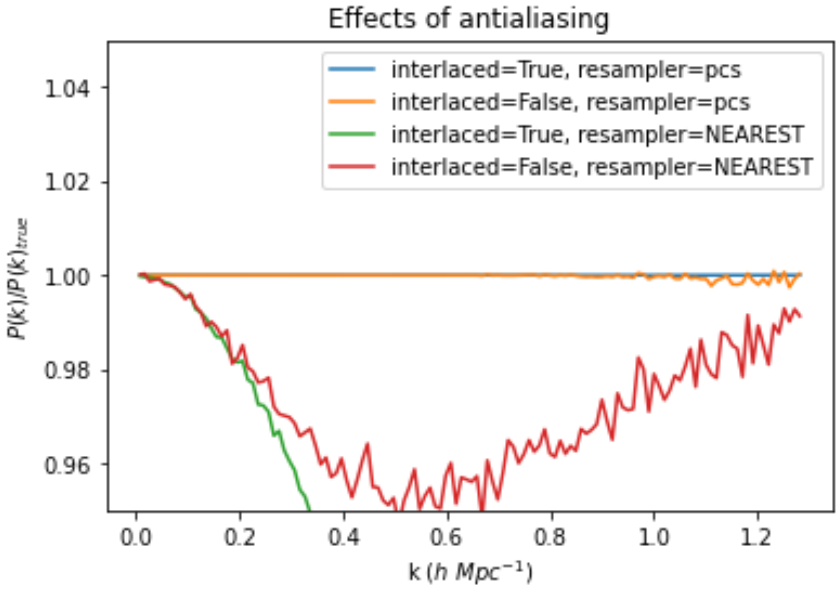

Figure 1.6: Effects of antialiasing on the power spectrum. The "true" power spectrum was computed using a much higher resolution mesh grid, effectively increasing the Nyquist limit. The "NEAREST" resampler assigns particles to the nearest grid cell, and the "pcs" resampler convolves the particle position with equation 1.51 using $p = 4$. The combination of the PCS resampler and interlacing technique allows the power spectrum to be recovered almost perfectly up to the Nyquist limit.

Chapter 2

Super-Sample Covariance

This chapter covers the effect of Super-Sample Covariance (SSC) on the covariance matrix estimator, and my investigation into how to best capture the super-sample effect in N-body simulations. In section 2.1, I discuss how modes of the power spectrum longer than the width of a survey or simulation volume modulates the density field within. Section 2.2 describes how these long-wavelength modes couple to modes within the simulation volume, contributing an additional term to the covariance matrix known as SSC. In section 2.3 I describe how SSC can be included in an ensemble of simulations by either splitting a large simulation into smaller sub-volumes or by running simulations in the presence of a non-zero background overdensity. In section 2.4, I demonstrate how the variance of the background mode of these simulations can be computed. In section 2.5, I present two methods for running separate universe simulations with non-zero average overdensity: the perturbed parameter approach introduced in [66], and the Spherical Collapse approach presented in [81]. In section 2.6, I show how separate universe simulations can be used to recover the effects of SSC either by using the addition or ensemble method [44]. In section 2.7, I discuss a number of subtleties involved in correctly running separate universe simulations and computing their power spectra. In section 2.9, I present my work in comparing the how well the different methods of running separate universe simulations and computing the SSC effect perform in recovering the correct covariance matrix. I also contextualize the results of each individual method's covariance calculation by comparing them to the results of previous studies.

2.1 Power spectrum response to larger-than-box modes

A galaxy survey constitutes a mapping of the matter density field (or more precisely, tracers of the field) within the survey volume V. The overdensity field within the survey can then be computed with $\overline{\rho}$ being the average density within the survey. However, since galaxy surveys can only cover a finite volume, the $\overline{\rho}$ estimated from the survey will be different from the average density of the entire universe. This introduces a question of whether to normalize δ (and thus, $P(k)$) relative to the mean density of the survey (the "local" mean), or relative to the mean density of the universe (the "global" mean). The overdensity normalized to the global and local mean density can be related by

$$\tilde{\delta}_{global} = \tilde{\delta}_{local}(1 + \delta_b), \tag{2.1}$$

where $\delta_b = \frac{\overline{\rho}_{local} - \overline{\rho}_{global}}{\overline{\rho}_{global}}$ is the average overdensity of the survey relative to the entire universe, also called the background overdensity. It is useful to then define $P_{local}(k)$ and $P_{global}(k)$ that correspond to whether the local or global mean densities are used in the normalization. These different normalizations are related by

$$P_{global}(k) = P_{local}(k)(1 + \delta_b)^2. \tag{2.2}$$

If one is interested in the global $P(k)$ as is common in large-scale structure surveys, the variance in the background density contributes greatly to the covariance matrix and thus must be accounted for. This results in a distinction between C_{local} and C_{global} depending on whether the local or global power spectrum is being considered.

2.2 Mode coupling

The normalization of the power spectrum is not the only way the variance of the background density affects the covariance of the power spectrum. In quasilinear and nonlinear regimes, non-linear gravitational evolution of large scale structure causes the Fourier modes of the power spectrum to become correlated with one another [64, 39, 19]. Any mode of the power spectrum corresponding to length scales undergoing nonlinear evolution are thus correlated with the other modes of the power spectrum, including very large modes that are themselves still within the linear regime. This can be readily visualized via the peak-background split [41, 18]. In figure 2.1, the presence of a long wavelength mode modulates the amplitude of the overdensity field enough to change the number of regions with density high enough to collapse under their own self-gravity. Since growth of structure in these

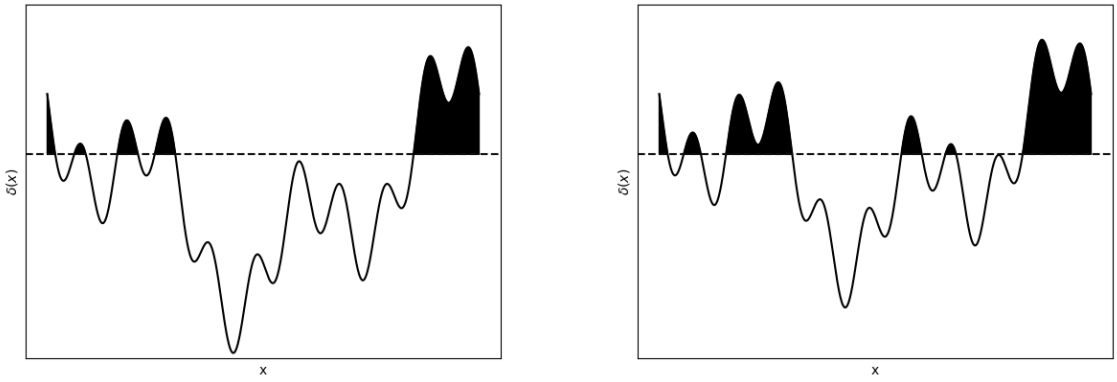

Figure 2.1: A demonstration of the peak-background split in the presence of a long wavelength mode. Shaded areas represent regions with sufficient overdensity to undergo collapse. In absence of a long wavelength mode (left panel), the overdensity crosses the critical density for collapse (dashed line) in fewer places than it would if such a mode were present (right panel).

regions is nonlinear and occurs on small scales, it is evident that large-scale modes are correlated to small scale modes.

In the case of a galaxy survey, the survey volume will contain parts of power spectrum modes with wavelengths greater than the width of the survey (called super-sample modes). While the super-sample modes themselves cannot be measured from the survey, they correlate nontrivially with modes that are present within the survey volume, thus contributing to the power spectrum covariance. This contribution is called Super-Sample Covariance (SSC). A theoretical model for the effects of SSC can be derived using either perturbation theory [21] or halo modeling [71, 44], yielding

$$C_{local}(k_i, k_j) = C_{Gauss}(k_i, k_j) + \frac{1}{V}\overline{T}(k_i, k_j) + \frac{676}{(21)^2}P(k_i)P(k_j)\sigma_b^2 \tag{2.3}$$

and

$$C_{global}(k_i, k_j) = C_{Gauss}(k_i, k_j) + \frac{1}{V}\overline{T}(k_i, k_j) + 16\left(\frac{17}{21}\right)^2 P(k_i)P(k_j)\sigma_b^2, \tag{2.4}$$

where σ_b^2 is the variance of the background density within the survey volume (see sec. 2.4).

2.3 SSC in simulations

Computing the covariance matrix analytically in the nonlinear regime has so far proven extraordinarily challenging beyond low-order perturbations. A simpler solution is to instead estimate the covariance matrix using a large number of N-body simulations whose geometry and cosmology match that of the survey whose covariance is desired. An ensemble consisting of N_s of these simulations provides N_s samples of the power spectrum within the survey volume. These power spectrum samples can be used to compute a brute-force estimate of the covariance matrix

$$C(k_i, k_j) = \frac{1}{N_s - 1} \sum_{m=1}^{N_s} \left[P_m(k_i) - \overline{P}(k_i)\right] \left[P_m(k_j) - \overline{P}(k_j)\right], \tag{2.5}$$

where $P_m(k_i)$ is the i^{th} bin of the m^{th} power spectrum sample, and $\overline{P}(k_i)$ is the average power among all N_s samples in the i^{th} bin.

To get an accurate estimate of a covariance matrix from simulations, they need to include the effects of SSC. The simplest way to do this is by running a simulation with a very large volume and dividing it into subsamples. Each of these sub-volumes will have evolved in the presence of modes too large to fit within the sub-volume but still smaller than the large volume. While this technique has the advantage of being straightforward to implement, it comes with some major drawbacks. Suppose one wants to run a large volume mock to be split into 8 smaller sub-volumes with a desired average number density of particles $\overline{n}$. The large volume mock must then contain 8 times as many particles as would be contained in a sub-volume with number density $\overline{n}$. Since the computation time of N-body simulations scales as $\mathcal{O}(N^2)$ (or $\mathcal{O}(NlogN)$ for PM codes), it will take more time to run the one large mock than it would to run the eight smaller mocks. Another issue is that the SSC effect within these sub-volumes will be underestimated compared to what would be seen in surveys. No matter how large of a simulation the sub-volumes are drawn from, there still exist modes of the power spectrum even larger than the width of the box whose effects are missing. This requires the box volume to be sufficiently large such that the SSC effect from the missing modes is negligible compared to the effect from the modes within the box.

Alternatively, the SSC effect of a background mode larger than a simulation or survey volume can be well approximated by the presence of a non-zero background overdensity δ_b corresponding to a "DC mode" (analogous to a direct current applied to an electrical circuit). This means we can run an ensemble of simulations with non-zero δ_b to capture the SSC effect. When running N-body simulations, it is typical to use a cubic simulation

volume and employ Periodic Boundary Conditions (PBCs) to handle cases when particles would leave the boundaries of the simulation. If a particle would pass the boundary of the simulation, its position is shifted by the side length of the box such that it is within the simulation box again. Unfortunately, employing PBCs implicitly sets the power in all modes larger than the simulation box to zero, meaning the covariance calculated from these simulation will be missing the SSC effect. Due to the strength of the SSC effect at quasilinear and nonlinear scales, it must be reintroduced into the simulations if a satisfactory estimate of the covariance in these regimes is to be obtained.

2.4 Calculating the variance of the background density

If the background mode is within the linear regime, its amplitude is Gaussian distributed with zero mean and has variance σ_b^2 that can be computed by integrating the linear power spectrum $P_{lin}(k)$:

$$\sigma_b^2 \equiv \frac{1}{2\pi^2} \int |\tilde{W}^2(k\mathbf{x})| P_{Lin}(k) k^2 dk, \tag{2.6}$$

where $\tilde{W}^2(k\mathbf{x})$ is the window function of the survey or simulation. Running an ensemble of simulations in a manner that can account for variance in the background overdensity would then reintroduce the effects of SSC.

For a cubic N-body simulation, the window function in equation 2.6 would be the Fourier transform of the cubic window representative of the simulation volume. However, this calculation can be simplified by instead using a window corresponding to the Fourier transform of a spherical top-hat function

$$\tilde{W}(kR) = \frac{3\left[\sin(kR) - (kR)\cos(kR)\right]}{(kR)^3}, \tag{2.7}$$

where R is the radius of the sphere in real space. If R is chosen such that the sphere has the same volume as the simulation (i.e. $\frac{4}{3}\pi R^3 = L^3$, where L is the side length of the simulation box), then the calculated σ_b^2 will closely match that calculated from a cubic window. The validity of this claim is shown at the end of sec. 2.9.1.

Realistic galaxy redshift surveys do not have a simple cubic geometry. When using simulations to estimate the covariance of such surveys, it is convenient to run simulations with a cubic window function and then apply the survey window function to the resulting

overdensity field afterwards. The modes of the power spectrum within the survey window evolved in the presence of modes that are larger than the survey window but smaller than the cubic simulation window, meaning the covariance estimated from these simulations will already have part of the SSC contribution. To recover the covariance matrix with the correct SSC contribution, one only needs to add the SSC term arising from modes larger than the *simulation* volume. Thus, regardless of the form of the survey window, calculation of σ_b^2 should be done using the simulation window.

2.5 Running separate universe simulations

The problem to solve is now this: how can an N-body simulation be run with both PBCs and nonzero δ_b? This problem has been an active field of study for many years, even before being contextualized by its application in calculating SSC [35, 22, 66, 48, 33, 46, 32, 7, 71, 44, 77, 20, 75, 17, 81]. While there exist multiple methods to run a simulation with nonzero δ_b, the fundamental principle is the same in each: the change in the average density of the simulation can be interpreted as though the simulation was a Separate Universe (SU) with different cosmological parameters compared to the parameters of a fiducial background universe with zero background density. This work focuses on the methods presented in [66, 71, 44] and [81] for computing the SU parameters.

2.5.1 Perturbed parameter approach

The perturbed parameter approach follows the derivations used in [66, 71, 44]. In this approach, the derivation begins by considering δ_b to be a perturbation in the mean matter density $\overline{\rho}(t)$ of the simulation:

$$\overline{\rho}_{uni}(t)(1+\delta_b(t)) = \overline{\rho}_{box}(t), \tag{2.8}$$

where the subscripts *uni* and *box* refer to the parameters of the background universe and the SU parameters, respectively. From here, the mass density parameters of the background universe and the SU can be related by

$$\frac{\Omega_{m,0,uni}H_{0,uni}^2}{a_{uni}^3(t)}\,(1+\delta_b(t)) = \frac{\Omega_{m,0,box}H_{0,box}^2}{a_{box}^3(t)}. \tag{2.9}$$

At early times, the physical content of each SU should be the same. This means $\lim_{t\to 0} a_{uni} = a_{box}$, which allows equation 2.9 at early times to simplify to

$$\Omega_{m,0,uni}H_{0,uni}^2 = \Omega_{m,0,box}H_{0,box}^2, \tag{2.10}$$

which in turn implies

$$a_{box}(t) = \frac{a_{uni}}{(1+\delta_b(t))^{1/3}} \approx a_{uni}(t)\left(1 - \frac{\delta_b(t)}{3}\right). \tag{2.11}$$

Using the definition $H \equiv \dot{a}/a$, one can take the time derivative of equation 2.11 and divide both sides by $a_{box}(t)$ to yield

$$H_{box} = H_{uni} - \left(\frac{\dot{\delta}_b/3}{1-\delta_b/3}\right). \tag{2.12}$$

If δ_b is drawn from a large enough mode to evolve under linear growth, then $\delta_b = \frac{D(a)}{D(1)}\delta_{b,0}$ and $\dot{\delta}_b = \frac{\dot{D}(a)}{D(1)}\delta_{b,0}$, where $D(a)$ is the linear growth factor

$$D(a) = \frac{5\Omega_{m,0}H_0^2}{2}\frac{\dot{a}}{a}\int_0^a \frac{da'}{\dot{a'}^3} \tag{2.13}$$

normalized such that $lim_{a\to 0}D(a) = a$. Substituting equation 1.12 into 2.13 and taking the time derivative gives

$$\dot{D}(a) = \frac{\Omega_{m,0}H_0^2}{2\dot{a}a}\left[5 - \frac{3D}{a} - \frac{2\Omega_{k,0}D}{\Omega_{m,0}}\right]. \tag{2.14}$$

Using equations 2.9, 2.11, 2.12, 2.14 and 1.12, it can be shown that the box cosmological parameters are of the form

$$H_{0,box} = H_{0,uni}(1+\phi)^{-1} \tag{2.15}$$
$$\Omega_{m,box} = \Omega_{m,uni}(1+\phi)^2 \tag{2.16}$$
$$\Omega_{\Lambda,box} = \Omega_{\Lambda,uni}(1+\phi)^2 \tag{2.17}$$
$$\Omega_{k,box} = 1 - \Omega_{m,box} - \Omega_{\Lambda,box}, \tag{2.18}$$

where

$$\phi = \frac{5}{6}\frac{\Omega_m}{D(1)}\delta_b. \tag{2.19}$$

2.5.2 Spherical collapse approach

Another way to interpret the presence of a background mode in the SU simulations is to consider how the change in background density perturbs the curvature of space. This approach follows the methodology and derivations presented in [55, 81].

