



THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

Cosmological constraints with large-scale structure: from Λ CDM to light relics

Abbé May Whitford

Bachelor of Science (Honours, 1st Class)



0000-0001-5829-8637

*A thesis submitted for the degree of Doctor of Philosophy at
The University of Queensland in 2025
School of Mathematics and Physics*

Abstract

The large-scale structure of the Universe as we see it today is a direct consequence of the Universe's expansion history; this depends on the constituents that drive changes in the expansion, and gravity. As such, studying the properties of the large-scale structure (LSS) – via correlation functions of observed density fields, or summarizing statistics such as the bulk flow – allows us to perform tests of gravity and cosmological models that capture the Universe's fundamental constituents. Cosmological data also captures the physics of the Universe on the largest length scales, allowing us to test fundamental principles and physical laws in different parts of space and time. In this thesis, I study the large-scale structure to improve tests of the concordance model of cosmology, Λ CDM, and extensions to this model, through parameters that capture the effects of neutrinos in the early Universe.

Firstly, in Chapter 2 I introduce the fundamental principles and observations behind Λ CDM, and some of the current problems for this model. In Chapter 3, I focus on the theory underpinning observations of galaxy and peculiar velocity surveys, and relate these to observables. I also describe the observational impacts of neutrinos and non-standard light relics.

My research using peculiar velocities to estimate a bulk flow and test the performance of bulk flow estimators starts in Chapter 4. Bulk flows allow a test of the Λ CDM model using the LSS in the local Universe. I perform a detailed analysis of the ability of two popular estimators to accurately and precisely measure the bulk flow from mock data to better understand measurements from real data. Given the anomalous bulk flows seen in the literature that are in apparent tension with Λ CDM, this analysis is important to determine the true level of tension in these results. I find that the geometric properties of galaxy peculiar velocity surveys are unlikely to be responsible for unexpectedly large bulk flow measurements. The estimators themselves are generally accurate at estimating the bulk flow, albeit less so when there are significant non-linearities present. However, the precision of these estimators is often overestimated, and the true variance in the estimates needs to be accounted for when making measurements with real data. I used the two estimators I studied to measure the bulk flow from the CosmicFlows-4 dataset, and found the amplitude of it to be $428 \pm 108 \text{ km s}^{-1}$ at a depth of $173 h^{-1} \text{ Mpc}$, which has a 0.11% chance of occurring within the Λ CDM model; further research is needed to determine if this result is a statistical fluctuation or new physics.

In Chapter 5 and Chapter 6, I study the phase shift that free-streaming relics, such as neutrinos, induce in the Fourier-space signal of baryon acoustic oscillations (BAO), and also in the cosmic microwave background (CMB) power spectra. The phase shift is a direct consequence of the free-streaming properties of neutrinos or light relics that behave in this manner, and thus provides a unique measurement of the cosmic neutrino background (CvB) as well as a test of neutrino properties, complementary to particle experiments. In Chapter 5, I focus on the phase shift amplitude β_ϕ that is modulated by Standard Model free-streaming relics and is directly related to the effective number of neutrino species, N_{eff} . I show a measurement of β_ϕ in the Dark Energy Spectroscopic Instrument data release 1 (DR1) BAO data after ensuring the robustness of the measurement pipeline with realistic mocks. This measurement finds $\beta_\phi = 2.7 \pm 1.7$ when applied to the DESI data alone, consistent with

the Standard Model expectation of $\beta_\phi = 1$ for $N_{\text{eff}} = 3.044$. When including a prior on parameters degenerate with β_ϕ by including information from *Planck* CMB data, the error bars shrink to $2.7^{+0.60}_{-0.67}$, creating tension with the expectation that $N_{\text{eff}} = 3.044$. I repeat the measurement with DESI data release 2 (DR2) data and find a more consistent measurement with $\beta_\phi = 2.00 \pm 0.78$. In this work, it is identified that higher precision measurements of β_ϕ from future data will require improvements in the modelling of the broadband power spectrum signal for the BAO analysis.

In Chapter 6, I construct a template for neutrinos with self-interactions that delay free-streaming compared to Standard Model neutrinos; this requires additional parameters to model the phase shift when it also depends on G_{eff} , the strength of self-interactions in a universal self-interaction model. When neutrinos self-interact their free-streaming is delayed, altering the functional form and amplitude of the phase shift. My analysis finds it is unlikely BAO data will ever be able to strongly constrain G_{eff} from the phase shift alone, but both BAO and CMB data may be able to at least show or disfavour a preference for neutrinos with strong self-interactions. Combining CMB and BAO data can result in the best constraints. This will be useful when applied to future data, to aid in removing bimodality in constraints on G_{eff} , which suffers from degeneracies with other cosmological parameters.

I summarize these analyses and conclude with recommendations for future work in Chapter 7. With upcoming peculiar velocity data from surveys that will cover the sky in the Southern Hemisphere, we will be able to get a more complete picture of the bulk motion of galaxies in the local Universe and further explore the tension seen in bulk flow measurements. Additionally, applying further study of the phase shift in the context of both Standard Model neutrinos (to the full DESI five year dataset) and neutrinos with self-interactions (to CMB data) may reveal interesting constraints, and allow us to better understand the hints of non-standard neutrino properties seen in cosmological data.

Declaration by Author

This thesis is composed of my original work, and contains no material previously published or written by another person, or generated by Artificial Intelligence (AI), except where due reference has been made in the text or the preliminary pages of this thesis. I have clearly stated the contribution by others to jointly authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, use of AI tools, financial support and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my higher degree by research candidature and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that an electronic copy of my thesis must be lodged with the University Library and, subject to the policy and procedures of The University of Queensland, the thesis be made available for research and study in accordance with the Copyright Act 1968 unless a period of embargo has been approved by the Dean of the Graduate School.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material. Where appropriate I have obtained copyright permission from the copyright holder to reproduce material in this thesis and have sought permission from co-authors for any jointly authored works included in the thesis.

Publications Included in this Thesis

1. Whitford et al. (2023) **Abbé M. Whitford**, Cullan Howlett, Tamara M. Davis, Evaluating bulk flow estimators for CosmicFlows-4 Measurements, *Monthly Notices of the Royal Astronomical Society*, 526, 2, 2023
2. Whitford et al. (2025) **Abbé M. Whitford**, Hugo Rivera-Morales, Cullan Howlett et al Constraining the phase shift of relativistic species in DESI BAOs, *Monthly Notices of the Royal Astronomical Society*, 538, 3, 2025

Manuscripts

1. **Abbé M. Whitford**, Cullan Howlett, Tamara M. Davis, Camarena, David and Cyr-Racine, Francis-Yan, *Limits of self-interacting neutrinos from the BAO and CMB phase shift*, prepared for submission to JCAP

Other Publications During Candidature

Journal articles

1. Whitford et al. (2022) **Abbé M. Whitford**, Cullan Howlett, Tamara M. Davis, Using peculiar velocity surveys to constrain neutrino masses, *Monthly Notices of the Royal Astronomical Society*, 513, 1, 2022
2. Elbers et al. (2025) **DESI Collaboration, Willem Elbers et al (Including Abbé M. Whitford)** Constraints on Neutrino Physics from DESI DR2 BAO and DR1 Full Shape, 2025

Conference proceedings

1. Whitford (2024) **Abbé M. Whitford** Using the peculiar velocities of galaxies to test cosmology and particle physics, *Rencontres du Vietnam 2023: Windows on the Universe, Quy Nhon, Vietnam, August 06-12 2023*

Contributions by Others to the Thesis

All contributions by collaborating researchers are described prior to each chapter.