Consider the cosmology of two large spherical regions containing equal amounts of mass: one following the background cosmology with $\delta_b = 0$ and scale factor a_{uni}, and another with $\delta_b \neq 0$ and scale factor a_{box} as given in equation 2.11. Assuming the dark energy component of each region is negligible at early times, their early time Friedmann equations can be written in the form

$$\left(\frac{1}{H_{0,uni}}\frac{da_X}{dt}\right)^2 = \frac{\Omega_{m,uni}}{a_X} + \epsilon_X, \tag{2.20}$$

where the subscript X is a placeholder for the *uni* or *box* subscripts and labels which parameters vary between the two spheres, and the curvature term ϵ_X is equal to $\Omega_{k,uni}$ for the background sphere and can take any real value for the perturbed sphere. In the limit $t \to 0$, the solution to equation 2.20 has the form

$$a_X = \left(\frac{9\Omega_{m,uni}}{4}\right)^{1/3}(H_{0,uni}t)^{2/3} + \frac{3\epsilon_X}{20}\left(\frac{12}{\Omega_{m,uni}}\right)^{1/3}(H_{0,uni}t)^{4/3}. \tag{2.21}$$

Combining equations 2.11 and 2.21, δ_b in the limit of $t \to 0$ gives

$$\delta_b(t) = \frac{3}{5}\left(\frac{3}{2\Omega_{m,0,uni}}\right)^{2/3}(\Omega_{k,uni} - \epsilon_{box})(H_{0,uni}t)^{2/3}. \tag{2.22}$$

This early time density can be linked to the present day density $\delta_{b,0}$ by linear growth

$$\delta_{b,0} = \frac{3D(1)}{5\Omega_{m,0,uni}}(\Omega_{k,uni} - \epsilon_{box}). \tag{2.23}$$

For any given ϵ_{box}, a matching value of δ_b can be matched to it

$$\epsilon_{box} = \Omega_{k,uni} - \frac{5\Omega_{m,uni}\delta_b}{3D_0} = \Omega_{k,uni} - 2\phi, \tag{2.24}$$

where ϕ is as defined in equation 2.19. The cosmological parameters of the patch can be computed by solving the Friedmann equation

$$\left(\frac{1}{H_{0,uni}}\frac{da_{box}}{dt}\right)^2 = \frac{\Omega_{m,uni}}{a_{box}} + \epsilon_{box} + \Omega_{\Lambda,uni}a_{box}^2 \tag{2.25}$$

It is important to remember that here $\epsilon_{box} \neq 1 - \Omega_{m,uni} - \Omega_{\Lambda,uni}$. A useful renormalization of equation 2.25 comes from dividing both sides by $1 - 2\phi$ and defining a new set of

renormalized cosmological parameters

$$H_{0,box} = H_{0,uni}\sqrt{1-2\phi} \tag{2.26a}$$

$$\Omega_{m,box} = \frac{\Omega_{m,uni}}{1-2\phi} \tag{2.26b}$$

$$\Omega_{k,box} = \frac{\epsilon_{box}}{1-2\phi} \tag{2.26c}$$

$$\Omega_{\Lambda,box} = \frac{\Omega_{\Lambda,uni}}{1-2\phi}. \tag{2.26d}$$

Under this renormalization, the Friedmann equation becomes

$$\left(\frac{1}{H_{0,box}}\frac{da_{box}}{dt}\right)^2 = \frac{\Omega_{m,box}}{a_{box}} + \Omega_{k,box} + \Omega_{\Lambda,box}a_{box}^2. \tag{2.27}$$

Under this new renormalization, $\Omega_{k,box} = 1 - \Omega_{m,box} - \Omega_{\Lambda,box}$. Defining the cosmological parameters in this way is useful for use with N-body codes that infer the value of Ω_k from the input values of Ω_m and Ω_Λ as L-PICOLA does.

It is worthy of note that the parameter perturbations from the SC method are identical to the results of the derivations presented in [7] and [77]. This is due to the root of each method being the same: perturbing the FLRW metric (i.e. the curvature of space) within the simulation volume in the presence of a long wavelength background mode that is well approximated by a nonzero background overdensity.

2.6 Calculating the SSC effect

Once a method of translating from δ_b to cosmological parameters is chosen, the next step is to introduce the effects of SSC. In this section, I will describe the ensemble method and addition method of computing the SSC effect.

2.6.1 Ensemble method

One method of including the effects of SSC in the covariance of an ensemble of simulations is to draw the background overdensity of each simulation from a normal distribution with zero mean and variance given by equation 2.6. This background overdensity corresponds to a perturbation in the cosmological parameters of each simulation as prescribed by either the

perturbed parameter approach (see section 2.5.1) or the SC approach (see section 2.5.2). This variation in the background overdensity allows the effect of SSC to be effectively baked in to the ensemble. The covariance with the SSC effect can then be calculated as normal using equation 2.5.

2.6.2 Addition method

Another way to include the effects of SSC is to consider that, to first order, the covariance matrix with the SSC effect included can be approximated as

$$C_{SSC}(k_i, k_j) = C_{noSSC}(k_i, k_j) + \sigma_b^2 \frac{\partial P(k_i)}{\partial \delta_b} \frac{\partial P(k_j)}{\partial \delta_b}, \tag{2.28}$$

where C_{noSSC} is the covariance matrix without any SSC correction. The partial derivative term can be computed using finite differences of simulations with nonzero background overdensity. In principle, only two simulations are required to compute this derivative: one with positive δ_b, one with negative δ_b. In practice, the initial conditions of the simulations affects the response to δ_b, so it is useful to run several pairs of simulations with different initial conditions and compute the finite difference of the average power spectrum at each δ_b.

Since equation 2.28 is a first order approximation of the ensemble method, the addition method expected to have reduced accuracy. Also, if C_{noSSC} is being computed from an ensemble of simulations, the addition method requires a slightly higher number of simulations than the ensemble method to recover the full covariance to the same accuracy. However, the addition method has the advantage of splitting the covariance into a no-SSC and SSC term, allowing for clear visualization of how significant the SSC effect is at different scales.

2.7 Comparing apples to apples

When running SU simulations, there are a number of subtleties that must be accounted for to ensure that the resulting power spectra are accurate. Care needs to be taken with handling the input linear power spectrum $P_{lin}(k)$, the units of $h^{-1}Mpc$ when referring to distances within a SU simulation, and the binning of the power spectrum. When N-body code generates the initial particle positions in a simulation, they are distributed such that the initial power spectrum matches a user-specified linear power spectrum $P_{lin}(k)$. The

linear power spectrum given by the user is defined at $a = 1$, so the N-body code scales the linear power back to an early time (typically $a = 0.001$) using the linear growth factor

$$P_{lin}(k, a = 0.001) = \left(\frac{D(a = 0.001)}{D(a = 1)}\right)^2 P_{lin}(k, a = 1). \tag{2.29}$$

A SU simulation with nonzero δ_b will have different cosmological parameters and a different expansion history $a_{box}(t)$ than the background universe, and thus have different values for the growth factor $D_{box}(a_{box})$. In order for the SU simulation to have the same initial conditions as the background, it must be given a $P_{lin,box}(k, a_{box} = 1)$ that, when scaled back to early redshift using $D_{box}(a_{box})$, matches the results of equation 2.29. To achieve this, the input linear power spectrum for the SU simulations must be

$$P_{lin,box}(k, a_{box} = 1) = \left(\frac{D_{box}(a_{box} = 1)}{D_{box}(a_{box} = 0.001)}\right)^2 \left(\frac{D(a = 0.001)}{D(a = 1)}\right)^2 P_{lin}(k, a = 1). \tag{2.30}$$

Attention must also be paid to the units on $P_{lin,box}$. Typically, N-body code works in units of $h^{-1}Mpc$, which means that any conversion from $h_{uni^{-1}}Mpc$ to $h_{box}^{-1}Mpc$ carries a factor of h_{box}/h_{uni}. The input $P_{lin,box}(k)$ and the range of k values at which it is defined must be converted into $h_{box}^{-1}Mpc$ units

$$P_{lin,box}(k)\left[\left(\frac{Mpc}{h_{box}}\right)^3\right] = P_{lin,box}(k)\left[\left(\frac{Mpc}{h_{uni}}\right)^3\right] * \left(\frac{h_{box}}{h_{uni}}\right)^3 \tag{2.31}$$

$$k\left[\frac{h_{box}}{Mpc}\right] = k\left[\frac{h_{uni}}{Mpc}\right] * \left(\frac{h_{uni}}{h_{box}}\right). \tag{2.32}$$

Finally, the choice of a_{box} given in equation 2.11 ensures that the SU simulations are output at the same cosmological time. However, when it comes time to compute the power spectra and covariance matrix of the simulations, care must be taken to ensure that modes of the same physical size are being compared to one another. One way to ensure this is to choose the comoving size in $h_{box}^{-1}Mpc$ of each simulation such that their proper sizes in Mpc are equal at the desired output epoch

$$\frac{L_{uni}}{h_{uni}}a_{uni} = \frac{L_{box}}{h_{box}}a_{box}. \tag{2.33}$$

Under this convention (called the "total derivative method" in [44]), the proper size of each simulation is equal at the cosmological time corresponding to scale factor a_{box}. Then,

the coordinates of all of the particles in the box can be converted from comoving Mpc to physical Mpc, and the power spectrum can be directly computed from there. This method has the advantage of being intuitive and simple to implement, but the physical scale these simulations will only be equal at one scale factor a_{box}; if you want to use this method to evaluate multiple different epochs, you will need to run a new set of simulations for each epoch.

Another method is to set the comoving box size in Mpc of each simulation to be the same

$$\frac{L_{uni}}{h_{uni}} = \frac{L_{box}}{h_{box}}. \tag{2.34}$$

Under this convention (the "growth-dilation method" in [44]), the comoving sizes of the simulations are the same at the output epoch, but their proper sizes are all different. Just as in the total derivative method, the particle positions can be converted to proper units and the power spectrum can be computed from there.

2.8 Error on the covariance

The error on the covariance matrix ΔC would be an 8-point function, which is far beyond what is practical to compute analytically. Instead, an estimate of the error on the covariance matrix can be calculated by resampling the ensemble power spectra. In the bootstrap resampling technique, the resamples are generated by randomly sampling the power spectra with replacement. A covariance matrix estimate can be computed from each of these resamples, forming a distribution of covariances. The variance of this distribution then represents the estimate of the error on the covariance.

Another way to estimate error on the covariance is using the Wishart distribution [79]. The Wishart distribution represents the distribution of sample covariance matrices from a multivariate normal distribution. Thus, the variance of the Wishart distribution of a covariance matrix is equal to the variance of the covariance matrix itself. The variance of the Wishart distribution is given by

$$Var(W_p) = \frac{1}{n}(C_{i,j}^2 + C_{i,i}C_{j,j}), \tag{2.35}$$

where n is the degrees of freedom ($n = 1$ for the covariance of the power spectrum). The error of the covariance matrix can then be calculated as

$$\Delta C = \sqrt{\frac{Var(W_p)}{N_s}}, \tag{2.36}$$

where N_s is the number of power spectrum samples. To test whether the bootstrap and Wishart distribution give the same error on the covariance, I applied both methods to the ensemble of small boxes described in sec. 2.9.1. The error on the covariance calculated from the Wishart distribution and bootstrap resampling methods match to within 5 percent or better for all k bins of the power spectrum (fig. 2.2). Since the Wishart distribution is much less computationally expensive than bootstrap resampling, I will be using it to compute all errors on covariance in this work.

2.9 Comparison of methods

The different methods of computing the effects of SSC presented thus far have been verified in previous works to recover the SSC correction with reasonable accuracy [44, 10, 77, 38]. However, these methods have yet to be directly compared to one another with the goal of determining which yields the most accurate results. In section 2.9.1, I compare the effectiveness of the Sirko and SC parameter choices in matching the covariance estimated from an ensemble of subsampled boxes. In section 2.9.2, I similarly compare the addition and ensemble methods of computing the SSC contribution. Section 2.9.3 demonstrates that the individual covariance matrices I calculated for these comparisons are consistent with the results of previous works.

2.9.1 Comparison of parameter choices

To demonstrate the significance of SSC and compare the effectiveness of the Sirko and SC parameter choices in recovering the SSC effect, I ran ensembles with the following parameters using L-PICOLA:

- *Small boxes*: 9728 $L = 625\ h^{-1}Mpc$, $N = 256^3$ simulations with identical cosmological parameters, used to compute the covariance matrix without any SSC correction.
- *Sub boxes*: 8 $L = 5000\ h^{-1}Mpc$, $N = 2048^3$ simulations with identical cosmological parameters. Each simulation was subdivided into 512 sub boxes with $L = 625\ h^{-1}Mpc$, resulting in a total of 4096 sub boxes. These simulations innately include the effect of SSC and will serve as the benchmark that the other ensembles will try to match the covariance of.
- *Sirko addition method*: 128 $L = 625\ h^{-1}Mpc$, $N = 256^3$ simulations, where the first set of 64 simulations has been generated with cosmological parameters corresponding

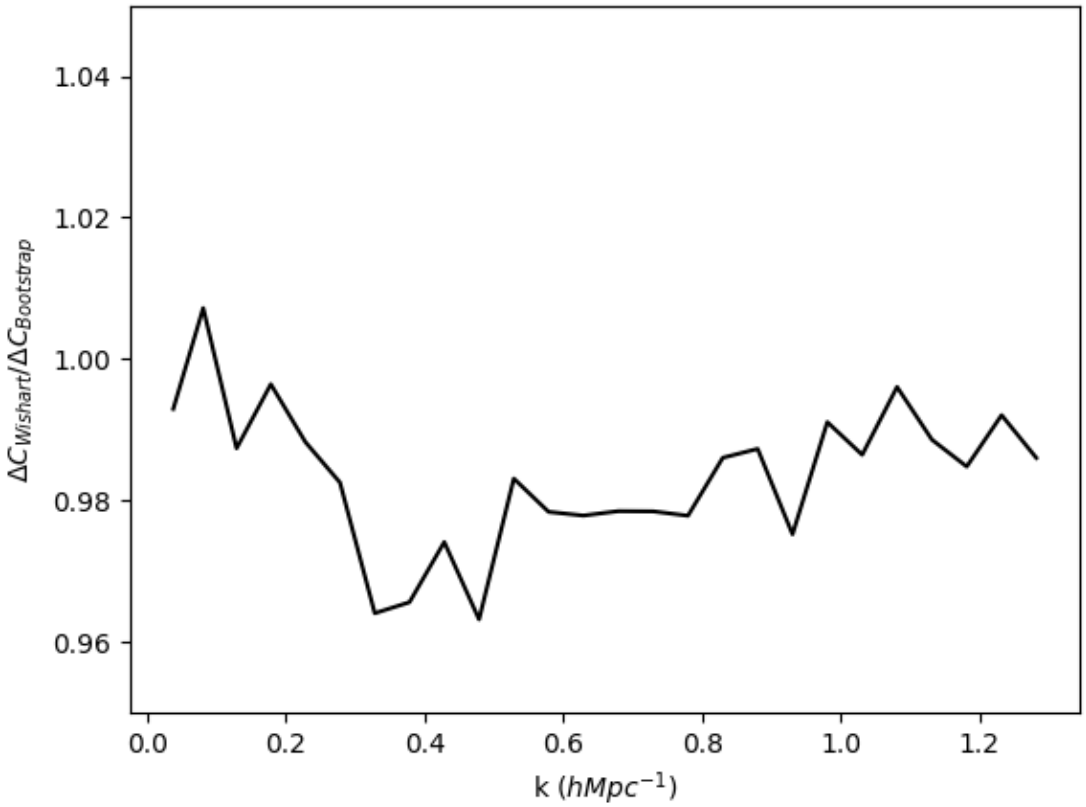

Figure 2.2: Ratio of diagonal elements of ΔC as estimated by the Wishart distribution and bootstrap resampling methods.