Use of Artificial Intelligence

No writing of the text or scientific results produced in this thesis involved the use of artificial intelligence. I have used Github co-pilot (an AI tool for coding) for automatic code completion, and I have

used ChatGPT at times to assist in the search of relevant information from literature articles and usage of coding packages that I have cited in this thesis. For example, “Can you use the CLASS Boltzmann solver version 3.1 to solve for the Silk Damping scale?” “Can you give me the expression for σ_8 in terms of the power spectrum?” I have also occasionally used ChatGPT to help suggest fixes for bugs in code by sharing an error message. Additionally, ChatGPT was used by my coauthors of the draft of the manuscript included in Chapter 6 to make suggestions for clarity or conciseness in a few paragraphs of the text.

Editorial Assistance

No editorial assistance.

Financial Support

This research was supported by an Australian Government Research Training Program Scholarship and a Federation of Women University Fellowship top-up Scholarship. During the completion of this research funding support was provided to attend conferences and meetings through the Australian Research Council’s Laureate Fellowship funding scheme (FL180100168), and through the Discovery Project (project DP20220101395) funding scheme and Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), through project number CE230100016. Additionally, financial support was provided by the Dark Energy Spectroscopic Instrument Collaboration to attend their collaboration meetings.

Inclusion of Work Submitted for Another Degree

No works submitted towards another degree have been included in this thesis.

Research involving Human or Animal Subjects

No animal or human subjects were involved in this research.

Acknowledgments

I thank the various scientists I met through collaboration meetings and schools in various parts of the world throughout the completion of this degree. I am grateful for all the support and education provided by the DESI collaboration, the Galileo Galilei Institute of Theoretical Physics in Florence, Italy, the OzGrav Collaboration, the Astronomical Society of Australia, the ICISE in Qhoy Nhon, Vietnam and the ICTP in Sao Paulo, Brazil.

The Astrobites collaboration gave me the chance to write about fascinating research articles through their platform as a graduate student while completing this degree. This has been a fun and useful opportunity to learn more from the academic literature and about scientific communication which has supported my thesis.

I thank my collaborators Hugo Rivera-Morales and Mariana Vargas-Magaña for their teamwork on a project that makes up a significant portion of this thesis. They contributed greatly to complete this project, and worked hard to communicate with us despite the significant time difference between our locations in Australia and Mexico.

I thank my supervisors Tamara and Cullan for their support and their feedback throughout the completion of every part of this degree and the writing of this thesis; without their support, this would not have been possible.

I have also indirectly received support from numerous individuals since beginning this degree. My family, in particular my grandparents Adele and John, have supported me to follow anything I have wanted and have always believed in my ability to succeed. My siblings, Lauren, Madeline, Jacob, Nicholas and my father Cameron have also supported me while I completed this degree. My late mother Rachel supported me in my studies leading up to commencing this degree. My partner Adam has also supported me endlessly throughout my time as a student. My numerous friends, both those in my office and outside of it, have supported me with their friendship over the years and allowed me to maintain some sanity.

This thesis is dedicated to Adele and John.

Contents

Abstract	ii
Contents	ix
List of Figures	xii
List of Tables	xv
List of Abbreviations and Symbols	xvii
1 Introduction	1
2 Modern cosmology	5
2.1 The expanding Universe	5
2.2 Distances in cosmology	9
2.3 Cosmological probes	11
2.3.1 Supernovae	11
2.3.2 The Cosmic Microwave Background	12
2.3.3 Galaxy surveys	13
2.3.4 Peculiar velocities	15
2.3.5 Other probes	16
2.4 The status of modern cosmology	16
2.4.1 Beyond Λ CDM	17
3 Theory of LLS with galaxy and peculiar velocity surveys	23
3.1 Galaxy surveys	23
3.1.1 The matter density field and velocity field	23
3.1.2 Galaxy correlation functions	27
3.1.3 The galaxy power spectrum	27
3.1.4 Galaxy surveys in redshift space	31
3.1.5 Analysis with baryon acoustic oscillations	35
3.1.6 Galaxy bulk flows	37
3.2 Neutrinos in cosmology	40

3.2.1	Impact of neutrinos in the early Universe	40
3.2.2	Impact of light relics on structure growth	40
3.2.3	Massive neutrinos in cosmology	43
3.2.4	Neutrino self-interactions in cosmology	45
4	Evaluating bulk flow estimators	51
4.1	Overview	52
4.2	Datasets and mock catalogues	54
4.2.1	The Sloan Digital Sky survey peculiar velocity catalogue	54
4.2.2	The CosmicFlows IV Tully–Fisher peculiar velocity catalogue	55
4.2.3	The 6-degree Field Galaxy survey PV catalogue	55
4.2.4	The combined CF4 dataset and mocks	56
4.3	Systematics in bulk flow measurements	58
4.3.1	Discussion of methods	60
4.4	Evaluation of bulk flow estimators	64
4.4.1	Tests on simplistic mocks	65
4.4.2	Tests on mocks with linear structure growth	68
4.4.3	Tests on realistic mocks for CF4 data	72
4.4.4	Summary: mock performance evaluation	79
4.5	Measured bulk flows from real data	80
4.5.1	Results: CosmicFlows 4 dataset	80
4.5.2	Comparison to bulk flows from previous literature	86
4.6	Conclusions	87
5	Constraining the phase shift in DESI BAOs	93
5.1	Overview	94
5.2	Methodology	98
5.2.1	Isotropic BAO analysis	98
5.2.2	Anisotropic BAO analysis	99
5.2.3	Anisotropic BAO analysis with β_ϕ	100
5.2.4	Implementation	102
5.3	Mock validation	105
5.3.1	Fitting DESI first-generation mocks	106
5.3.2	Fitting DESI second-generation mocks with DESILIKE	111
5.4	Results with DR1 DESI data	115
5.4.1	Constraints from individual tracers	115
5.4.2	Combined constraints and cosmological interpretation	118
5.5	Conclusions	124
5.6	Extended results with DR2 DESI data	126
5.6.1	Constraints from individual tracers	126