Cosmological parameter	Value
h	0.6736
$\Omega_b h^2$	0.02237
$\Omega_{cdm} h^2$	0.12
n_s	0.9649
σ_8	0.8111

Table 2.1: Flat ΛCDM parameters chosen for the background universe.

to $\delta_b = -0.01$, and the remaining 64 generated with $\delta_b = 0.01$. The two sets of simulations are run with the same set of 64 initial seeds. These simulations are used to compute the power spectrum derivative needed for the SSC term in equation 2.28. The no-SSC term is computed using the "Small boxes" ensemble covariance.

- *SC addition method*: $128L = 625\ h^{-1}Mpc$, $N = 256^3$ simulations identical in setup to the "Sirko addition method" simulations except using the SC method parameters from equation 2.26d.

Each ensemble has fiducial background cosmological parameters given by table 2.1 and each mock run to an output redshift corresponding to $a_{uni} = 1$. The power spectrum estimator for each simulation was calculated using the processing pipeline described in section 1.6.4. For the addition method simulations with nonzero δ_b, the output particle positions were first converted from comoving $h_{box}^{-1}Mpc$ units to proper $h_{uni}^{-1}Mpc$ units before computing $P_{global}(k)$ using equation 2.28. In the case of the sub boxes, P_{local} was calculated first. The background overdensity in each sub box was computed as

$$\delta_b^{sub} = \frac{n_{sub}}{\overline{n}_{sub}}, \tag{2.37}$$

where n_{sub} is the number density of particles in the box and $\overline{n}_{sub} = (\frac{256}{625} hMpc^{-1})^3$ is the average number density of particles in all of the sub boxes. These background overdensities were then used to calculate P_{global} using equation 2.2.

To ensure that the use of a spherical window in equation 2.6 does not bias the calculation of σ_b^2, I compared the variance on the background densities of the sub boxes $\left\langle |\delta_b^{sub}|^2 \right\rangle$ to the value of σ_b^2 computed for a sphere matching the volume of the sub boxes. I found the sub box background variance to be $\left\langle |\delta_b^{sub}|^2 \right\rangle = 6.53 \times 10^{-5}$ compared to the calculated $\sigma_b^2 = 6.49 \times 10^{-5}$. These values agree within 0.6% of one another, indicating that using a spherical window for determining σ_b^2 has not significantly biased the SSC calculation.

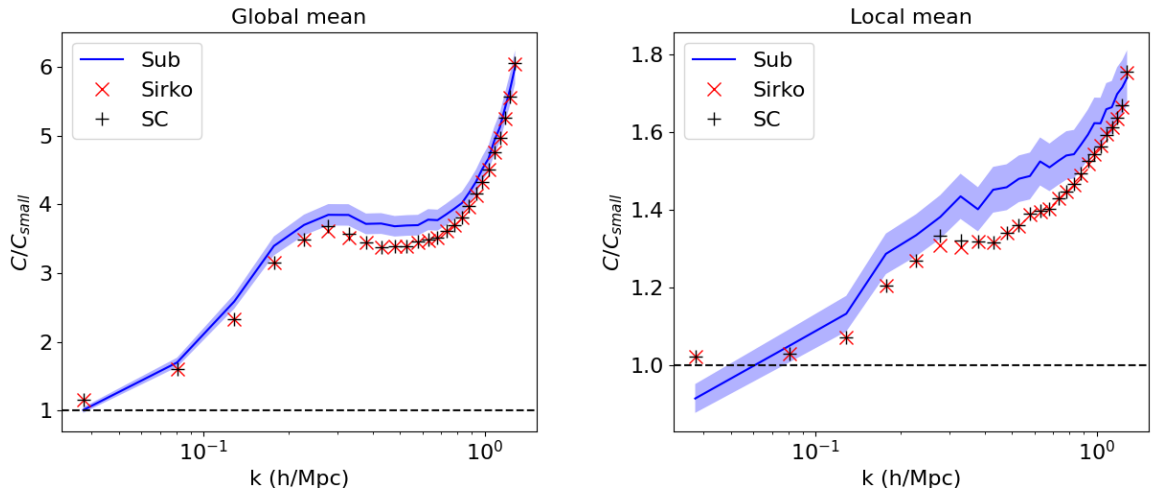

Figure 2.3: Comparison of Sirko (red x) and SC (black plus) method variance over small box variance. The blue curve and shaded region show the subbox variance and the 2σ confidence interval. The left and right panels show the variance relative to the global and local mean density, respectively.

Figure 2.3 shows the results of the comparison between the Sirko and SC method. The SSC effect on the variance relative to the global mean is to increase it by over 100% for modes of $k \sim 0.1\ hMpc^{-1}$ or greater compared to simulations without SSC. The effect grows larger as k increases, reaching a maximum of almost 500% at $k = 1.28\ hMpc^{-1}$. This increase in variance is primarily driven by the normalization of the power spectrum with different δ_b. In the addition method, the global and local power spectrum derivatives can be related by taking the derivative of equation 2.2

$$\frac{dP_{local}(k)}{d\delta_b} = \frac{dP_{global}(k)}{d\delta_b} - 2\overline{P}(k), \tag{2.38}$$

where $\overline{P}(k)$ is the average small box power spectrum. This increase in the power spectrum derivative drives the increase in the effect of SSC. In the local mean case, the SSC contribution is still significant, reaching almost 80% at $k = 1.28\ hMpc^{-1}$.

Both the Sirko and the SC methods perform well in modeling the SSC effect and were able to recover the same variance as the subboxes to within 10% or better for most bins in both the local and global cases (figure 2.4). The two methods were consistent with one another, with the SC method only marginally outperforming the Sirko method in some of the bins. The variance from both methods was biased low compared to the subboxes, and

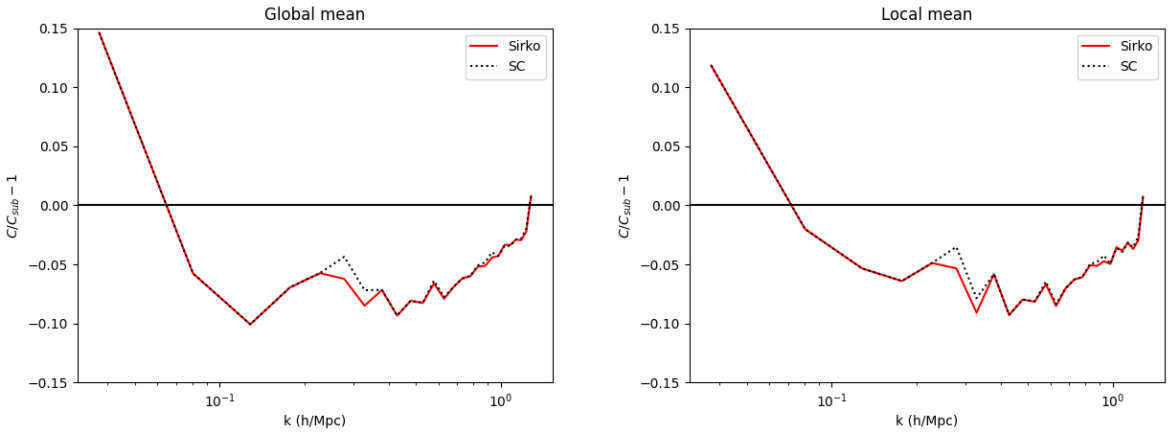

Figure 2.4: Accuracy of the Sirko and SC method variances in recovering the subbox variance. The left and right panels show the global and local mean results, respectively.

most bins fell outside of the 95 percent confidence interval of the subbox variance. This bias could be a product of using the addition method, which itself is only a first order approximation of the ensemble method.

The shape of the covariance ratio plots in figure 2.3 can be attributed to the fact that the SSC contribution to the total covariance arises from three different sources: the Beat Coupling (BC) effect, the Linear Dilation (LD) effect, and the Halo Sample Variance (HSV) effect [35, 21, 71, 44]. The BC effect arises from the fact that short wavelength modes grow faster in the presence of a background overdensity due to coupling with super-sample modes in the trispectrum. This enhances the covariance on quasi-linear scales where linear theory breaks down and the contribution of the trispectrum to the total covariance becomes significant. The LD effect occurs due to the nonzero background overdensity changing the Hubble parameter of the survey or simulation volume, adjusting the expansion rate and modifying the physical sizes of modes within the volume. This has a dampening effect on the covariance, partially cancelling out the BC contribution. The HSV effect is what is captured in the peak-background split discussed in sec. 2.2; the presence of a long wavelength mode pushes more of the initial density perturbations over the threshold for collapse, causing more halos to form. This effect substantially increases the covariance in the deeply nonlinear regime.

Figures 2.5 and 2.6 show the results of the SSC correction on the off-diagonal terms of

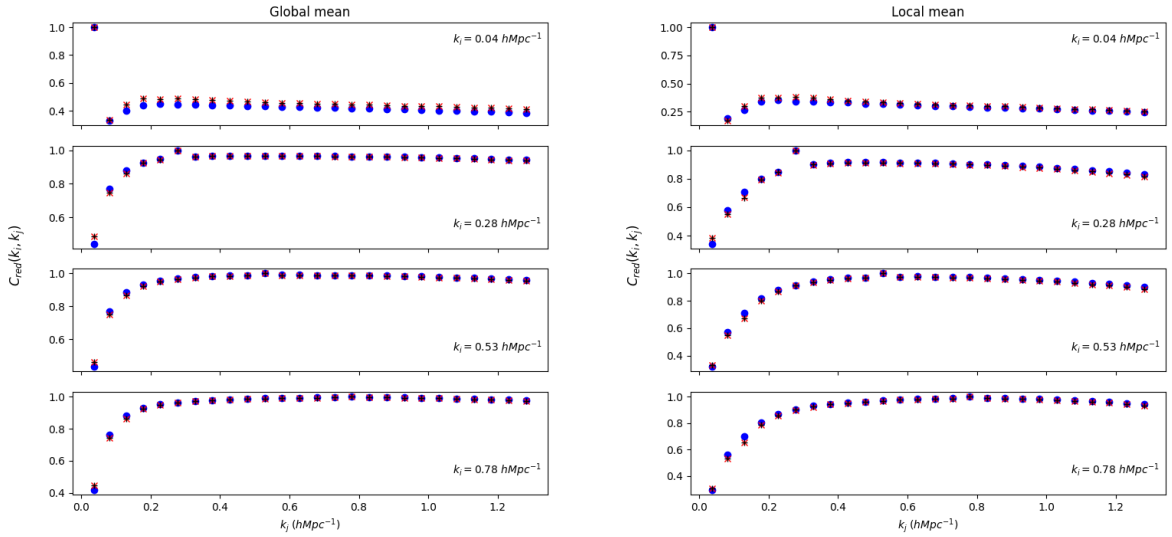

Figure 2.5: Comparison of select off-diagonal elements of the correlation matrices. Blue dots represent sub boxes, black pluses represent SC addition boxes, and red xs represent Sirko addition boxes. The left and right panel show the global and local mean results, respectively.

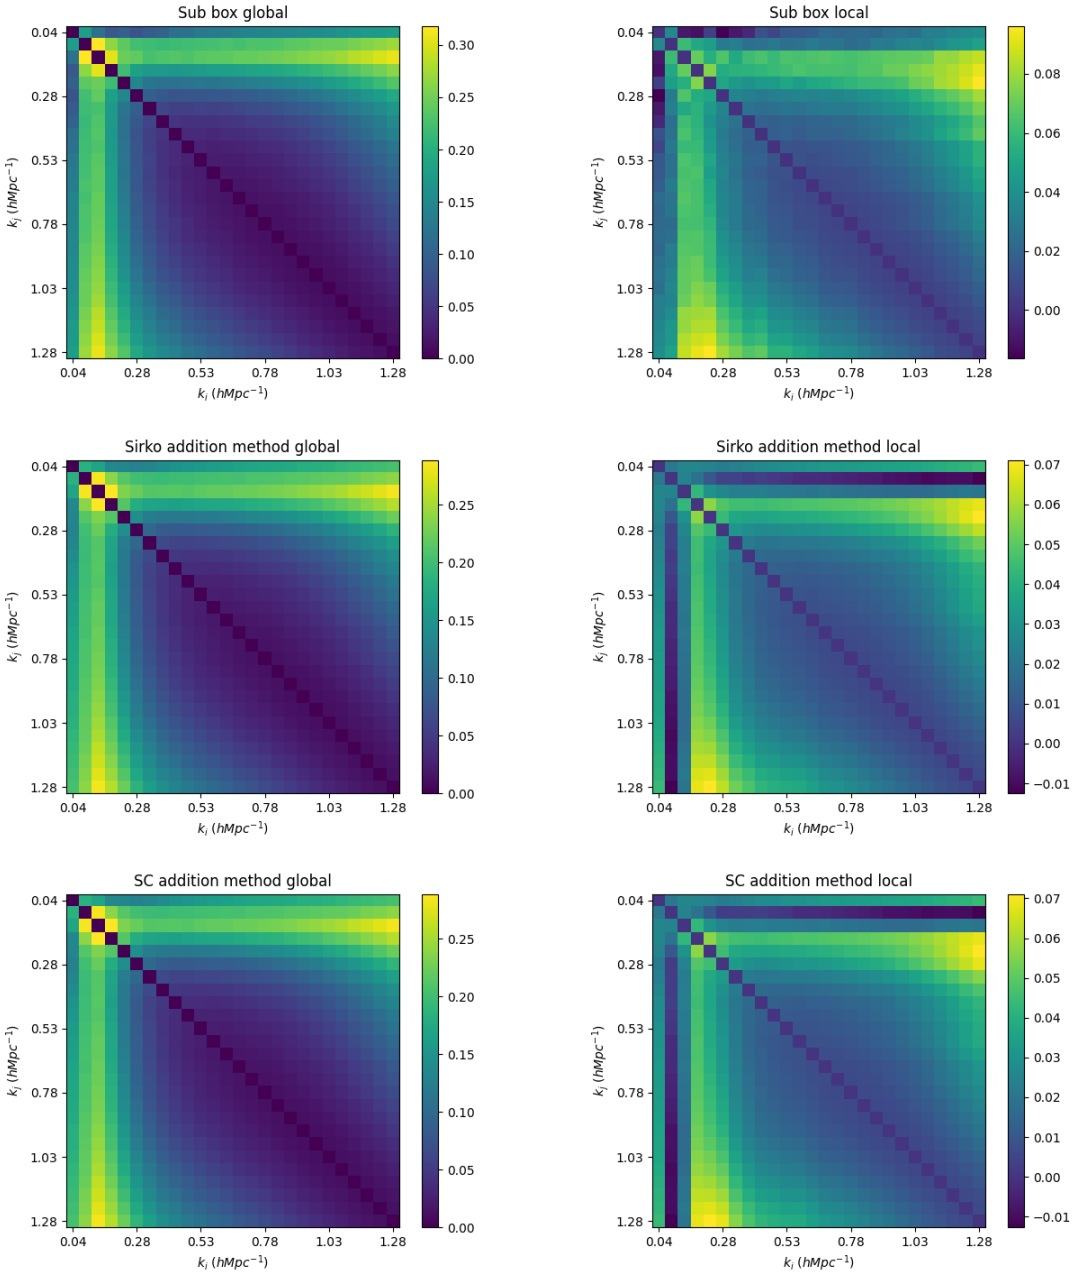

Figure 2.6: Difference in correlation coefficients $C_{red} - C_{red,small}$ for sub boxes (top row), Sirko addition method (middle row), and SC addition method (bottom row). Left and right columns show global and local mean results, respectively.

the covariance matrix. Here, the elements are plotted in terms of the correlation matrix

$$C_{red}(k_i, k_j) = \frac{C(k_i, k_j)}{\sqrt{C(k_i, k_i)C(k_j, k_j)}}. \tag{2.39}$$

The nonlinear gravitational evolution of the N-body simulations can be readily seen in the extremely high off-diagonal correlation for $k > 0.1\ hMpc^{-1}$. Both the Sirko and SC methods are able to successfully recover the correct off-diagonal covariance. The two methods are highly consistent with one another, with neither one showing a significant improvement over the other.

2.9.2 Comparison of addition and ensemble method implementation

It is also worth comparing the performance of the addition and ensemble implementations of the cosmological parameters. To test this, I ran a new ensemble of simulations called "Ensemble SC", containing 4096 simulations of $L = 625\ h^{-1}Mpc$ with background overdensities drawn from a Gaussian distribution with zero mean and variance given by equation 2.6. The cosmological parameters for each simulation were computed using the SC method using the same fiducial background cosmology given in table 2.1.

For both the global and local normalizations of the power spectra, figure 2.7 shows that the addition and ensemble methods are consistent with one another. The ensemble method estimates the covariance a few percent higher than the addition method in most k bins, causing it to match more closely to the subbox covariance. This is due to the fact that the addition method is a first order approximation of the ensemble method. While this first order approximation is good enough to recover the majority of the SSC effect, the higher order terms make a few percent improvement to the overall accuracy.

Figures 2.8 and 2.9 show the off-diagonal terms of the correlation matrices using the addition and ensemble methods. Once again, both methods are able to accurately recover the off-diagonal covariance. As with the variance, the ensemble method slightly outperforms the addition method in matching the sub box covariances, with the difference being most noticeable in the low k bins (top panels of figure 2.8).

2.9.3 Comparison to previous results

This level of agreement between the subbox covariance and the different SSC methods is comparable to previous studies [21, 39, 44, 73, 10, 35]. To compare my results to

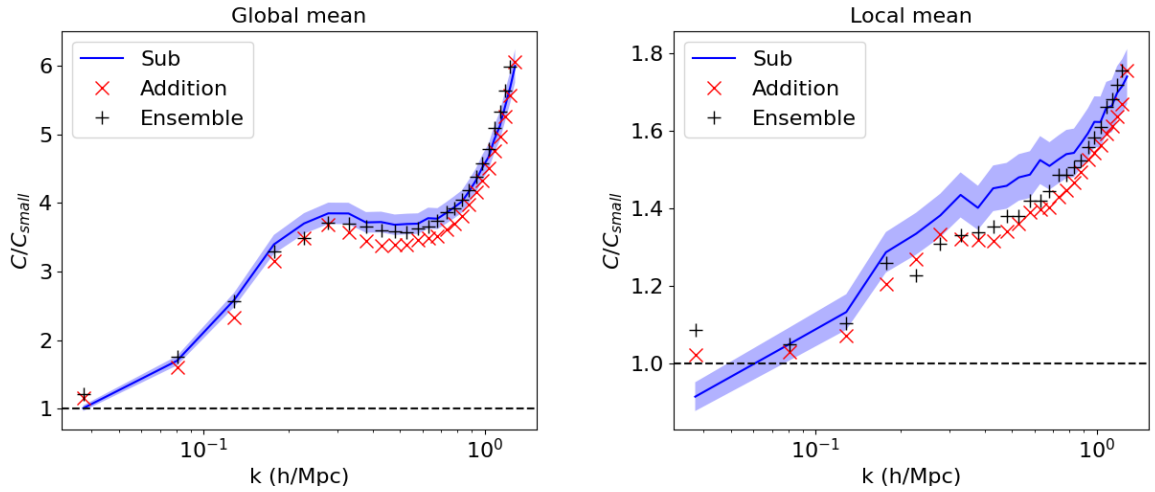

Figure 2.7: Comparison of SC addition (red x) and ensemble (black plus) variance over small mock variance.

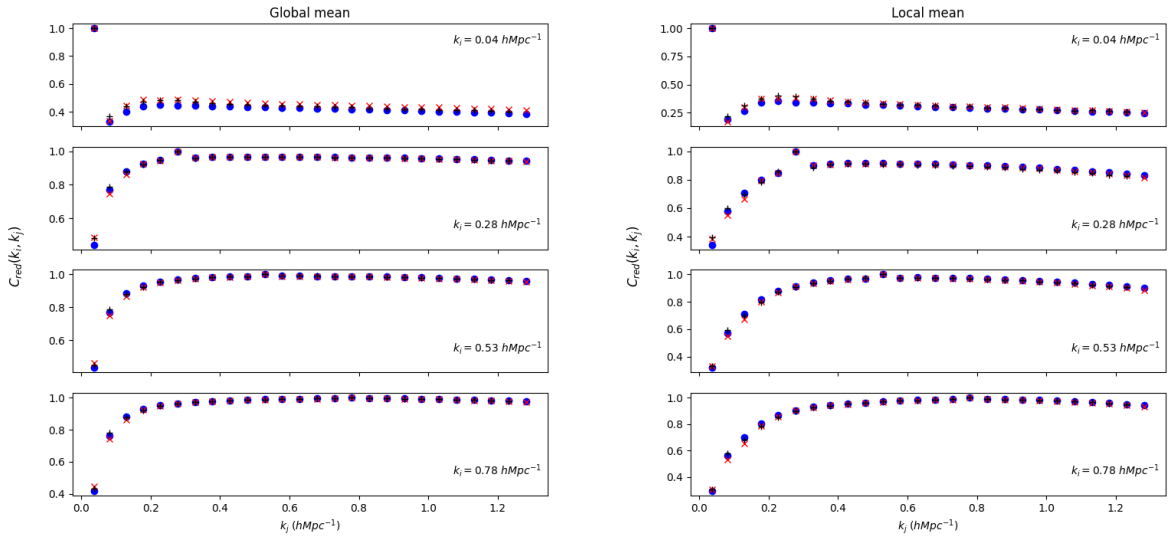

Figure 2.8: Comparison of select off-diagonal elements of the correlation matrices. Blue dots represent sub boxes, black pluses represent SC ensemble method boxes, and red xs represent SC addition method boxes. The left and right panel show the global and local mean results, respectively.

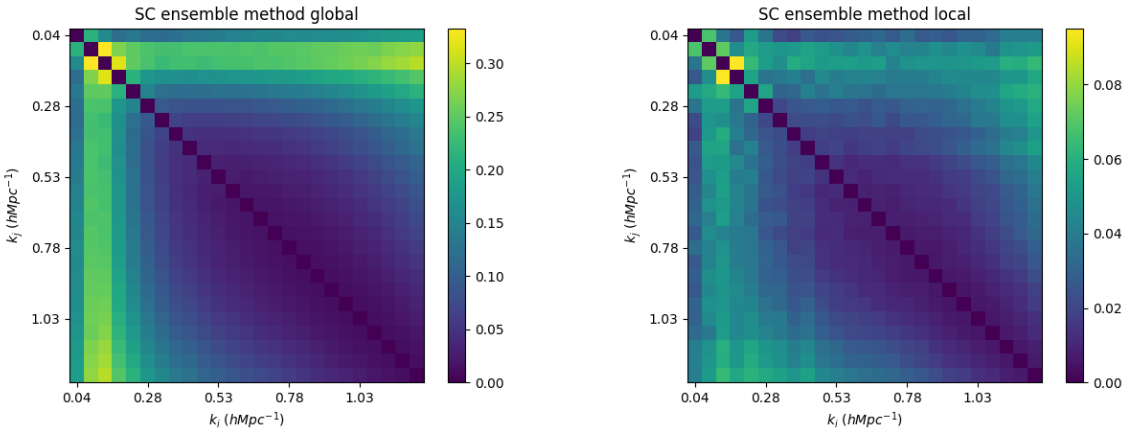

Figure 2.9: Difference in correlation coefficients $C_{red} - C_{red,small}$ for SC ensemble method boxes. Left and right panels show global and local mean results, respectively.

Cosmological parameter	Value
h	0.7
Ω_b	0.047
Ω_m	0.286
n_s	0.96
σ_8	0.82

Table 2.2: Flat ΛCDM parameters used by [44].

these studies, I computed the variance over the Gaussian expectation C_{Gauss} for my power spectra and the data from [10] to match the format of [44]. The results of this are shown in figure 2.10. The SC ensemble method I used was able to capture the effect of SSC with a comparable level of success to the results from [44] and [10], both of which also found agreement between subbox and SU variance to within 10 percent or better.

It should be noted that the covariances measured by each study are different from one another. This is likely driven by differences in the methods used to set up and run the simulations. In [44], the simulations were run with cosmology given in table 2.2 using L-Gadget2 with a (Tree)-PM grid [67]. The (Tree)-PM method of running N-body simulations has the advantage of performing full N-body calculations for motion of particles on small scales while using a PM grid to handle large scale evolution, effectively enhancing accuracy at small scales at the cost of higher computing time. The subboxes were drawn from 7 large-volume mocks with $L = 4\ h^{-1}Gpc$, 2048^3 particles and a 3072^3 grid. To compute the

Cosmological parameter	Value
h	0.6774
Ω_m	0.3089
Ω_b	0.0486
n_s	0.9667
σ_8	0.8159

Table 2.3: Flat ΛCDM parameters used by [10].

subbox power spectra, they assigned the particles using a Clouds-in-Cells (CIC) window to a $(8 \times 1920)^3$ mesh grid, subdivided each large volume mock into 512 subboxes with $L = 500\ h^{-1}Mpc$, computed the subbox power spectra via FFT, and deconvolved the CIC window. Their small box simulations without SSC were run with $L = 500\ h^{-1}Mpc$, 256^3 particles, a 512^3 (Tree)-PM grid, and their power spectra were computed on a 1920^3 mesh grid. They noticed a bias between the average power spectra of the subbox and small box simulations at low k due to the convolution of the subbox power with the window function. To correct for this, they rescale each of the subbox power spectra

$$P_{sub} \rightarrow \frac{\overline{P}_{small}}{\overline{P}_{sub}} P_{sub}, \tag{2.40}$$

where $\overline{P}_{sub}$ and $\overline{P}_{small}$ represent the mean power spectra of the subboxes and small boxes, respectively. For calculating the SSC effect, they perturbed the cosmological parameters in a manner similar to the method of Sirko

$$a_{box} = a_{uni}\left(1 - \frac{\delta_b}{3}\right) \tag{2.41}$$

$$H_{0,box} = H_{0,uni}(1 - \phi) \tag{2.42}$$

$$\Omega_{m,box} = \Omega_{m,uni}(1 + 2\phi) \tag{2.43}$$

$$\Omega_{\Lambda,box} = \Omega_{\Lambda,uni}(1 + 2\phi) \tag{2.44}$$

$$\Omega_{k,box} = 1 - \Omega_{m,box} - \Omega_{\Lambda,box} \tag{2.45}$$

$$\tag{2.46}$$

and used the addition method implementation of these parameters.

The authors of [44] also account for possible correlations between sub boxes drawn from the same large box. They do this by creating $s = 1, ..., N_s/N_L$ independent estimates of the sub box covariance C_s, where N_s is the total number of sub boxes drawn from the N_L

large boxes

$$C_s(k_i, k_j) = \frac{N_L}{N_L - 1}\left[\frac{\sum_{a=1}^{N_L} P_{sa}(k_i)P_{sa}(k_j)}{N_L} - \frac{\sum_{a=1}^{N_L} P_{sa}(k_i)\sum_{b=1}^{N_L} P_{sb}(k_j)}{N_L^2}\right]. \tag{2.47}$$

Each C_s calculated using equation 2.47 is an estimate of the covariance between one sub box and all of the sub boxes drawn from the other large boxes. These independent, uncorrelated C_s values can then be averaged to get an unbiased sub box covariance estimate at the expense of slightly increased ΔC.

In [10], the simulations were run using the cosmology listed in table 2.3 using the FastPM Particle Mesh code [29]. They created their ensemble of 512 sub boxes by subdividing a single $L = 5000\ h^{-1}Mpc$ box containing 2048^3 particles. Their small box ensemble consisted of 512 simulations with $L = 625\ h^{-1}Mpc$ each containing 256^3 particles. To compute the SSC correction to be applied to their small box ensemble, they use Sirko's parameter perturbation (equation 2.15) and the addition method implementation. Their power spectra were computed similarly to the method used by [44], but using a 256^3 grid to assign particles.

The discrepancy between the covariance measured at high k by these different studies is likely primarily driven by the different choice of simulation code. The PM codes used by [10] and this work are expected to have reduced accuracy in modeling small-scale interactions between particles compared to full N-body code or the (tree)-PM code used by [44]. On scales of $k \sim 1\ hMpc^{-1}$, PM simulations using the COLA method has been shown to underestimate the power spectrum and bispectrum by $\sim 10\%$ and $\sim 15\%$, respectively, when compared to results from full N-body simulations [37]. It is likely that the trispectrum is also underestimated relative to full N-body simulations, though this remains to be studied. Due to the strong dependence of the covariance matrix at high k on these higher-order statistics, the covariance would then be underestimated on these scales when using PM simulation codes.

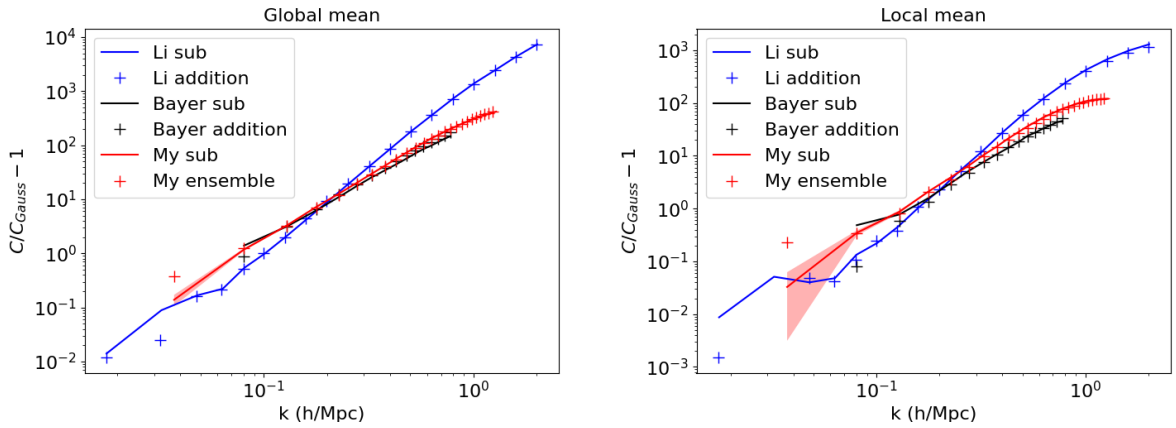

Figure 2.10: Variance over the Gaussian expectation of subbox (solid lines) and SU (pluses) simulations from [44] (blue), [10] (black), and this work (red). The left and right panels show global and local mean results respectively.