5.6.2	Combined constraints	127
6	Limits on self-interacting neutrinos from the phase shift	131
6.1	Overview	131
6.1.1	Searching for neutrino physics beyond the Standard Model	132
6.1.2	Studying the phase shift of self-interacting neutrinos: motivation	133
6.2	Cosmological impacts of neutrino self-interactions	135
6.3	The phase shift of Standard Model neutrinos	139
6.3.1	Baryon acoustic oscillations	139
6.3.2	Cosmic microwave background	140
6.4	The phase shift with non-standard model neutrino physics	142
6.4.1	Baryon Acoustic Oscillations	143
6.4.2	Cosmic microwave background	146
6.5	Can we constrain G_{eff} from the phase of the BAO?	148
6.5.1	Profile of χ^2	149
6.5.2	Fisher matrix forecasts	150
6.6	Can we constrain G_{eff} from the phase of the CMB power spectra?	153
6.6.1	Profile of χ^2	153
6.6.2	Fisher matrix forecasts	154
6.6.3	Combined BAO and CMB forecasts	156
6.7	Conclusion	157
7	Conclusion	161
7.1	Conclusion	161
7.2	Future work	164
7.2.1	Bulk flows	164
7.2.2	The phase shift induced by free-streaming relics	165
	Bibliography	167
A	Appendices	187
A.1	Linear theory velocity covariance matrix	187
A.2	Biases in galaxy surveys	187
A.3	Validating the anisotropic pipeline by comparison to Baumann et al (2019)	188
A.4	Sampling methods in cosmology	189
A.5	Additional comparisons from <i>Barry</i>	190
A.6	Fisher information and forecasts	190
A.7	The effect of N_{eff} on the phase shift	191

List of Figures

1.1	A slice of the galaxies in the DESI DR1 dataset, up to $z \sim 0.2$	2
2.1	The relative densities of radiation, matter and dark energy over the history of the Universe.	8
2.2	A plot of comoving distance, $d_C(z)$, luminosity distance, $d_L(z)$ and angular distance, $d_A(z)$, as a function of redshift z	10
2.3	A plot of the matter power spectrum and correlation function.	14
2.4	A summary plot showing bulk flow amplitudes measured from various authors in the literature.	18
2.5	Measurements of the ratio $\frac{d(z)H(z)}{cz}$ obtained from DESI DR2 BAO measurements, plotted as a function of the effective redshift z_{eff}	19
2.6	Left panel: Constraints on the ‘effective’ sum of neutrino masses, taken directly from Elbers et al. (2025). Right panel: Bimodal constraints on the strength of universal neutrino self-interactions in data from various CMB experiments, taken directly from Kreisch et al. (2024).	22
3.1	The growth rate of structure $f(a)$, shown as a function of scalefactor a	26
3.2	The impact of varying cosmological parameters on the matter power spectrum.	30
3.3	The impact of redshift space distortions on the correlation function, due to the Kaiser effect and the FoG (Finger-of-God) effect.	33
3.4	The oscillatory BAO feature, computed from a Λ CDM model matter power spectrum with $\Omega_m = 0.31$, $H_0 = 67 \text{ km s}^{-1} \text{ Mpc}^{-1}$	35
3.5	A demonstration of the Alcock-Paczynski (AP) Effect, shown in the galaxy correlation function.	37
3.6	The CMB temperature power spectrum with varying properties of neutrinos.	42
3.7	A summary plot of constraints in the $m_\phi - g_\nu$ plane for neutrino self-interactions, taken directly from Berryman et al. (2023).	49
4.1	Individual selection functions for SDSS, CF4TF, SDSS and CF4.	56
4.2	The sky coverage of the CF4 dataset in galactic coordinates.	57
4.3	The theoretical bulk flow V_p as a function of the survey radius R for two survey geometries.	59
4.4	The average of the recovered bulk flows from simplistic mocks with a cone geometry.	67

4.5	The difference between the recovered bulk flow components from 512 simplistic mocks using two different estimator methods.	67
4.6	The impact of distance on the scatter in the observed values of log-distance ratios.	72
4.7	Observed trend between sample kurtosis of the observed peculiar velocities and the reduced chi-squared best fit to the bulk flow.	73
4.8	Recovered bulk flow moments for realistic mocks using the Kaiser MLE and Peery MVE.	74
4.9	Bulk flow residuals when applying the Kaiser MLE to Zeldovich mocks vs mocks with non-linear structure growth.	75
4.10	The recovered bulk flow moments from realistic mocks using the Kaiser MLE when a global zero-point offset is applied.	78
4.11	The recovered bulk flow moments from realistic mocks using the Peery MVE when a relative zero-point offset is applied to CF4TF data.	78
4.12	Estimated bulk flow components in Supergalactic coordinates and bulk flow amplitude from the CF4 data and subsets using the Kaiser MLE.	82
4.13	Estimated bulk flow components in Supergalactic coordinates and bulk flow amplitude from the CF4 data and subsets using the Peery MVE.	83
4.14	A map of the CF4 dataset in Supergalactic coordinates, each panel showing a 2D projection of two coordinates.	83
4.15	Comparison of different bulk flows of datasets shown in the literature, converted to Supergalactic coordinates and listed in order of decreasing effective depth.	87
5.1	The effect of altering N_{eff} on the matter power spectrum, the BAOs, and the isolated phase shift.	96
5.2	The effect of altering N_{eff} on the correlation function.	97
5.3	Comparison of fits to first-generation mocks for BARRY and DESILIKE.	107
5.4	Fits to the mean of 25 first-generation CV mocks using BARRY, comparing spline and polynomial broadband methodologies.	108
5.5	Fits to the mean of 25 first-generation CV mocks with the polynomial broadband methodology, when the fiducial cosmology is set to c003.	111
5.6	Fits to the mean of 25 second-generation mocks and the mock mean, for the 6 DESI tracers.	114
5.7	The fits to β_ϕ from the combinations of the LRGs and ELGs for the second-generation mocks.	116
5.8	Fits to the BAO distortion and phase shift parameters for the tracers in DESI DR1.	117
5.9	The best fit models for the DESI tracers in the DR1 data with β_ϕ varying in the analysis, for comparison to the data.	120
5.10	Posterior distributions for α , α_{AP} , β_ϕ , Ω_m and $r_s h$ for the baseline fits to the DESI DR1 data.	121
5.11	Combined constraints on β_ϕ and cosmological parameters with DESI tracers, including fits with priors from <i>Planck</i>	122

5.12	Posteriors for fits to the BAO and phase shift parameters in the DESI DR2 data for individual tracers.	127
5.13	Combined constraints to β_ϕ for the DESI DR2 data, in comparison to the DR1 data and the BOSS DR12 fits.	128
6.1	The impact of G_{eff} (the interaction strength of neutrinos) on the matter power spectrum.	136
6.2	The effect of neutrino self-interactions and thus delayed free-streaming on the CMB temperature-temperature power spectrum.	137
6.3	The phase shift function $f(k)$ extracted from the BAO signal, for Standard Model neutrinos.	140
6.4	The scale dependent function $f(\ell)$ that captures the phase shift in the CMB power spectra, for Standard Model neutrinos.	141
6.5	The phase shift signal in the BAO, for different fixed values of G_{eff}	143
6.6	The fit to the phase shift amplitude and exponential damping parameters A and B in the BAO signal, as a function of $\log_{10}(G_{\text{eff}})$	145
6.7	The phase shift signal for different fixed values of G_{eff} in the CMB power spectra.	147
6.8	The fit to A and B in the CMB power spectra as a function of G_{eff}	147
6.9	a) The $\Delta\chi^2$ from a fit of BAO data to Λ CDM vs models with varying G_{eff} , and b) the ratio of the BAO oscillations in models with varying G_{eff} , with error bars for a volume-limited survey.	149
6.10	a) The difference between the χ^2 goodness-of-fit with Λ CDM and the χ^2 goodness-of-fit for a model with neutrino self-interactions parameterized by G_{eff} , in a cosmic-variance limited CMB experiment. b) A plot of the CMB TT spectrum after subtracting spectra with varying choices of G_{eff} or β_ϕ which changes the phase shift.	154
A.1	Comparison of fits to first-generation mocks for the pipeline of Baumann et al. (2019), an extended version of this pipeline, and <i>desilike</i>	188
A.2	Comparison of fits when the phase shift is allowed to vary vs. when it is fixed, with various modelling choices, using <i>barry</i>	190
A.3	The phase shift $\Delta\ell$, demonstrated in the CMB temperature-temperature power spectrum, when varying $\log_{10}(G_{\text{eff}})$, for four fixed choices of N_{eff}	192
A.4	The neutrino-induced phase shift in the CMB, for Standard Model neutrinos.	193

List of Tables

4.1	Properties of each peculiar velocity dataset within CF4.	56
4.2	Results for the reduced χ^2 goodness-of-fit for the bulk flow when applying the Peery MVE for 512 mocks for the CF4 data, compared to the chosen radius R of the ideal survey used for the method.	76
4.3	Estimated bulk flow components in Supergalactic coordinates from the Peery MVE and the Kaiser MLE.	81
5.1	Summary of DESI DR1 data properties.	105
5.2	The difference in the best fit to β_ϕ , when different broadband methodologies are used, for first-generation mocks.	109
5.3	A comparison of the absolute difference to the best fits to the BAO and phase shift parameters for the mean of the first-generation CV mocks, when the fitting pipeline is varied.	110
5.4	Average of the difference in the fits to α , α_{AP} and β_ϕ between Barry and desilike, for second-generation mocks.	113
5.5	Summary of the choices of priors and ranges used for fitting the second-generation mocks and DR1 data for each DESI tracer.	113
5.6	Constraints on β_ϕ , α and α_{AP} from the DR1 data baseline results.	119
5.7	The χ^2 and degrees-of-freedom (dof) for each of the Baseline fits to the different DESI tracers when β_ϕ is allowed to vary.	119
5.8	Constraints on β_ϕ from the combined DESI DR1 BAO tracers, alongside cases when a prior on the BAO α parameters has been included from the <i>Planck</i> data.	123
5.9	Summary of DESI DR2 data tracers properties.	126
5.10	Fits to the BAO and phase shift parameters in the DESI DR2 data for individual tracers.	127
6.1	Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a volume-limited BAO survey.	152
6.2	Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for the DESI 5-year BAO dataset.	152
6.3	Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a cosmic-variance limited CMB measurement of the TT, EE and TE power spectra.	156

6.4	Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a <i>Planck</i> -like CMB measurement of the TT, EE and TE power spectra.	156
6.5	Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a <i>Planck</i> -like CMB measurement combined with BAO data from DESI YR5.	157

List of Abbreviations and Symbols

Abbreviations	
LSS	Large-scale structure
CMB	cosmic microwave background
CvB	cosmic neutrino background
BAOs	baryon acoustic oscillations
CLASS	the Cosmic Linear Anisotropy Solving System (a boltzmann-solver code)
CAMB	Code for Anisotropies in the Microwave Background (a boltzmann-solver code)
DESI	the Dark Energy Spectroscopic Instrument
BOSS/eBOSS	the Baryon Oscillation Spectroscopic Survey (and extended BOSS)
BF	bulk flow
Λ CDM	the current concordance cosmological model of cosmology
CDM	cold dark matter
PV	peculiar velocity
ISW, SW	the Integrated Sachs-Wolfe Effect, Sachs-Wolfe effect
Ly α	the Lyman- α emission line
QSOs	quasars (or quasi-stellar objects)
ELGs	emission line galaxies
LRGs	luminous red galaxies
BGS	the DESI Bright Galaxy Survey
RSDs	redshift-space distortions
AP effect	the Alcock-Paczyński effect
CF4	the CosmicFlows-4 PV catalogue of galaxies
MLE	Maximum Likelihood estimator (for bulk flows)
MVE	the Minimum Variance estimator (for bulk flows)
MCMC	Markov Chain Monte Carlo
PDF	probability distribution function
recon/ pre-recon	post-reconstruction, pre-reconstruction
DR	data-release
SM, BSM	the Standard Model, Beyond the Standard Model

Symbols	
c	the speed of light
G	the universal gravitational constant
k_B	the Boltzmann constant
ρ_i	the energy density of the quantity labelled by i
$\rho_{\text{cr.}}$	the critical energy density that halts expansion or collapse in a purely matter dominated Universe
Ω_i	the energy density of the quantity labelled by i , normalized by $\rho_{\text{cr.}}$
$a(t)$	the dimensionless scale factor of the Universe, defined such that $a(t=0) = 0$ and $a(t_{\text{today}}) = 1$
w_i	the equation of state of the quantity labelled by i
H_0	the Hubble constant
h	‘little h’, a dimensionless quantity defined such that $H_0 = 100h\text{kms}^{-1}$
$H(t)$	the Hubble constant at time t
z	redshift
Λ	a cosmological constant (dark energy)
$\sum m_\nu$	the sum of neutrino masses
N_{eff}	the effective number of neutrino species
τ_{re}	the optical depth to reionization (optical depth of Thomson scattering)
A_s	the amplitude/normalization of the matter power spectrum
n_s	the spectral index or ‘tilt’ of the matter power spectrum
$\sigma_8(z)$	the dispersion of matter fluctuations on scales of $8\text{Mpc}h^{-1}$ at z
$\delta(x)$	the relative matter density of space compared to the average (or density contrast or density field)
$u(x)$	the peculiar velocity field
$\theta(x)$	the velocity divergence field
$P(k)$	the power spectrum
$\xi(r)$	the correlation function
a_{eq}	the scale factor at the epoch of matter-radiation equality
$D(z)$	the linear growth factor at z
$f(k, z)$	the growth rate of structure at z and wavenumber k
K	the dimensionless parameter for the curvature of the Universe
b_g	the linear galaxy bias
$\alpha, \alpha_{\text{AP}}$	the isotropic BAO distortion parameter, the anisotropic distortion parameter
β_ϕ	the BAO phase-shift amplitude
χ^2	the goodness-of-fit
B_i	the i th component of the BF (bulk flow)
$d(z)$	the comoving radial distance to a galaxy at a cosmological redshift z
C_{ij}	the i th, j th component of the covariance matrix
$r_s(z_d)$	the BAO scale (sound horizon scale at the baryon-drag epoch)
$\mathcal{P}_l(\mu)$	the l th Legendre polynomial
G_{eff}	the ‘effective’ strength of neutrino self-interactions

Chapter 1

Introduction

The distribution and motions of galaxies and matter in the Universe contain a wealth of information to be extracted; from the overall amount and evolution of different materials that are present (matter, radiation, dark energy) to the properties of particles that comprise the Standard Model (such as the masses of neutrinos). The growing understanding of how we can use this information to learn about the Universe and its expansion history has led to a field known as *large-scale structure (LSS) cosmology*. At the time of writing, surveys such as the Sloan Digital Sky Survey (SDSS), the Baryon Oscillation Spectroscopic Survey (BOSS) and the Dark Energy Spectroscopic Instrument (DESI) (Aghamousa et al., 2016a,b; Abareschi et al., 2022; Adame et al., 2024; Dawson et al., 2012; York et al., 2000) have collected measurements of the spatial distribution of over 10 million galaxies. Figure 1.1 shows the LSS in galaxies up to $z \sim 0.2$ in the DESI DR1 (Data Release 1). The cosmology community has also collected information on the motions of over $\sim 50,000$ objects in our local patch of the Universe (Tully et al., 2023). With this information and that we expect to obtain with continuing and future galaxy surveys, we can try answer questions like ‘What are dark matter and dark energy?’; ‘How fast is the Universe expanding?’; ‘How strongly are galaxies tugging on each other?’; ‘How does matter in galaxies clump together, and what does this tell us about the early Universe?’ In my thesis, I attempt to make a contribution to understanding questions such as these, using measurements of the net flows of galaxies in our local Universe, and imprints of the early Universe sound waves in the clustering of distant galaxies.