Chapter 3

Volume Scaling of the Covariance Matrix

In this chapter, I detail the method by which the covariance of an arbitrary survey volume can be estimated using an ensemble of simulations smaller than the survey volume. Section 3.1 provides the theoretical basis for this technique in the volume scaling of the covariance matrix as presented in [38]. In section 3.2, I provide a description of the ensembles I ran to test how well the volume scaling of the covariance matrix holds. In section 3.3, I demonstrate how well the volume scaling of the covariance matrix holds in the absence of any SSC corrections and discuss some observed departures from the expected analytic behaviour. In section 3.4, I show that the SSC term of the covariance does not scale with volume. In section 3.5, I show how to include the effects of SSC such that an accurate estimate of the full covariance matrix can be recovered after volume scaling. In section 3.6 I derive a way to estimate the smallest volume simulation that can be used and still accurately recover the covariance matrix at scales of interest. I summarize the results from this chapter in section 3.7 by providing a step-by-step description of how to optimally use the volume scaling technique when estimating the covariance matrix.

3.1 Background

The power spectrum covariance matrix $C_{i,j}$ is defined as

$$C_{i,j} = \langle P(k_i)P(k_j)\rangle - \langle P(k_i)\rangle\langle P(k_j)\rangle. \tag{3.1}$$

Considering bins of the covariance matrix corresponding to modes much smaller than the size of the simulation box ($k << 1/L$), the covariance in the absence of super-sample modes can be written as [71, 64, 49]

$$C_{i,j} \simeq C^G_{i,j} + C^{T0}_{i,j}, \tag{3.2}$$

where the Gaussian expectation of the covariance $C^G_{i,j}$ is given by equation 1.40 and the non-Gaussian piece $C^{T0}_{i,j}$ that captures the coupling of modes within the box is

$$C^{T0}_{i,j} = \frac{1}{V}\int_{V_{k_i}} \frac{d^3\mathbf{k}}{V_{k_i}} \int_{V_{k_j}} \frac{d^3\mathbf{k}'}{V_{k_j}} T(\mathbf{k}, -\mathbf{k}, \mathbf{k}', -\mathbf{k}'). \tag{3.3}$$

Here, $T(k_1, k_2, k_3, k_4)$ is the matter trispectrum

$$\langle \delta(\mathbf{k}_1)\delta(\mathbf{k}_2)\delta(\mathbf{k}_3)\delta(\mathbf{k}_4)\rangle = \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 + \mathbf{k}_4) T(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4). \tag{3.4}$$

It is crucial to note here that both $C^G_{i,j}$ and $C^{T0}_{i,j}$ scale inversely with the volume of the survey or simulation. In principle, one could estimate the covariance of a large volume V_L by running an ensemble of simulations with smaller volume V_S, computing the covariance matrix, and scaling it by a factor of $\frac{V_S}{V_L}$.

Now consider the error on the estimated covariance matrix ΔC. It has been shown in [73, 72] that the error on covariance estimated from an ensemble of simulations scales with the inverse of the square root of the number of simulations run N:

$$\Delta C \propto \frac{1}{\sqrt{N}}. \tag{3.5}$$

The error on covariance is also proportional to the covariance matrix, which itself scales with the inverse of the volume of the simulations. Equation 3.5 can be modified to include this volume dependence

$$\Delta C \propto \frac{1}{\sqrt{N}}\frac{1}{V}. \tag{3.6}$$

Given this relation and the volume scaling of the covariance matrix, it can be shown that running smaller simulations and applying the volume scaling results in a reduced error on the covariance matrix given the same amount of computing time. To demonstrate this, suppose one wants to estimate the covariance matrix of a survey with volume V_L. If one runs an ensemble of N_L simulations with volume V_L, the error on the resulting covariance matrix estimate will be proportional to

$$\Delta C_L \propto \frac{1}{\sqrt{N_L}}\frac{1}{V_L}. \tag{3.7}$$

In the same amount of time it takes to run these large volume simulations, a second ensemble consisting of a larger number N_S of simulations with volume $V_S < V_L$ can be run. This ensemble would have error on its covariance given by

$$\Delta C_S \propto \frac{1}{\sqrt{N_S}} \frac{1}{V_S}. \tag{3.8}$$

Since V_S is less than the survey volume, the covariance C_S estimated from the small mocks will be larger than the true covariance of the survey. ΔC also scales more strongly with volume than number of mocks run, ΔC_S will be larger than ΔC_L. However, the volume scaling can be applied to C_S to create a scaled covariance $C_{S,scaled}$ that matches the survey covariance. Since the error on the covariance is directly proportional to the covariance itself, this volume scaling also applies to ΔC

$$\Delta C_{S,scaled} \propto \frac{1}{\sqrt{N_S}} \frac{1}{V_S} \frac{V_S}{V_L} \propto \frac{1}{\sqrt{N_S}} \frac{1}{V_L} \propto \sqrt{\frac{N_L}{N_S}} \Delta C_L. \tag{3.9}$$

Thus, the scaled covariance measured from the small simulations has reduced error compared to the covariance from the large mocks for the same computational cost.

In practice, there are some complicating factors preventing one from being able to scale the covariance from arbitrarily small simulations. To estimate covariance on large scales, the simulations must be large enough that enough modes are present to measure the power spectrum on these scales. Not only must the simulations be large enough to accommodate the larges mode of interest, there must be enough of these modes present in each simulation for the power spectrum to obey the central limit theorem. If not enough large-scale modes are present, the power on these scales will be skewed rather than Gaussian distributed, rendering the measured covariance inaccurate [6]. A problem also arises on non-linear scales due to the behaviour of SSC. The SSC term contributes substantially to the total covariance even on weakly nonlinear scales, but it does not scale as $1/V$ (see section 3.4). It is then necessary to include the SSC term in such a way that it matches the survey SSC after volume scaling (see section 3.5).

3.2 Ensembles run

It is worth investigating how well the theoretical volume scaling of the covariance matrix holds for simulations of different volumes. To this end, I have the following ensembles with different box side lengths using L-PICOLA:

- $L = 1250\ h^{-1}Mpc$ *boxes*: 9728 $N = 512^3$ simulations with identical cosmological parameters. These simulations are used to calculate the large volume covariance that the re-scaled smaller volume covariances will be compared to. Values with subscript L refer to this ensemble.

- $L = 625\ h^{-1}Mpc$ *boxes*: 9728 $N = 256^3$ simulations with identical cosmological parameters. These simulations have volume $V_S = \frac{V_L}{8}$, but are otherwise identical to the $L = 1250\ h^{-1}Mpc$ mocks.

- $L = 312.5\ h^{-1}Mpc$ *boxes*: 9728 $N = 128^3$ simulations with identical cosmological parameters. These simulations have volume $V_S = \frac{V_L}{64}$, but are otherwise the same as the $L = 1250\ h^{-1}Mpc$ mocks.

- $L = 1250\ h^{-1}Mpc$ *addition method boxes*: 128 $N = 512^3$ simulations used for computing the power spectrum derivative for the addition method SSC. 64 of the simulations were generated with $\delta_b = 0.01$, and the remaining 64 were generated with $\delta_b = -0.01$.

- $L = 625\ h^{-1}Mpc$ *addition method boxes*: 128 $N = 256^3$ simulations used for computing the power spectrum derivative.

- $L = 312.5\ h^{-1}Mpc$ *addition method boxes*: 128 $N = 128^3$ simulations using for computing the power spectrum derivative.

Each ensemble has fiducial cosmological parameters given in table 2.1. Each mock has the same number density of particles $\overline{n} = (\frac{256}{625}\ h\ Mpc^{-1})^3$. The power spectra for each mock was computed using the processing pipeline described in section 1.6.4 and identical mesh grid resolutions. In principle, the power spectrum derivative should be independent of simulation volume, meaning it can be calculated from just one of the "addition method" ensembles. I found the $L = 1250\ h^{-1}Mpc$ and the $L = 625\ h^{-1}Mpc$ power spectrum derivatives to be equal to within less than 1% in all k bins (see fig. 3.1). For the $L = 312.5\ h^{-1}Mpc$, the first k bin of the power spectrum derivative was underestimated compared to the larger ensembles. This is caused by skewness of the distribution of power in this bin for the $L = 312.5\ h^{-1}Mpc$ simulations (see section 3.3). Due to this, I chose to use the $L = 625\ h^{-1}Mpc$ power spectrum derivative to compute the SSC term of all three ensembles.

I used each ensemble to estimate the covariance within a $(1250\ h^{-1}Mpc)^3$ volume survey. The $L = 1250\ h^{-1}Mpc$ ensemble served as the baseline to compare the other ensembles to. The covariances of the $L = 625\ h^{-1}Mpc$ and $L = 312.5\ h^{-1}Mpc$ ensembles were re-scaled

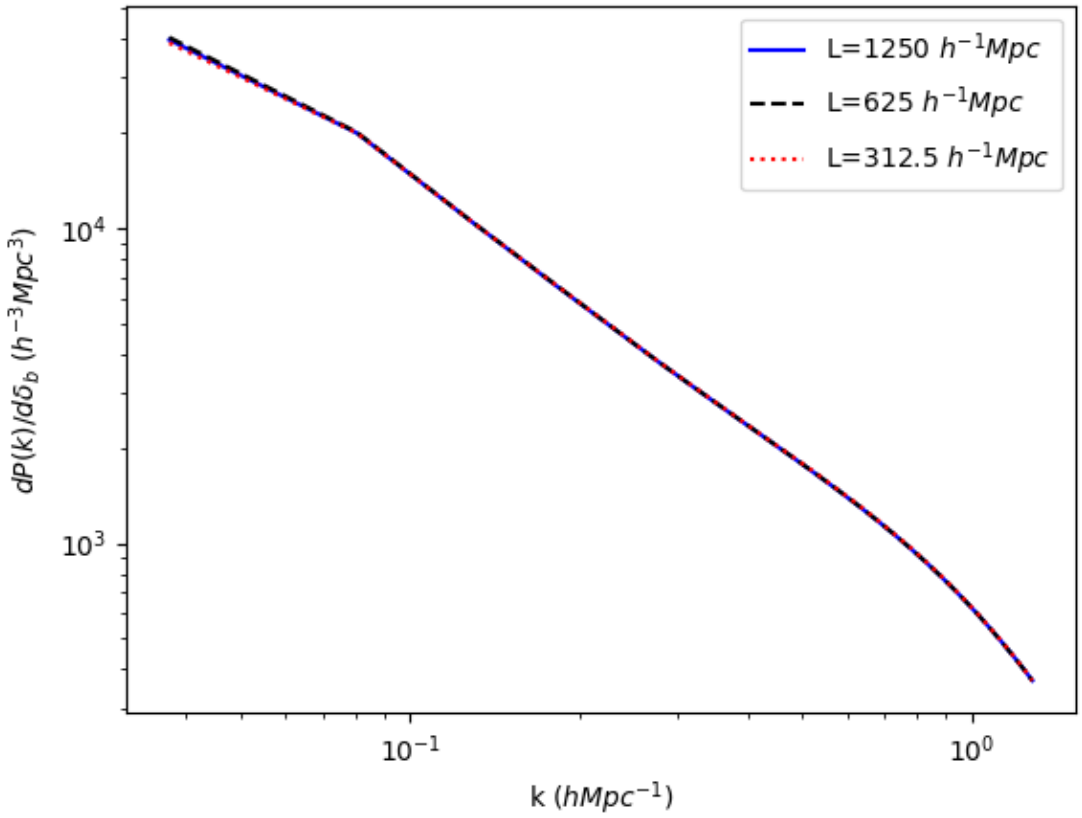

Figure 3.1: Measured power spectrum derivative relative to the global mean for different sized mocks.

by a factor of 1/8 and 1/64, respectively, to match the ratio of their volume to the large mock volume. For both of the smaller ensembles, I computed the ratio of the scaled small box variance $C_{S,scaled}$ and the large box variance C_L. In any k bin where volume scaling holds, the ratio $C_{S,scaled}/C_L$ is expected to equal 1. The error on C_L and $C_{S,scaled}$ can be calculated using the Wishart distribution (see equation 2.35 in section 2.8). The error on $C_{S,scaled}/C_L$ can be found by propagating the uncertainty

$$\Delta\left(\frac{C_{S,scaled}}{C_L}\right) = \frac{C_{S,scaled}}{C_L}\sqrt{\left(\frac{\Delta C_{S,scaled}}{C_{S,scaled}}\right)^2 + \left(\frac{\Delta C_L}{C_L}\right)^2}. \tag{3.10}$$

For the ensembles I ran, $\Delta\left(\frac{C_{S,scaled}}{C_L}\right)$ was on the order of a few percent. To reduce the uncertainty, I further binned the variance into four coarse k bins; these k bins spanned the ranges [0.04,0.23], [0.28,0.48], [0.53,0.73], and [0.78,0.98] $hMpc^{-1}$. Each coarse bin contained 5 fine k bins, resulting in $\sim 1\%$ error on the $C_{S,scaled}/C_L$.

3.3 Volume scaling of covariance without SSC

To test the volume scaling of the covariance without any SSC correction, the scaled covariance of the smaller mocks was calculated as

$$C_{S,scaled} = \frac{V_S}{V_L}C_S, \tag{3.11}$$

where $V_S = L_S^3$ is the volume of the smaller mocks, $V_L = (1250\ h^{-1}Mpc)^3$ is the volume of the large mocks, and C_S is the covariance of the smaller mocks before volume scaling is applied. Figure 3.2 shows the ratio of the scaled small mock variance to the large mock variance. The general trend is that the scaled small mock variance slightly underestimates the large mock variance, with the underestimation becoming more significant in higher k bins and for lower small mock volume. The $L = 625\ h^{-1}Mpc$ mocks recover the correct large mock variance for $k < 0.28\ hMpc^{-1}$ while underestimating the variance by $\sim 2\%$ for $k \geq 0.28\ hMpc^{-1}$. The difference becomes more drastic for the $L = 312.5\ h^{-1}Mpc$ mocks, underestimating the variance by $\sim 6\%$ in the highest k bin. The underestimation of the variance here is to be expected because the small simulations are missing the contribution to the covariance from modes that are small enough to fit in the $L = 1250\ h^{-1}Mpc$ mocks, but too large to be present in the $L \leq 625\ h^{-1}Mpc$ mocks. Figure 3.3 shows the covariance ratio for the off-diagonal elements without the coarse binning described in sec. 3.2. While the scaled $L = 625\ h^{-1}Mpc$ mock covariance is able to recover the correct large mock

covariance on all scales within a few percent, the $L = 312.5\ h^{-1}Mpc$ mock covariance is significantly underestimated at small k and slightly underestimated at high k. The underestimation at high k is consistent with the explanation given for the on-diagonal terms due to the covariance on these scales being dominated by the underestimated SSC.

With the general trend for the scaled small mock variance to better estimate the large mock variance in lower k bins, it is noteworthy then that the low k bin of the $L = 312.5\ h^{-1}Mpc$ mocks underestimates the variance more than the other k bins for those simulations. This is caused by the fact that each simulation at this size contains very few modes in the low k bins. Since so few of these low k modes are present in each simulation, the power in these bins is skewed rather than Gaussian distributed (figure 3.4). This skewness results in a decreased variance compared to what the volume scaling predicts.

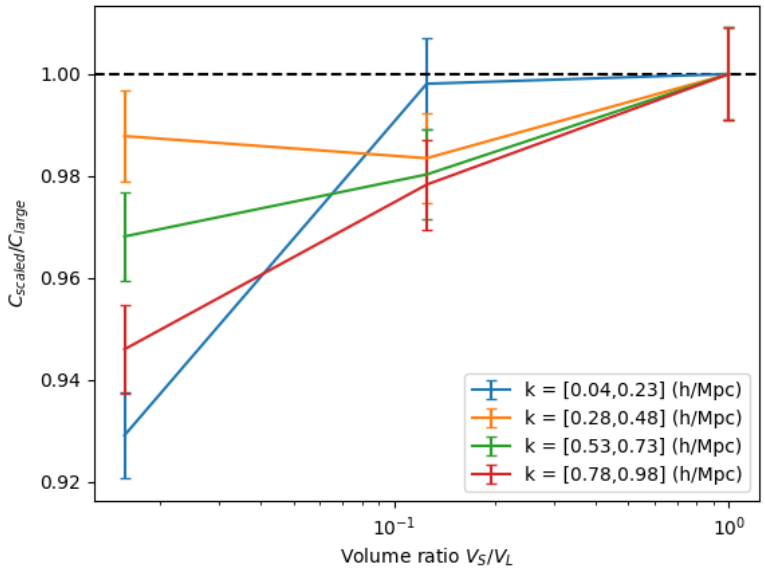

Figure 3.2: Ratio of volume scaled small box covariance and large box covariance with no SSC correction.

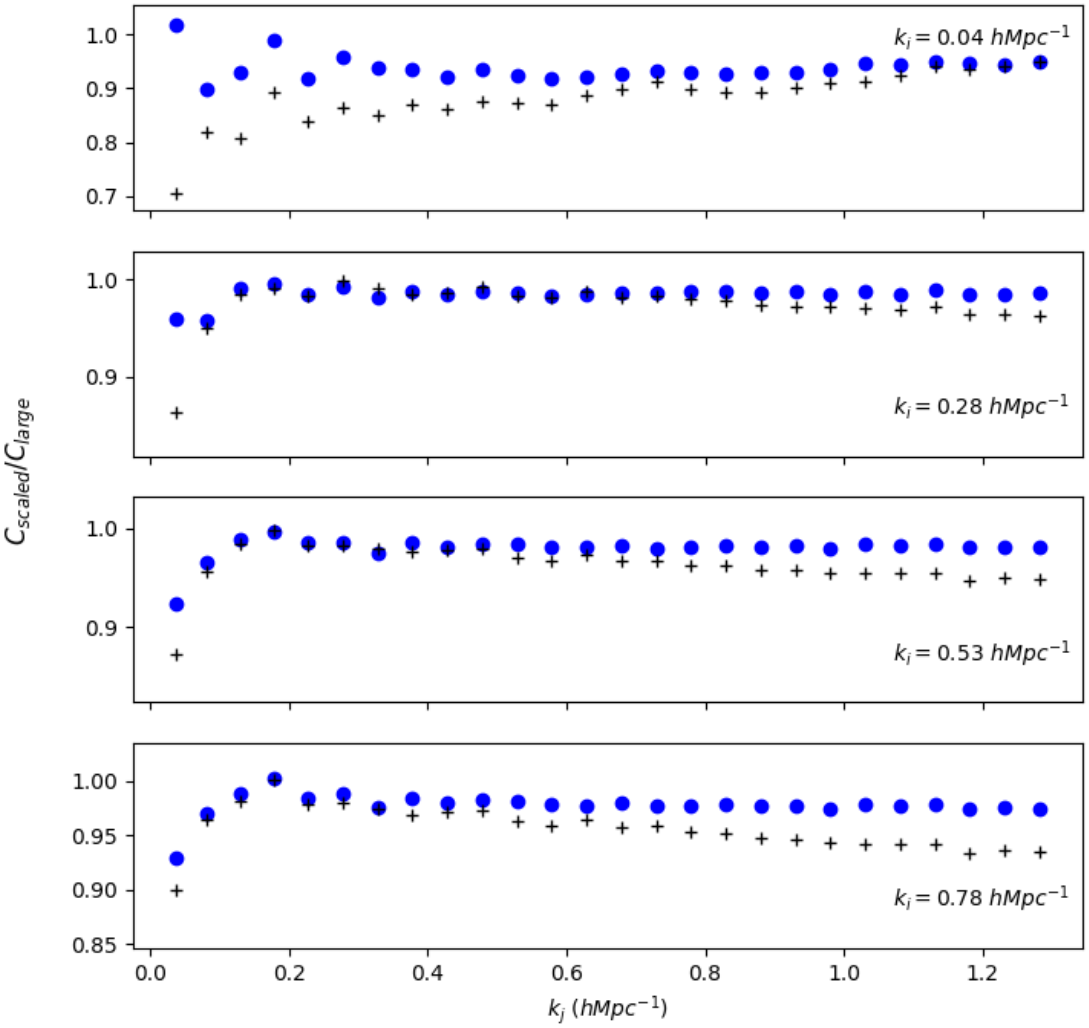

Figure 3.3: Ratio of select off-diagonal elements of the scaled small mock and large mock covariances. The blue dots correspond to the scaled covariance of the $L = 625\ h^{-1}Mpc$ boxes, and the black pluses correspond to the scaled covariance of the $L = 312.5\ h^{-1}Mpc$ boxes.

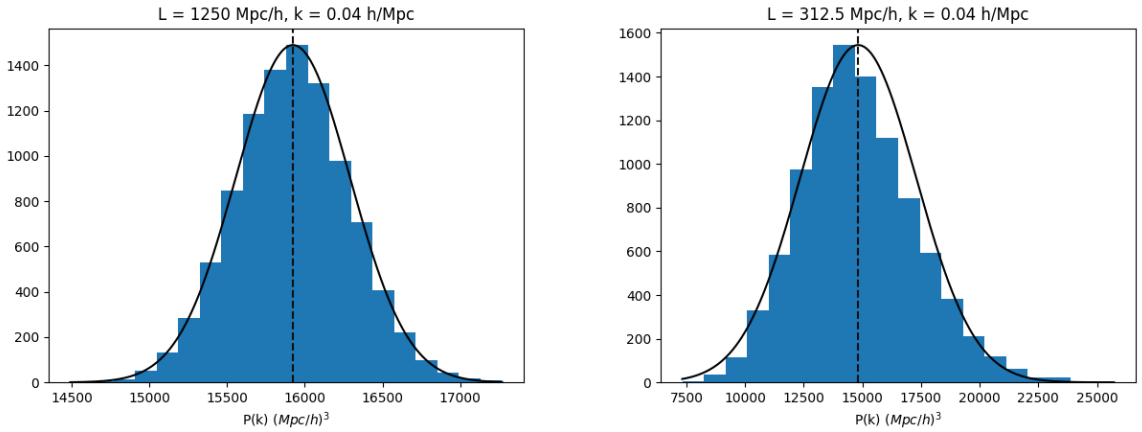

Figure 3.4: Histograms of power spectrum amplitudes in the $k = 0.04\ hMpc^{-1}$ bin. The left panel shows the power spectrum amplitudes in the $L = 1250\ h^{-1}Mpc$ ensemble, and the right panel shows the amplitudes in the $L = 312.5\ h^{-1}Mpc$ ensemble. The black dashed line shows the average power, and the solid black curve shows a Gaussian distribution with mean and variance matching those computed from the power spectra. In the $L = 312.5\ h^{-1}Mpc$ ensemble, there is a noticeable skewness in the distribution compared to a Gaussian.

3.4 Volume scaling SSC

For practical applications of the volume scaled covariance, the SSC contribution must be included. Naively, one might attempt to include the SSC term in the small mock covariance before applying the volume scaling. When using the addition method, the scaled covariance would then be

$$C_{S,scaled}(k_i, k_j) = \frac{V_S}{V_L}\left[C_S(k_i, k_j) + \sigma_b^2\left(\frac{dP(k_i)}{d\delta_b}\right)\left(\frac{dP(K_j}{d\delta_b}\right)\right], \tag{3.12}$$

where the window function used to calculate σ_b^2 corresponds to the volume of the small mocks. However, the SSC term does not scale with volume in the same manner as the rest of the covariance. The power spectrum derivative $\frac{dP(k_i)}{d\delta_b}$ remains constant with respect to volume and σ_b^2 has a nontrivial dependence on volume from the window function in equation 2.7. This can be seen in figures 3.5 and 3.6 where, for both the global and local normalizations of the power, the covariance is greatly overestimated on all scales. Since the SSC term is most significant at high k, this overestimation becomes greater as k increases.

The failure of the volume scaling of SSC can be linked to ergodicity. In the context of cosmology, a field is considered ergodic if the volume average of that field within a survey

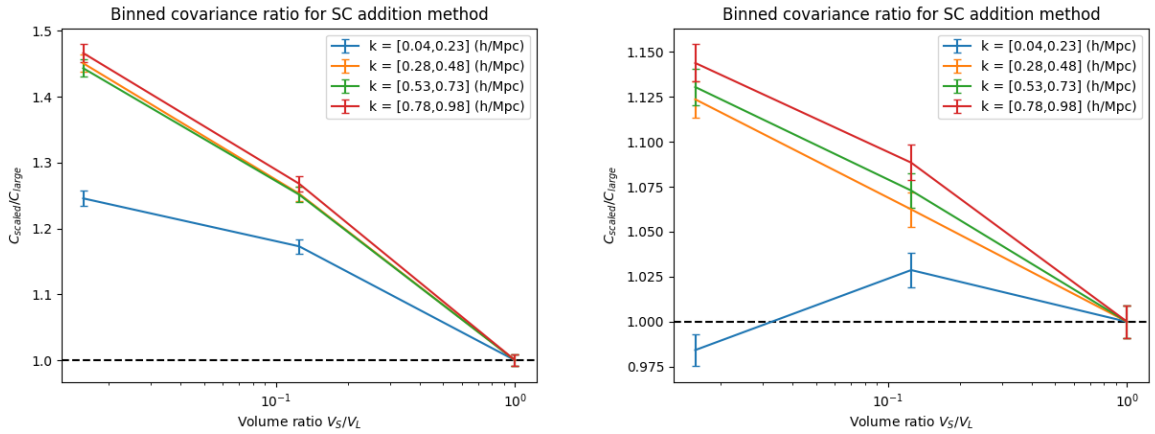

Figure 3.5: Ratio of volume scaled small box covariance and large box covariance with SSC correction. The small mock covariance had its SSC term added before being scaled by the volume ratio. The left and right panels show covariance ratios relative to global and local mean densities, respectively. The error bars represent the 1σ confidence interval as estimated from the Wishart distribution.

or simulation is equal to the ensemble average of that field. Within a sufficiently large box, small patches within it that are sufficiently far apart are causally disconected and thus uncorrelated. These patches can be considered separate universes, and the modes of the power spectrum within each patch will be independent of other patches. The volume average of these modes within the large box would be effectively the same as the average of many small, independent separate universes. This is what gives rise to the volume scaling of the power spectrum covariance for modes smaller than the box width; the number of modes N_{modes} within a simulation is proportional to the volume, and the covariance scales as $1/N_{modes}$ (see equation 2.5). The SSC, conversely, is not ergodic. The ensemble average of δ_b is not equal to the volume average of δ_b; the ensemble average of δ_b is zero, and each simulation only has one value of δ_b from which to draw a "volume average". Increasing the volume of a simulation does not increase the number of background modes present, so the SSC term does not scale as $1/V$.

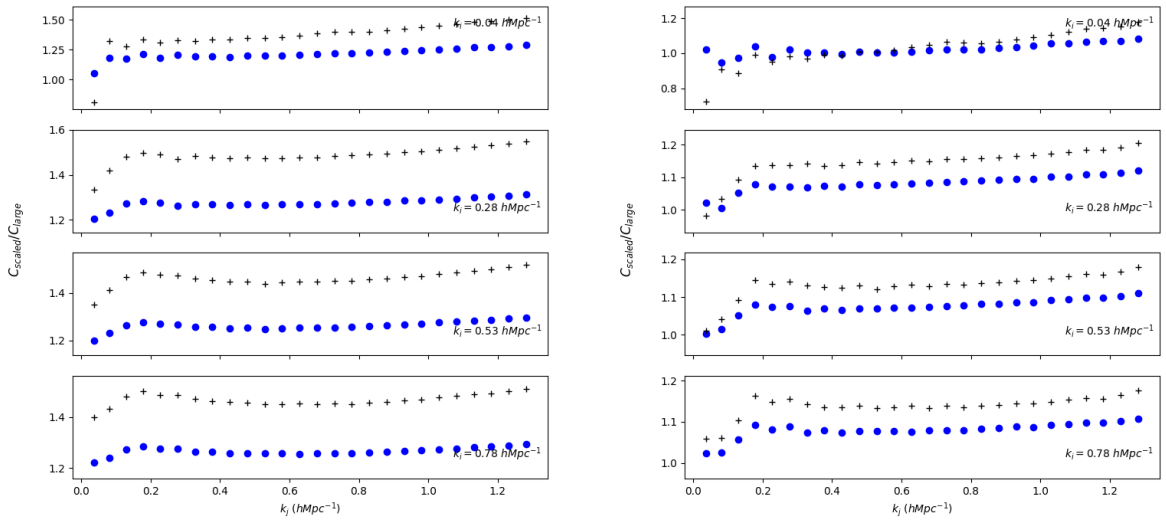

Figure 3.6: Ratio of select off-diagonal elements of the scaled small mock and large mock covariances with the small mock SSC correction as described in equation 3.12. The blue dots correspond to the scaled covariance of the $L = 625\ h^{-1}Mpc$ boxes, and the black pluses correspond to the scaled covariance of the $L = 312.5\ h^{-1}Mpc$ boxes. The left and right panels show the covariance ratios relative to global and local mean densities, respectively.

3.5 Volume scaling before SSC correction

So far we have seen that the non SSC term scales as $\frac{1}{V}$, but the SSC term scales as a highly nontrivial function of V. To take advantage of the volume scaling when computing the large volume covariance, we can calculate the non SSC term from the small mocks and then add the SSC term appropriate for the large volume. For the addition method, the covariance would be

$$C_{S,scaled}(k_i, k_j) = \frac{V_S}{V_L} C_S(k_i, k_j) + \sigma_b^2 \left(\frac{dP(k_i)}{d\delta_b} \right) \left(\frac{dP(k_j)}{d\delta_b} \right), \tag{3.13}$$

where σ_b^2 is being calculated for the large volume (unlike in equation 3.12). When using the ensemble method, this corresponds to drawing the δ_b from a Gaussian distribution with variance $\frac{V_L}{V_S}\sigma_b^2$ [38]. This ensures that the SSC term that gets baked into this ensemble will match that of the large volume after the covariance is re-scaled.

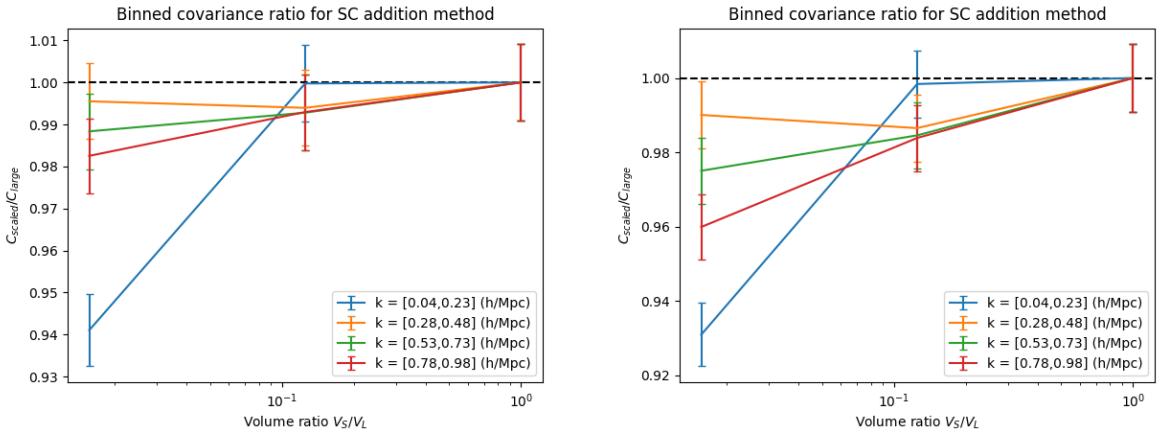

Figure 3.7: Ratio of volume scaled small box covariance and large box covariance with SSC correction. The SSC term was added to the scaled small mock covariance after volume scaling. The left and right panels show covariance ratios relative to global and local mean densities, respectively. The error bars represent the 1σ confidence interval as estimated from the Wishart distribution.