In Chapter 2, I introduce the fundamental mathematics and theory that supports modern cosmology, including the Cosmological Principle, and discuss the Friedmann equations that describe the expansion history of the Universe starting from the Friedmann-Lemaître-Robertson-Walker (FLRW) metric. These equations allow one to compare measurements of the positions of objects over the history of the Universe to a model and make measurements of the amount of each component that comprises the total material in the Universe. I then describe distances in cosmology and how different approaches to measuring distances should be compared theoretically to a model. Furthermore, I introduce some of the observations that have led to the currently accepted concordance model, Λ CDM, that describes our Universe, and some of the cracks that have begun appearing in this model with increasingly precise




Figure 1.1: A slice of the galaxies in the DESI DR1 dataset, up to $z \sim 0.2$. I produced this figure using code provided in tutorials to access the DESI DR1 data by the DESI collaboration, which can be found https://github.com/desihub/dr1paper/blob/main/nb/DESI_Y1_butterfly.ipynb. The different coloured points represent galaxies with different angular positions in a wedge of finite width.

observations.

Chapter 3 introduces the theory that is relevant to my research on the large-scale structure of the Universe. Particularly, I discuss the observations and statistical information that are drawn from galaxy surveys that measure redshifts of galaxies or peculiar motions. These probes contain information about the expansion history and distribution of matter close to the present day, which can be extracted in various ways from redshift space distortions, the baryon acoustic oscillations, the Alcock-Paczynski effect, and the bulk motions of galaxies. I also explain the signatures of massive neutrinos and free-streaming particles in observations of large-scale structure data.

In Chapter 4 I discuss my research on measurements of the bulk motions of galaxies, which is an average of the motion of many galaxies in a volume, due to local gravitational interactions. The bulk flow can be used to test cosmological models by comparing it to the expected distribution of bulk flows in a given model. I give an overview of different biases that can arise in bulk flow measurements, issues in comparing the measurements to theory and pros and cons of different methods to measure the bulk flow. For the main focus of this chapter I assess the ability of popular bulk flow *estimators* in the literature to accurately and precisely measure the bulk flow with simulations. This is interesting in

the context of the current bulk flow *tension*, which refers to a number of bulk flow measurements in the literature that are in tension with predictions from Λ CDM (Kashlinsky et al., 2008; Watkins et al., 2009; Feldman et al., 2010; Watkins and Feldman, 2015b; Peery et al., 2018; Howlett et al., 2022; Watkins et al., 2023; Whitford et al., 2023). My findings suggest this tension is unlikely to be due to problems related to estimators.

In Chapter 5, we pivot our discussion to measurements of the baryon acoustic oscillations (BAOs), which are the sound waves that propagated through the hot plasma of baryons and photons in the early Universe. These sound waves leave a detectable signature in the distribution of galaxies in the Universe, as an excess clustering of galaxies on scales corresponding to the largest possible wavelength that could be present in the waves. The BAOs have historically been used to tell us about the Universe's expansion dynamics over time by measuring their signature at different times in the Universe's history. However, they are also sensitive to the contribution of neutrinos to the radiation density in the very early Universe via a phase shift that is induced in the signal of the BAOs in Fourier space (Bashinsky and Seljak, 2004; Green and Ridgway, 2020). This phase shift can allow one to derive constraints on the number of neutrino species (or presence of sterile neutrinos that free-stream) in the early Universe. In this chapter, I explain how this signal arises, discuss my research to validate fitting codes to measure this signal on mock data, before applying the pipeline to the most precise measurement of the BAOs to date in data collected by DESI. We find the phase shift in the DESI data is consistent with the Standard Model expectation for free-streaming neutrinos given the measurement precision, although there is some preference for a larger N_{eff} .

In Chapter 6 I further explore whether the phase shift signal in the BAO signal and in Cosmic Microwave Background (CMB) measurements can also be used to look for signatures of beyond-the-standard model neutrinos; this is interesting given cosmological measurements in the literature that have found a preference for neutrinos with self-interactions (Kreisch et al., 2020; Camarena et al., 2023a). I explore the change in the phase of these signals using existing codes that predict cosmological observables for a Universe that contains neutrino with self-interactions in a simple model. I explore whether it may eventually be possible to use the phase shift measured in CMB or BAO data to put constraints not only the presence of free-streaming species but also the strength of their interactions, and the possibility of using this data to complement constraints from existing terrestrial experiments.

Finally, I conclude in Chapter 7 with a summary of the future directions for the research topics covered. It is expected upcoming data from the 4MOST Hemisphere Survey, which covers more of the Southern sky than previously measured, will help us to further test Λ CDM by covering more regions that are not currently probed by galaxy velocity measurements for measuring the bulk flow. For light relics, it will be interesting to further explore the possibilities for signatures that can be constrained with cosmological data, bridging cosmology and particle physics to enhance our understanding of the Universe and the Standard Model.

Chapter 2

Modern cosmology

Cosmology is the field of the astronomy concerned with the study of the expansion history of the Universe and the events that led to our present day Universe. Cosmologists seek knowledge of its constituent components and how they influence the expansion dynamics and formation of structures such as stars and galaxies. As such, cosmology allows us to probe fundamental physics; studying the motions of galaxies in clusters and individual galaxy rotation curves led astronomers to discover the presence of *dark matter* (Zwicky, 1933, 1937; Rubin and Ford Jr, 1970), while studying the expansion history allowed cosmologists to discover *dark energy* (Riess et al., 1998; Perlmutter et al., 1999). The existence of dark matter and dark energy are the pillars of the concordance model of cosmology, Λ CDM model, where the Λ represents dark energy in the form of a cosmological constant and CDM stands for cold dark matter.

In this chapter I give an overview of the fundamental principles of modern cosmology and the observations that led to the development of the Λ CDM model. This then leads into a discussion of using galaxy redshift and peculiar velocity surveys with cosmic microwave background (CMB) information to inform on the nature of the Universe and its formation history in Chapter 3, providing the required background for the following chapters.

2.1 The expanding Universe

Prior to the discoveries of dark matter and dark energy, Lemaître (1927) and Hubble (1929) discovered the first evidence that the Universe is not static, which was the consensus at the time, but is in fact expanding; Friedman (1922) and Lemaître showed that an expanding Universe is a prediction of General Relativity. The observations used by Hubble showed that the spectra of distant galaxies are *redshifted* (the wavelengths of their spectral emission lines or absorption lines are shifted to larger wavelengths), and this redshift increases with the distance to the galaxy being observed. As a redshift indicates motion away from an observer, this result suggests that more distant galaxies have a larger recession velocity; this is true for an observer situated at any point in space. Initially, the redshift of each galaxy was interpreted to be a result of the *Doppler* effect due to the galaxy itself moving.