The covariance ratio $\frac{C_{S,scaled}}{C_L}$ using equation 3.13 is shown in figures 3.7 and 3.8. With the exception of the low k bin of the $L = 312.5\ h^{-1}Mpc$ ensemble, the scaled covariance is able to recover the large volume covariance to within 2% or better in the global mean case and within 4% or better for the local mean. There is an improvement in how well the

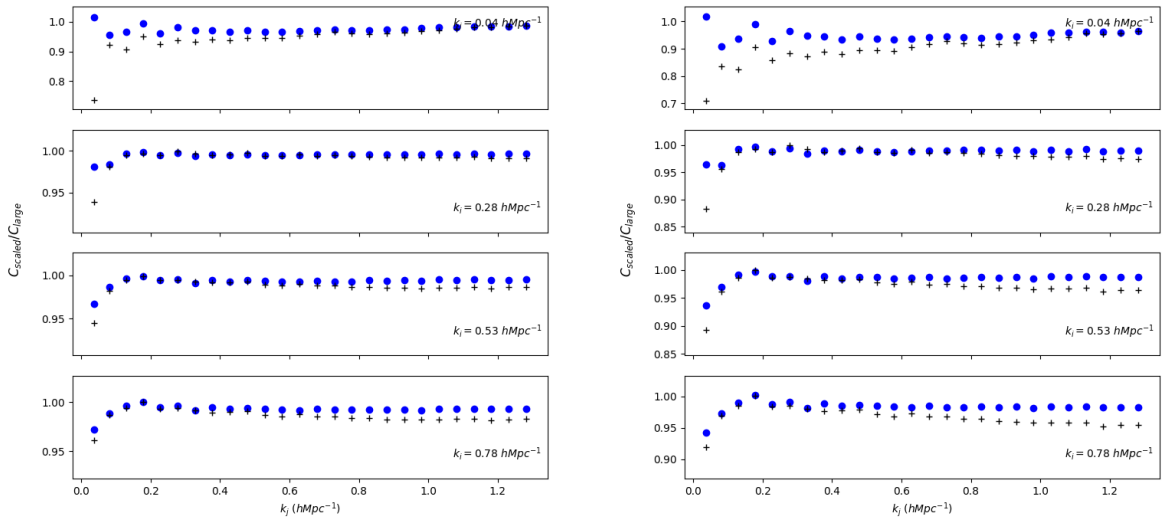

Figure 3.8: Ratio of select off-diagonal elements of the scaled small mock and large mock covariances with the large mock SSC correction as described in equation 3.13. The blue dots correspond to the scaled covariance of the $L = 625\ h^{-1}Mpc$ boxes, and the black pluses correspond to the scaled covariance of the $L = 312.5\ h^{-1}Mpc$ boxes. The left and right panels show the covariance ratios relative to global and local mean densities,respectively.

volume scaling holds here over the no SSC case seen in figure 3.2, especially in the high k bins. This improvement comes from the fact that the SSC terms being added to both $C_{S,scaled}$ and C_L are identical, bringing the ratio $\frac{C_{S,scaled}}{C_L}$ closer to 1, especially in higher k bins where the SSC term is a larger fraction of the total covariance. Despite this, the trend for $C_{S,scaled}$ is to be slightly lower than C_L, especially for higher k and lower V_S. This can again be explained by the missing contribution to the covariance by modes small enough to fit in the large mocks but too large to fit in the small mocks. These modes would contribute to the SSC term of $C_{S,scaled}$.

3.6 A prescription for avoiding skewness

As observed in section 3.3, the minimum effective simulation volume that can be used in the volume scaling technique is limited by there being too few large-scale modes in the small volume simulations for them to obey the central limit theorem. For the volume scaling to be applied in a realistic scenario, the simulation volume must be chosen carefully to ensure scales of interest are well populated with modes.

Consider the power spectrum of a cubic survey or simulation volume with side length L. The volume in k-space of a bin of the spherically averaged power spectrum is given by

$$V_{k_i} = \frac{4}{3}\pi \left[(k_i + \Delta k)^3 - (k_i - \Delta k)^3\right], \tag{3.14}$$

where k_i is the central k value of the bin and Δk is the bin width. The k-space volume of a single mode of the power spectrum is k_F^3, where

$$k_F = 2\pi/L \tag{3.15}$$

is the fundamental mode of the box. The number of modes one would expect to find in a given k bin is thus

$$N_{k_i} = \frac{4\pi \left[(k_i + \Delta k)^3 - (k_i - \Delta k)^3\right]}{3k_F^3}. \tag{3.16}$$

When actually counting the number of modes in a given k bin, the discrete nature of the Fast Fourier Transform adds a shotnoise component to equation 3.16. However, the estimated number of modes recovers the correct order of magnitude when $N_{k_i} \sim 100$, and the shotnoise becomes negligible for $N_{k_i} \gtrsim 1000$. Substituting equation 3.15 into equation 3.16 and expanding the cubed terms yields

$$N_{k_i} = \frac{L^3 k_i^2 \Delta k}{2\pi^2} + \frac{L^3 \Delta k^3}{24\pi^2}. \tag{3.17}$$

The number of modes in a k bin intuitively depends on the volume of the mocks in configuration space, the central value of the k bin, and the width of the k bin.

This equation can be used to estimate the number of modes present in bins of the covariance matrix that show skewness. The $k_i = 0.037\ hMpc^{-1}$ bins of the $L = 312.5\ h^{-1}Mpc$ mock power spectra are expected to contain ~ 140 modes. It is clear then that 140 modes is too few for the power spectrum at this scale to obey the central limit theorem. Examining the off-diagonal elements of the scaled covariance as presented in figure 3.3 shows that, for $k_i \geq 0.13\ hMpc^{-1}$, the bias induced by skewness has vanished for the $L = 312.5\ h^{-1}Mpc$ mocks and the volume scaled covariance recovers the large mock covariance to within a few percent. The $k_i = 0.13\ hMpc^{-1}$ bin of the power spectra in these simulations is expected to have $N_{k_i} = 1300$. This indicates that having $N_{k_i} \gtrsim 1300$ is enough for the central limit theorem to hold, allowing the volume scaling of the covariance to hold well at scales of k_i. The minimum simulation volume at which the volume scaling of the covariance holds can then be estimated by plugging the lowest k_i scale of interest, the desired k-bin width Δk, and $N_{k_i} = 1300$ into equation 3.17 and solving for L.

3.7 How to optimally use the volume scaling technique

I will now present a summary of the most effective way to leverage the volume scaling of the covariance to optimize the computation of survey covariance. First, the minimum k_i scale of interest and desired power spectrum binning scheme Δk must be identified and used to calculate the smallest viable simulation volume. Note that there may be advantages to choosing a larger volume. For example, when using PM simulation code, matching the survey resolution using the minimum simulation volume may require a simulation mesh grid with a number of grid cells not equal to some power of 2. Using such a mesh grid introduces inefficiency in computing the fast Fourier transform of the mesh, so increasing the simulation size may be preferable to avoid this inconvenience. Once a simulation volume is chosen, run an ensemble of simulations using the desired fiducial cosmological parameters and compute their power spectra. If you are interested in the behaviour of high k modes, use a CIC or higher-order resampler when painting the overdensity field to a mesh and use the interlacing technique to reduce the effects of aliasing. Compute the covariance matrix of these power spectra and multiply it by a factor of V_{small}/V_{large}. The appropriate SSC correction for the *survey* volume can be computed as described in section 2.6.2 and added to the scaled covariance to get the final result. This methodology for

computing the covariance matrix extracts the greatest boost in computational efficiency that the volume scaling can provide.

Chapter 4

Conclusions

This chapter is a summary of the work presented in this thesis. It will review the key findings of each chapter and present future directions for research on the subject of SSC and the volume scaling of the covariance matrix.

Chapter 1 introduced the framework for modern cosmology, including the FLRW metric and the ΛCDM model. I presented the observables and statistics relevant for discerning between different cosmological models and quantifying clustering of matter in the universe. The covariance matrix was introduced as a powerful tool for evaluating the likelihood of a cosmological model given a set of observations. Finally, I described the process for running N-body simulations and their role in accurately estimating the covariance matrix at scales where analytic treatment becomes difficult.

Chapter 2 investigated the problem of including Super-Sample Covariance in N-body simulations. The theoretical basis of how a mode of the power spectrum with wavelength longer than the width of a survey or simulation can couple nontrivially to modes within the survey, substantially increasing the covariance especially in the nonlinear regime. The separate universe approach is presented as a way to model the effects of these super-sample modes in simulations that otherwise would explicitly set the amplitude of these modes to zero. In particular, the derivations of the perturbed parameter approach taken by Sirko in [66] and the Spherical Collapse approach presented in [81] are given. I then describe the addition and ensemble methods of using a set of perturbed parameters to model the effect of SSC. By running ensembles of simulations using the Sirko and SC approaches to perturbing the cosmological parameters, I showed a side-by-side comparison of these different methods to determine which performed the best in modeling SSC. I found that both the Sirko and SC approaches were able to model the SSC effect to within 10% or

better, with neither model significantly outperforming the other. I ran further N-body simulations to perform a similar comparison of the addition and ensemble methods. While both methods were able to model the SSC to within 10% or better, the ensemble method outperformed the addition method by a few percent on most scales while requiring less computation time to run the required simulations. After determining which methods most accurately capture the effect of SSC, I compared the covariance matrices measured from my ensembles with the results of other independent studies of these methods in isolation.

Chapter 3 investigated the capability of the volume scaling technique presented in [38] to ease the cost of computing the covariance for large-volume surveys. Taking advantage of the fact that all but the SSC term of the covariance matrix scales as the inverse of volume, it is possible to use a larger number of small volume simulations to compute the covariance of a large volume survey with the same degree of accuracy using less computation time. I generated ensembles of simulations at different volumes to evaluate how well the theoretical volume scaling of the covariance holds. As long as the SSC is correctly accounted for, I found that the large-volume covariance can be estimated from the small-volume simulations to within 4% accuracy or better. Still, the volume-scaled covariance was consistently biased low compared to the large mock covariance, with the bias generally becoming greater as simulation volume decreased and the wavenumber k of the covariance increased. The only exception to this was for the large-scale covariance estimated from the smallest volume simulations, where the extremely limited number of modes of that scale in the volume caused the power in that bin to become skewed, substantially decreasing the estimated covariance in those bins. This skewness demonstrates where and how the volume scaling of the covariance matrix breaks down, effectively setting a minimum simulation volume required to obtain accurate results in the lowest k bins. By estimating the number of modes required to eliminate the skewness problem, I derived a prescription for determining the minimum simulation volume that can be used with the volume scaling technique and still recover an accurate estimate of the covariance on all scales of interest. I concluded the chapter with a prescription for how to optimally use the volume scaling technique and SSC correction to estimate the covariance in a survey volume as efficiently as possible.

The methods presented in this work show great promise for reducing the cost of computing the covariance matrix without incurring a heavy trade-off in accuracy. N-body simulations are notoriously computationally expensive to run owing to the fact that, at best, their runtime scales as $\mathcal{O}(NlogN)$. Next generation survey data from instruments such as the Dark Energy Spectroscopic Instrument (DESI) [2] and the Euclid space telescope [42] will require $\mathcal{O}(1000)$ of simulations run with large volume and high mass resolution to get sufficiently accurate estimates of their covariance matrices. Simulation suites such as AbacusSummit [45] designed to meet DESI requirements require thousands of node-hours

per simulation to run. Running an ensemble of separate universe simulations would eliminate the need for the simulations to be substantially larger than the survey volume to accurately capture the SSC effect, substantially reducing the computational cost. Use of the volume scaling technique presents a further reduction in cost; conservatively assuming one were to run simulations that are $1/8^{th}$ the volume of the survey, the error on the estimated covariance will be reduced by a factor of $\sim \sqrt{8}$ for the same number of node-hours used. Alternatively, this could represent a factor of $\sim \sqrt{8}$ reduction in node-hours needed to get a similar precision on the covariance matrix using large-volume mocks. The volume scaling can potentially be pushed even further at the cost of losing information about the large-scale covariance. If this large-scale information is still within the linear regime, it can potentially be recovered accurately by theoretical methods of computing covariance [52, 50, 16, 11, 9, 8, 51]. A key advantage to using the volume scaling technique is that it can be applied alongside any other techniques used to reduce computational cost. Any future improvements to N-body code efficiency, theoretical modeling of the covariance matrix, or treatment of super-sample modes can be applied to volume-scaled simulations to reduce computational time even further.

Redshift space distortions pose a challenge to the use of volume scaling technique in estimating covariance in realistically-shaped surveys. While a prescription for dealing with a non-trivial survey window function is provided in [38], they find that this prescription breaks down once a line-of-sight dependence is introduced to the power spectrum. A more thorough treatment of the survey window in the presence of redshift space distortions is needed for practical applications.

The volume-scaled covariance from small mocks tends to be biased low relative to the covariance from full-sized mocks. This is likely due to there being a missing contribution to the SSC of the small mocks by modes small enough to fit in the full-sized mocks but too large to fit in the small mocks. Inclusion of these modes in the calculation of the SSC could potentially remove this bias. With the wavelength of these modes being similar to the width of the simulation, they may not be well approximated solely by a correction to how σ_b^2 is calculated. Proper treatment of these modes may require accounting for super-survey tidal fields [63, 47, 5].

While the volume scaling has proven successful in recovering the matter power spectrum covariance, it is worth verifying that the covariance of other statistics of interest are recovered with similar accuracy. The covariance of the 2PCF is expected to scale with volume similarly to the power spectrum covariance and receive a similar improvement in accuracy when using the volume scaling technique. While the SSC correction has been found to accurately model the behaviour of the halo mass function, void size function, and matter bispectrum covariance in [10], the volume scaling of these statistics, as well as the

higher-order mutipoles of the 2-point statistics, should be investigated.

Continued improvements in the volume and resolution of cosmological surveys necessitates a race to keep the computational time of covariance calculations feasible. This thesis constitutes a step forward in that race by evaluating the accuracy of different methods of including SSC in simulations and by demonstrating at high precision the effectiveness of the volume scaling technique. If the potential of this technique can be fully realized, the ambitions of next generation surveys will become much closer to our grasp.

References

[1] TMC Abbott, M Aguena, A Alarcon, S Allam, O Alves, A Amon, F Andrade-Oliveira, J Annis, S Avila, D Bacon, et al. Dark energy survey year 3 results: Cosmological constraints from galaxy clustering and weak lensing. *Physical Review D*, 105(2):023520, 2022.

[2] Amir Aghamousa, Jessica Aguilar, Steve Ahlen, Shadab Alam, Lori E Allen, Carlos Allende Prieto, James Annis, Stephen Bailey, Christophe Balland, Otger Ballester, et al. The desi experiment part i: science, targeting, and survey design. *arXiv preprint arXiv:1611.00036*, 2016.

[3] Nabila Aghanim, Yashar Akrami, Frederico Arroja, Mark Ashdown, J Aumont, Carlo Baccigalupi, M Ballardini, Anthony J Banday, RB Barreiro, Nicola Bartolo, et al. Planck 2018 results-i. overview and the cosmological legacy of planck. *Astronomy & Astrophysics*, 641:A1, 2020.

[4] Nabila Aghanim, Yashar Akrami, Mark Ashdown, J Aumont, C Baccigalupi, M Ballardini, AJ Banday, RB Barreiro, N Bartolo, S Basak, et al. Planck 2018 results-vi. cosmological parameters. *Astronomy & Astrophysics*, 641:A6, 2020.

[5] Kazuyuki Akitsu, Yin Li, and Teppei Okumura. Cosmological simulation in tides: power spectra, halo shape responses, and shape assembly bias. *Journal of Cosmology and Astroparticle Physics*, 2021(04):041, 2021.

[6] Benedict Bahr-Kalus, David Parkinson, and Eva-Maria Mueller. Measurement of the matter-radiation equality scale using the extended baryon oscillation spectroscopic survey quasar sample. *Monthly Notices of the Royal Astronomical Society*, page stad1867, 2023.