However, it has since been established that the galaxy redshift is not so simply described by the Doppler shift equation for the galaxy velocity at an instant; instead, the redshift arises because photons lose energy as they traverse an expanding spacetime which gives the galaxy its recessional motion. The redshift can be calculated as an integral of infinitesimal Doppler shifts along the line of sight, as the photon's energy is measured with respect to receding comoving points in space (Bunn and Hogg, 2009). Today, the relationship between receding galaxies with distance is known as the *Hubble-Lemaître law*.

Under the assumption of *the Cosmological Principle*, which states the Universe is homogeneous and isotropic,¹ one can derive the Friedmann equations that give information about the dynamics and expansion history of the Universe. The derivation here is inspired by Carroll (2019). Firstly, one must start with the Friedmann-Lemaître Robertson Walker (FLRW) metric,

$$ds^2 = -c^2 dt^2 + R^2(t) \left(\frac{dr^2}{1-kr^2} + r^2 d\Omega^2 \right). \quad (2.1)$$

This metric is maximally symmetric, is evolving with time t , and is scaled by the dimensionful scale factor $R(t)$ that gives the relative size of the Universe. ds gives the spacetime interval which is invariant under coordinate transformations. The dimensionless coordinate r gives the radial distance to an object (the alteration to the distance with the expansion of the Universe, $R(t)$, is factored out)² Ω gives the solid angle the object has in the sky, in a spherical-polar coordinate system. k determines the geometry of the Universe by giving the spacetime *curvature*, and can take values of $\{+1, 0, -1\}$. One can then use the metric to calculate the components of the *Ricci-curvature tensor* $R_{\mu\nu}$ (a tensor that arises in General Relativity and dictates the geometry of spacetime) from the metric tensor $g_{\mu\nu}$ (this can be extracted from the FLRW metric, equation 2.1). The other required equation is the *energy-momentum tensor* $T^{\mu\nu}$, defined as the flux of 4-momentum μ across a 3D surface at a fixed coordinate x^ν ; this gives the important energy-like aspects of a system and is essential to the Einstein field equations that govern the dynamics of a system in General Relativity. For a universe that is composed by materials that are completely described by their energy density in the rest-frame, ρ , and isotropic rest-frame pressure, p , (a perfect fluid), it is possible to write

$$T^{\mu\nu} = \left(\rho + \frac{p}{c^2} \right) U^\mu U^\nu + p g^{\mu\nu}. \quad (2.2)$$

Here U^ν is a velocity four-vector and $g^{\mu\nu}$ is the inverse metric tensor. Using the conservation of the energy-momentum tensor with respect to the covariant derivative, that gives $\nabla_\mu T^\mu_\nu = 0$, and choosing the rest-frame, it is possible to write

$$\frac{\dot{\rho}}{\rho} = -3(1+w) \frac{\dot{R}}{R} = -3(1+w) \frac{\dot{a}}{a}. \quad (2.3)$$

Here, a dot above x is $\dot{x} = \frac{dx}{dt}$. We have assumed that p and ρ are related by an equation of state w , such that $p = w\rho c^2$. We have allowed $R(t)$ to be expressed as $R(t) = R_0 a(t)$ where $a(t)$ is dimensionless

¹This is believed to be true on large enough length scales, from observations of the CMB (Aghanim et al., 2020; Fixsen et al., 1996) and galaxy surveys (Cole et al., 2005; Scrimgeour et al., 2012).

² r is defined by the circumference of a circle, centered at the origin, such that the object lies on the circumference, then defined divided by 2π . This is not in general the same as the comoving distance that a photon will travel to reach the object relative to an origin in a curved Universe. The comoving distance will become apparent when we make a change of coordinates in section 2.2.

and defined such that $a(t_0) = 1$ where t_0 is the present day and R_0 is a constant. For a constant w one can find that for the perfect fluid $\rho \propto a^{-3(w+1)}$. For pressureless matter we expect $w = 0$, indicating the density of matter drops proportionally to a^{-3} (proportional to volume). For radiation that should drop in pressure proportionally to both volume and an additional dimension of distance due to the wavelength of the light stretching in an expanding spacetime, it can be expected that $\rho \propto a^{-4}$ so $w = \frac{1}{3}$. One can then take the *Einstein equation* from General Relativity that governs how spacetime is impacted by the energy-momentum tensor, thus relating $R_{\mu\nu}$ to $T_{\mu\nu}$,³

$$R_{\mu\nu} = \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right). \quad (2.4)$$

Substituting the trace $T = T_{\mu}^{\mu} = -\rho c^2 + 3p$, $T_{00} = \rho c^4$ and $R_{00} = -3\frac{\ddot{a}}{a}$ (the expression for $R_{\mu\nu}$ involves computing the Christoffel symbols for the $\mu, \nu = 0, 0$ terms – and applying a similar process for the $\mu, \nu = i, j$) it is possible to write the two *Friedmann equations* that determine the expansion dynamics of the Universe,⁴

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3} \left(\rho + \frac{3p}{c^2} \right). \quad (2.5)$$

$$\left(\frac{\dot{a}}{a} \right)^2 = \frac{8\pi G}{3} \rho - \frac{kc^2}{a^2 R_0^2}, \quad (2.6)$$

It is possible to define a number of new variables and substitute them into the second Friedmann equation. Firstly we define the Hubble parameter $H(t) = \frac{\dot{a}}{a}$. Then one can define a dimensionless energy density parameter $\Omega = \frac{8\pi G}{3H^2} \rho = \frac{\rho}{\rho_{\text{crit}}}$. ρ_{crit} is called the *critical* density, since in a matter dominated Universe this corresponds to the density that halts expansion or collapse. Making these substitutions, and assuming there is density due to matter, radiation and potentially other sources (and thus summing over these with the index i) gives

$$H^2(t) = \sum_i \Omega_i H(t)^2 - \frac{k^2 c^2}{a^2 R_0^2}. \quad (2.7)$$

Defining a new parameter that we will call the curvature density as $\Omega_k = \frac{-k^2 c^2}{H^2 a^2 R_0^2}$ gives

$$H^2(t) = H^2(t) \sum_i \Omega_i = H_0^2 \sum_i \Omega_{0,i} a^{-3(w_i+1)}, \quad (2.8)$$

where now Ω_k is included in the summation over i and we additionally take into account the equation of state w_i that varies for different sources. We have also defined $\Omega_{0,i}$ to be the present day dimensionless energy density; this involved multiplying each term by a factor of $\left(\frac{H_0}{H(t)} \right)^2$ since the factor of H_0 (the present day universal Hubble constant) has been absorbed into the definition of the present-day critical density. Equation 2.8 is also a differential equation, that can be rearranged for a relationship between time and scale factor as

$$t = \int \frac{da}{H(a)a} = \frac{1}{H_0} \int \frac{da}{\sqrt{(a^2 \sum_i \Omega_{0,i} a^{-3(w_i+1)})}}. \quad (2.9)$$

³ c is the speed of light, G is the universal gravitational constant.

⁴In general, $\ddot{x} = \frac{d^2x}{dt^2}$. Also, to obtain the second Friedmann equation, it helps to make a substitution with the first.

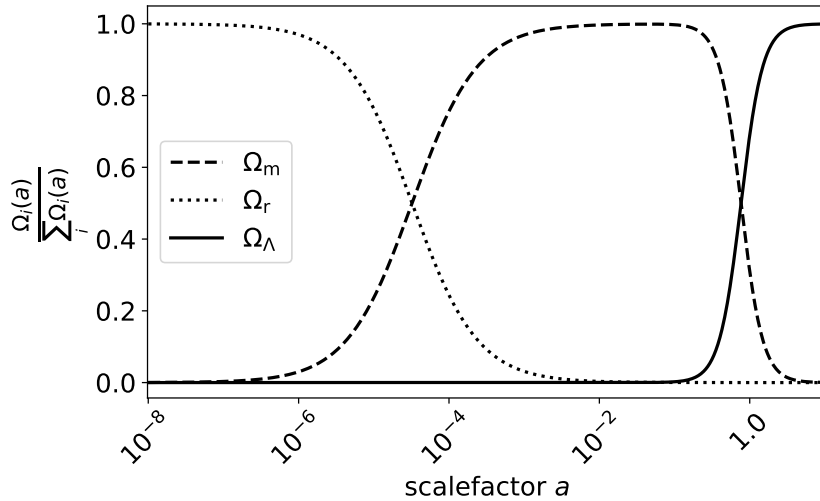


Figure 2.1: The relative contributions of Ω_m , Ω_r and Ω_Λ to the total energy density and thus expansion history as a function of the relative size of the Universe given by a , if their present-day densities are 0.31, $\sim 10^{-5}$ and 0.69 respectively.

The combination of equation 2.8 and the expression for the FLRW metric, equation 2.1 are also extremely useful for determining distances in cosmology and allowing us to compare observations to a cosmological model with different values for $\{\Omega_i, w_i, H_0\}$. Mathematically, we can also allow for additional energy density components in equation 2.8, such as *dark energy*, to be parameterized as $\Omega_{0,\Lambda}$. If this component is constant with time, which is represented by setting $w_\Lambda = -1$, this is called a cosmological constant. Dark energy will be revisited later in this chapter, as it is currently believed that a dark energy component of energy density is required to explain cosmological observations, hence the Λ in Λ CDM. Figure 2.1 shows how different components dominate in their contribution to the expansion rate given by equation 2.8 over time.

A period of radiation domination

Since the Big Bang and a period of rapid expansion called *inflation*, there has been three main phases in the Universe's expansion history. The first was a period in which the energy density of radiation was very high, and the pressure due to radiation was the primary driver of changes to the expansion of the Universe. At this time the baryonic matter constituted a hot dense plasma. The high temperature and pressure allowed for baryons and photons to stay coupled and this largely suppressed the gravitational collapse of matter. Later, sound waves were also able to propagate through the plasma.

A period of Matter domination

Eventually (~ 50000 years after the Big Bang) the energy density of matter (both baryonic and dark matter) began to dominate the expansion of the Universe. The sound waves mentioned above, which are imprinted upon the current-day matter distribution, are referred to as baryonic acoustic oscillations (BAOs) and were able to propagate during this time. Later, these sound waves left their imprint in the *cosmic microwave background* at the epoch called *recombination*, and in the matter distribution

when they ceased. As the Universe continued to expand and cool, so did the hot thermal plasma. Photons that were initially strongly coupled to the hot plasma decoupled and began to propagate freely. Consequently the ionised baryons eventually cooled enough for *recombination* to occur ($\sim 400,000$ years after the Big Bang); this refers to the event where protons could combine with electrons to form hydrogen and after which the BAOs could then eventually cease. The remnant photons from recombination are now what we measure as the cosmic microwave background (CMB, Penzias and Wilson, 1965).

In this epoch baryons could also begin collapsing gravitationally (by this point, dark matter, which does not interact with photons, had already begun collapsing) and started to form the first stars and galaxies. Due to the BAOs which froze in place after *recombination*, matter began preferentially collapsing on scales set by the largest distance sound waves had travelled in the age of the Universe at this time. The collapse of baryons, were dragged slightly by the dark matter which had already been able to collapse prior to recombination, but this process led to an excess clustering of matter on a scale called the BAO scale, and is a visible signature in the distribution of matter in the Universe today. We can probe this through the *matter correlation function* which we observe from the galaxy distribution, or equivalently its Fourier transform, the *matter power spectrum* in which signatures of the BAOs also have been detected (Cole et al., 2005; Eisenstein et al., 2005; Gil-Marín et al., 2016; Abbott et al., 2019). The correlation function represents the excess probability of finding matter overdensities separated by particular length scales from that of a random distribution of galaxies.

A period of Dark Energy domination

In more recent times (a few billion years ago) the density of *dark energy* began to dominate the expansion of the Universe. The discovery of Dark energy was confirmed by Type Ia SNe when astronomers observed that the Universe began to accelerate in its expansion in these more recent times (Riess et al., 1998; Perlmutter et al., 1999); however, tension in measurements of the total matter density from CMB measurements and number counts of galaxies hinted at this prior to 1998 (Ostriker and Steinhardt, 1995; Efstathiou et al., 1990; Yoshii and Peterson, 1995).

Overall, the rate of expansion during these different eras (radiation, matter then dark energy domination) over the age of the Universe has depended on the amount of these components present. The history of the expansion and time between these epochs has affected the suppression and growth of matter overdensities on different length scales and thus the formation of stars and galaxies. Gravity also plays a role in all these epochs. Thus we are able to test models of gravity and cosmology via the distribution and motions of galaxies.

2.2 Distances in cosmology

Using the FLRW metric equation 2.1 one can write an expression for the distance to a galaxy at a moment in time. However we will first rewrite this metric in a new parameterization that is more convenient for determining distances in cosmology. Allowing for a new quantity $d\chi^2 = \frac{dr^2}{1-kr^2}$, we can

write

$$ds^2 = -c^2 dt^2 + a^2(t)R_0^2 (d\chi^2 + S_k^2(\chi)d\Omega^2). \quad (2.10)$$

Solving the expression $d\chi^2 = \frac{dr^2}{1-kr^2}$ for r gives

$$r = S_k(\chi) = \begin{cases} \sin(\chi), & k = 1 \\ \chi, & k = 0 \\ \sinh \chi, & k = -1. \end{cases} \quad (2.11)$$

For light we expect $ds^2 = 0$ (a null geodesic), and if we pick a radial trajectory we can set $d\Omega^2 = 0$. This gives the following expression for the dimensionful radial comoving distance $d(a) = R_0\chi(a)$ to an object that emitted light at t_i reaching us today at $a(t_0) = 1$,

$$d(a) = R_0\chi(a) = \int \frac{cdt}{a(t)} = c \int_{a(t_i)}^1 \frac{da}{a^2 H(a)}. \quad (2.12)$$

The cosmological redshift of an object z can be related to the scale factor corresponding to the time the light was emitted as $a(t) = \frac{1}{(1+z)}$, allowing this to be more conveniently expressed as⁵

$$d(z) = R_0\chi(z) = c \int_0^z \frac{dz}{H(z)} = \frac{c}{H_0} \int_0^z \frac{dz}{\sqrt{(\sum_i \Omega_{0,i} z^{3(w_i+1)})}}. \quad (2.13)$$




Figure 2.2: A plot of comoving distance, $d_C(z)$, luminosity distance, $d_L(z)$ and angular distance, $d_A(z)$, as a function of redshift z . The calculations correspond to the same cosmological model as is shown in Figure 2.1. The distances are shown in dimensionless units by multiplying by $\frac{H_0}{c}$.