[7] Tobias Baldauf, Uroš Seljak, Leonardo Senatore, and Matias Zaldarriaga. Galaxy bias and non-linear structure formation in general relativity. *Journal of Cosmology and Astroparticle Physics*, 2011(10):031, 2011.

[8] Alexandre Barreira and Fabian Schmidt. Response approach to the matter power spectrum covariance. *Journal of Cosmology and Astroparticle Physics*, 2017(11):051, 2017.

[9] Alexandre Barreira and Fabian Schmidt. Responses in large-scale structure. *Journal of Cosmology and Astroparticle Physics*, 2017(06):053, 2017.

[10] Adrian E Bayer, Jia Liu, Ryo Terasawa, Alexandre Barreira, Yici Zhong, and Yu Feng. Super-sample covariance of the power spectrum, bispectrum, halos, voids, and their cross-covariances. *arXiv preprint arXiv:2210.15647*, 2022.

[11] Daniele Bertolini, Katelin Schutz, Mikhail P Solon, Jonathan R Walsh, and Kathryn M Zurek. Non-gaussian covariance of the matter power spectrum in the effective field theory of large scale structure. *Physical Review D*, 93(12):123505, 2016.

[12] A Blanchard and J Schneider. Gravitational lensing effect on the fluctuations of the cosmic background radiation. *Astronomy and Astrophysics (ISSN 0004-6361), vol. 184, no. 1-2, Oct. 1987, p. 1-6.*, 184:1–6, 1987.

[13] George R Blumenthal, SM Faber, Joel R Primack, and Martin J Rees. Formation of galaxies and large-scale structure with cold dark matter. *Nature*, 311(5986):517–525, 1984.

[14] Jordan Carlson, Beth Reid, and Martin White. Convolution lagrangian perturbation theory for biased tracers. *Monthly Notices of the Royal Astronomical Society*, 429(2):1674–1685, 2013.

[15] Sean M Carroll. The cosmological constant. *Living reviews in relativity*, 4(1):1–56, 2001.

[16] Julien Carron, Melody Wolk, and István Szapudi. On the information content of the matter power spectrum. *Monthly Notices of the Royal Astronomical Society*, 453(1):450–455, 2015.

[17] Shaun Cole. Adding long-wavelength power to n-body simulations. *Monthly Notices of the Royal Astronomical Society*, 286(1):38–47, 1997.

[18] Shaun Cole and Nick Kaiser. Biased clustering in the cold dark matter cosmogony. *Monthly Notices of the Royal Astronomical Society*, 237(4):1127–1146, 1989.

[19] Asantha Cooray and Wayne Hu. Power spectrum covariance of weak gravitational lensing. *The Astrophysical Journal*, 554(1):56, 2001.

[20] Liang Dai, Enrico Pajer, and Fabian Schmidt. On separate universes. *Journal of Cosmology and Astroparticle Physics*, 2015(10):059, 2015.

[21] Roland De Putter, Christian Wagner, Olga Mena, Licia Verde, and Will J Percival. Thinking outside the box: effects of modes larger than the survey on matter power spectrum covariance. *Journal of Cosmology and Astroparticle Physics*, 2012(04):019, 2012.

[22] Vincent Desjacques, Donghui Jeong, and Fabian Schmidt. Large-scale galaxy bias. *Physics reports*, 733:1–193, 2018.

[23] Eleonora Di Valentino, Olga Mena, Supriya Pan, Luca Visinelli, Weiqiang Yang, Alessandro Melchiorri, David F Mota, Adam G Riess, and Joseph Silk. In the realm of the hubble tension—a review of solutions. *Classical and Quantum Gravity*, 38(15):153001, 2021.

[24] Jaan Einasto, Ants Kaasik, and Enn Saar. Dynamic evidence on massive coronas of galaxies. *Nature*, 250(5464):309–310, 1974.

[25] Albert Einstein. Näherungsweise integration der feldgleichungen der gravitation. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, pages 688–696, 1916.

[26] Daniel J Eisenstein, Idit Zehavi, David W Hogg, Roman Scoccimarro, Michael R Blanton, Robert C Nichol, Ryan Scranton, Hee-Jong Seo, Max Tegmark, Zheng Zheng, et al. Detection of the baryon acoustic peak in the large-scale correlation function of sdss luminous red galaxies. *The Astrophysical Journal*, 633(2):560, 2005.

[27] Abdalla Elcio, Aboubrahim Amin, Agnello Adriano, Akrami Yashar, Alestas George, Aloni Daniel, Amendola Luca, Luis A Anchordoqui, Richard I Anderson, Arendse Nikki, et al. Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies. *Journal of High Energy Astrophysics*, 34:49–211, 2022.

[28] Hume A Feldman, Nick Kaiser, and John A Peacock. Power spectrum analysis of three-dimensional redshift surveys. *arXiv preprint astro-ph/9304022*, 1993.

[29] Yu Feng, Man-Yat Chu, Uroš Seljak, and Patrick McDonald. Fastpm: a new scheme for fast simulations of dark matter and haloes. *Monthly Notices of the Royal Astronomical Society*, 463(3):2273–2286, 2016.

[30] DJ Fixsen. The temperature of the cosmic microwave background. *The Astrophysical Journal*, 707(2):916, 2009.

[31] Alexander Friedmann. Über die krümmung des raumes. *Z. Phys.*, 10:377–386, 1922.

[32] Nickolay Y Gnedin, Andrey V Kravtsov, and Douglas H Rudd. Implementing the dc mode in cosmological simulations with supercomoving variables. *The Astrophysical Journal Supplement Series*, 194(2):46, 2011.

[33] David M Goldberg and Michael S Vogeley. Simulating voids. *The Astrophysical Journal*, 605(1):1, 2004.

[34] James E Gunn and J Richard Gott III. On the infall of matter into clusters of galaxies and some effects on their evolution. *Astrophysical Journal, vol. 176, p. 1*, 176:1, 1972.

[35] Andrew J. S. Hamilton, Christopher D. Rimes, and Roman Scoccimarro. On measuring the covariance matrix of the non-linear power spectrum from simulations. *Monthly Notices of the Royal Astronomical Society*, 371(3):1188–1204, sep 2006.

[36] Nick Hand, Yu Feng, Florian Beutler, Yin Li, Chirag Modi, Uroš Seljak, and Zachary Slepian. nbodykit: An open-source, massively parallel toolkit for large-scale structure. *The Astronomical Journal*, 156(4):160, 2018.

[37] Cullan Howlett, Marc Manera, and Will J Percival. L-picola: A parallel code for fast dark matter simulation. *Astronomy and Computing*, 12:109–126, 2015.

[38] Cullan Howlett and Will J Percival. Galaxy two-point covariance matrix estimation for next generation surveys. *Monthly Notices of the Royal Astronomical Society*, 472(4):4935–4952, 2017.

[39] Wayne Hu and Martin White. Power spectra estimation for weak lensing. *The Astrophysical Journal*, 554(1):67, 2001.

[40] YP Jing. Correcting for the alias effect when measuring the power spectrum using a fast fourier transform. *The Astrophysical Journal*, 620(2):559, 2005.

[41] N. Kaiser. On the spatial correlations of abell clusters. *The Astrophysics Journal*, 284:L9–L12, 1984.

[42] Rene Laureijs, J Amiaux, S Arduini, J-L Augueres, J Brinchmann, R Cole, M Cropper, C Dabin, L Duvet, A Ealet, et al. Euclid definition study report. *arXiv preprint arXiv:1110.3193*, 2011.

[43] Georges Lemaître. The expanding universe. *Monthly Notices of the Royal Astronomical Society, Vol. 91, p. 490-501*, 91:490–501, 1931.

[44] Yin Li, Wayne Hu, and Masahiro Takada. Super-sample covariance in simulations. *Phys. Rev. D*, 89:083519, Apr 2014.

[45] Nina A Maksimova, Lehman H Garrison, Daniel J Eisenstein, Boryana Hadzhiyska, Sownak Bose, and Thomas P Satterthwaite. Abacussummit: a massive set of high-accuracy, high-resolution n-body simulations. *Monthly Notices of the Royal Astronomical Society*, 508(3):4017–4037, 2021.

[46] Matthew C Martino and Ravi K Sheth. On the equivalence between the effective cosmology and excursion set treatments of environment. *Monthly Notices of the Royal Astronomical Society*, 394(4):2109–2112, 2009.

[47] Shogo Masaki, Takahiro Nishimichi, and Masahiro Takada. Anisotropic separate universe simulations. *Monthly Notices of the Royal Astronomical Society*, 496(1):483–496, 2020.

[48] Patrick McDonald. Toward a measurement of the cosmological geometry at z˜ 2: Predicting lyα forest correlation in three dimensions and the potential of future data sets. *The Astrophysical Journal*, 585(1):34, 2003.

[49] A Meiksin, Martin White, and JA Peacock. Baryonic signatures in large-scale structure. *Monthly Notices of the Royal Astronomical Society*, 304(4):851–864, 1999.

[50] Irshad Mohammed and Uroš Seljak. Analytic model for the matter power spectrum, its covariance matrix and baryonic effects. *Monthly Notices of the Royal Astronomical Society*, 445(4):3382–3400, 2014.

[51] Irshad Mohammed, Uroš Seljak, and Zvonimir Vlah. Perturbative approach to covariance matrix of the matter power spectrum. *Monthly Notices of the Royal Astronomical Society*, 466(1):780–797, 2017.

[52] Mark C Neyrinck. Removable matter-power-spectrum covariance from bias fluctuations. *The Astrophysical Journal*, 736(1):8, 2011.

[53] J Ostriker, P Peebles, and A Yahil. The size and mass of galaxies and the mass of the universe, 1974.

[54] PJE Peebles. Large-scale background temperature and mass fluctuations due to scale-invariant primeval perturbations, 1982.

[55] Will J Percival. Cosmological structure formation in a homogeneous dark energy background. *Astronomy & Astrophysics*, 443(3):819–830, 2005.

[56] WJ Percival, L Miller, and JA Peacock. An analytic model for the epoch of halo creation. *Monthly Notices of the Royal Astronomical Society*, 318(1):273–279, 2000.

[57] Saul Perlmutter, Goldhaber Aldering, Gerson Goldhaber, RA Knop, Peter Nugent, Patricia G Castro, Susana Deustua, Sebastien Fabbro, Ariel Goobar, Donald E Groom, et al. Measurements of ω and λ from 42 high-redshift supernovae. *The Astrophysical Journal*, 517(2):565, 1999.

[58] Adam G Riess, Alexei V Filippenko, Peter Challis, Alejandro Clocchiatti, Alan Diercks, Peter M Garnavich, Ron L Gilliland, Craig J Hogan, Saurabh Jha, Robert P Kirshner, et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The astronomical journal*, 116(3):1009, 1998.

[59] Adam G Riess, Wenlong Yuan, Lucas M Macri, Dan Scolnic, Dillon Brout, Stefano Casertano, David O Jones, Yukei Murakami, Gagandeep S Anand, Louise Breuval, et al. A comprehensive measurement of the local value of the hubble constant with 1 km s- 1 mpc- 1 uncertainty from the hubble space telescope and the sh0es team. *The Astrophysical journal letters*, 934(1):L7, 2022.

[60] Howard Percy Robertson. Kinematics and world-structure. *Astrophysical Journal, vol. 82, p. 284*, 82:284, 1935.

[61] Vera C Rubin, W Kent Ford Jr, and Norbert Thonnard. Rotational properties of 21 sc galaxies with a large range of luminosities and radii, from ngc 4605/r= 4kpc/to ugc 2885/r= 122 kpc. *Astrophysical Journal, Part 1, vol. 238, June 1, 1980, p. 471-487.*, 238:471–487, 1980.

[62] R. K. Sachs and A. M. Wolfe. Perturbations of a Cosmological Model and Angular Variations of the Microwave Background. *The Astrophysics Journal*, 147:73, January 1967.

[63] Andreas S Schmidt, Simon DM White, Fabian Schmidt, and Jens Stücker. Cosmological n-body simulations with a large-scale tidal field. *Monthly Notices of the Royal Astronomical Society*, 479(1):162–170, 2018.

[64] Roman Scoccimarro, Matias Zaldarriaga, and Lam Hui. Power spectrum correlations induced by nonlinear clustering. *The Astrophysical Journal*, 527(1):1, 1999.

[65] Emiliano Sefusatti, Martin Crocce, Roman Scoccimarro, and Hugh MP Couchman. Accurate estimators of correlation functions in fourier space. *Monthly Notices of the Royal Astronomical Society*, 460(4):3624–3636, 2016.

[66] Edwin Sirko. Initial conditions to cosmological N-body simulations, or, how to run an ensemble of simulations. *The Astrophysical Journal*, 634(2):728–743, dec 2005.

[67] Volker Springel. The cosmological simulation code gadget-2. *Monthly notices of the royal astronomical society*, 364(4):1105–1134, 2005.

[68] RA Sunyaev and Ya B Zeldovich. The observations of relic radiation as a test of the nature of x-ray radiation from the clusters of galaxies. *Comments on Astrophysics and Space Physics, Vol. 4, p. 173*, 4:173, 1972.

[69] RA Sunyaev and Ya B Zeldovich. The velocity of clusters of galaxies relative to the microwave background-the possibility of its measurement. *Monthly Notices of the Royal Astronomical Society, vol. 190, Feb. 1980, p. 413-420.*, 190:413–420, 1980.

[70] Rashid A Sunyaev and Ya B Zeldovich. Small-scale fluctuations of relic radiation. *Astrophysics and Space Science*, 7:3–19, 1970.

[71] Masahiro Takada and Wayne Hu. Power spectrum super-sample covariance. *Phys. Rev. D*, 87:123504, Jun 2013.

[72] Ryuichi Takahashi, Naoki Yoshida, Masahiro Takada, Takahiko Matsubara, Naoshi Sugiyama, Issha Kayo, Takahiro Nishimichi, Shun Saito, and Atsushi Taruya. Non-gaussian error contribution to likelihood analysis of the matter power spectrum. *The Astrophysical Journal*, 726(1):7, 2010.

[73] Ryuichi Takahashi, Naoki Yoshida, Masahiro Takada, Takahiko Matsubara, Naoshi Sugiyama, Issha Kayo, Atsushi J Nishizawa, Takahiro Nishimichi, Shun Saito, and Atsushi Taruya. Simulations of baryon acoustic oscillations. ii. covariance matrix of the matter power spectrum. *The Astrophysical Journal*, 700(1):479, 2009.

[74] Svetlin Tassev, Matias Zaldarriaga, and Daniel J Eisenstein. Solving large scale structure in ten easy steps with cola. *Journal of Cosmology and Astroparticle Physics*, 2013(06):036, 2013.

[75] Giuseppe Tormen and Edmund Bertschinger. Adding long-wavelength modes to an n-body simulation. *The Astrophysical Journal*, 472(1):14, 1996.

[76] Michele Trenti and Piet Hut. Gravitational n-body simulations. *arXiv preprint arXiv:0806.3950*, 2008.

[77] Christian Wagner, Fabian Schmidt, Chi-Ting Chiang, and Eiichiro Komatsu. Separate universe simulations. *Monthly Notices of the Royal Astronomical Society: Letters*, 448(1):L11–L15, 2015.

[78] Arthur Geoffrey Walker. On milne's theory of world-structure. *Proceedings of the London Mathematical Society*, 2(1):90–127, 1937.

[79] John Wishart. The generalised product moment distribution in samples from a normal multivariate population. *Biometrika*, pages 32–52, 1928.

[80] Ya B Zel'Dovich. Gravitational instability: An approximate theory for large density perturbations. *Astronomy and astrophysics*, 5:84–89, 1970.

[81] Zhongxu Zhai and Will J. Percival. Sample variance for supernovae distance measurements and the hubble tension. *Physical Review D*, 106(10), nov 2022.

[82] Fritz Zwicky. On the masses of nebulae and of clusters of nebulae. In *A Source Book in Astronomy and Astrophysics, 1900–1975*, pages 729–737. Harvard University Press, 1979.