The proper distance at time t can be computed by simply multiplying by $a(t)$. This expression above allows us to compare the distances to objects with their measured redshifts from spectra to test cosmological models. However we need to additionally account for impacts that alter the apparent distances we measure. If we use the angular size of an object in the sky to determine the distance to it,

⁵Often the radiation density is neglected in the $H(z)$ term in the denominator of this equation as it has little impact on the distance to objects (and the expansion rate) in the late Universe. However, for distances or calculations of $H(z)$ in the early Universe when radiation dominated, it must be included for accurate calculations.

there is an issue that the apparent angular size we measure is altered by the expansion of space; an object may look larger than it really is since the object was closer to us when the light was emitted. Likewise when using the measured flux from an object via a comparison with its intrinsic luminosity to determine the distance, the number of photons arriving in a detector per unit time will decrease due to expanding space making the object appear further away. Additionally there are curvature effects. As such, corrections need to be made; the angular diameter distance we measure is given theoretically as $d_A(z) = d(z)/(1+z)$ and for the luminosity distance it is $d_L(z) = d(z)(1+z)$, where now $d(z) = R_0 S_k(\chi)$. To make sense of this, in general a distance determined from angular size is defined using the radius r of a circle centered at the origin of the coordinate system and likewise a distance determined using flux and luminosity is defined by relating the total power from an object divided by the area of a sphere with radius r . As such, it is clear we should use $r = S_k(\chi)$ generally to obtain the correct answer when the Universe has curvature. The factor of $(1+z)$ for luminosity distance accounts for the decreased number of photons and energy of each photon per unit time in expanding space, while for angular diameter distance the factor $1/(1+z)$ accounts for the impact an expanding space has on the angular size seen in the sky. The difference between various definitions of distance discussed are shown in Figure 2.2.

2.3 Cosmological probes

Having covered the essential mathematics, we can now move on to discussing measurements of the Universal expansion rate and distances to objects from cosmological data that has led to the Λ CDM model.

2.3.1 Supernovae

Type Ia Supernovae (Type Ia SNe) are an explosion of a white dwarf star.⁶ Due to the physics that leads to these explosions, Type Ia SNe have a standardisable intrinsic luminosity, and thus are referred to as *standard candles* (Phillips, 1993). A measurement of the flux (or apparent magnitude m) of a SN, combined with information about their intrinsic luminosity (or absolute magnitude M) gives us a way to measure the *luminosity distance* $d_L(z)$ to the SN. If the redshift of the SN host galaxy can also be measured from spectroscopy (or photometry, these are explained in more detail in section 2.3.3) we can test cosmological models. We can write the distance modulus $\mu(z)$ as,⁷

$$\mu(z) = m - M = 5 \log_{10}(d_L(z)) + 25. \quad (2.14)$$

Since the luminosity distance is dependent on cosmological parameters this gives us a way to explore cosmological models that give the best fit to the data. In 1998 and 1999 two teams of cosmologists

⁶White dwarfs are remnant of stars with masses similar to our sun, that have undergone various nuclear fusion reactions before expelling their core leaving electron-degenerate matter; that is matter that has a temperature dictated by the Pauli Exclusion principle. When these stars accrete sufficient mass in a binary system it leads to the SN explosion.

⁷This equation holds given the distance is entered in megaparsecs.

(Riess et al., 1998; Perlmutter et al., 1999) used Type Ia SNe to measure the expansion rate of the Universe, allowing them to determine constraints primarily on $\Omega_m \propto a^{-3}$ the matter density in the Universe and $\Omega_k \propto a^{-2}$, the curvature density parameter. However, the results found that an additional parameter, a non-zero *cosmological constant*, was required in order for the data to give a good fit to the model. Very approximately these teams found that for a good fit to the data, it is required that $\Omega_m \sim 0.3$, the cosmological constant $\Omega_\Lambda \sim 0.7$ and consequently the curvature parameter defined such that $\Omega_{0,k} = 1 - \sum_j \Omega_{0,j}$ (where summing over j does not include $\Omega_{0,k}$) gives $\Omega_k \sim 0$. This result confirmed the existence of dark energy (although it was mentioned earlier there were some prior hints for it); the cosmological constant drives the Universe to accelerate at late times rather than eventually slow down, reach a constant expansion rate or collapse. These observations are still supported by SNe Ia data measured over twenty years later (at the time of writing); for example, by the Dark Energy Survey Collaboration (DES Abbott et al., 2024); these results also now give hints of a time-varying dark energy, which is discussed in section 2.4.

2.3.2 The Cosmic Microwave Background

The cosmic microwave background (CMB), first detected in 1965 (Penzias and Wilson, 1965), has a wealth of cosmological information. It refers to the radiation that initially was coupled to a hot dense plasma of baryonic matter in the early Universe, before decoupling and freely propagating after the Universe became both cool enough and reached a low enough density. As mentioned before, this epoch is known as *recombination* because electrons combined with protons to form hydrogen atoms since the photons were no longer energetic enough to ionise them. Today the photons are detected as extremely uniform microwave radiation corresponding to a temperature of $T \sim 2.725$ K that can be detected from all directions (Mather et al., 1994; Fixsen et al., 1996). The CMB is additionally a prediction of an expanding Universe that started initially at $a(t=0) = 0$ (with a Big Bang).

The highly uniform microwave radiation has tiny temperature fluctuations across the sky that correspond to about $\Delta T \sim 10^{-5}$ K. These temperature fluctuations exist primarily due to two effects; firstly, the photons that have reached us today climbed out of gravitational potential wells across space, that is a consequence of the matter distribution at the moment the photons decoupled. This is referred to as the Sachs-Wolfe effect. The matter distribution across space is a consequence of a number of variables including how quickly the Universe expanded, how much radiation and matter (including the different contributions to baryonic matter and dark matter) were present and the speed of sound up until this epoch. This effect dominates the anisotropies on *super-horizon* scales - length scales larger than the causal horizon scale at the epoch of recombination z_* , $r_s(z_*)$, which refers to the distance sound waves can travel in the age of the Universe at z_* .

Secondly, on characteristic length-scales (sub-horizon scales) photons were hotter due to the over-density peaks of sound waves that propagated through the plasma prior to recombination. Sound waves were able to propagate because of the high density and pressure of the plasma; these sound waves are the baryon acoustic oscillations (BAOs) mentioned previously. When the sound waves effectively froze

in place, the over-densities corresponding to compressions of matter and under-densities corresponding to rarefactions sourced some of the largest temperature fluctuations that can be seen in the CMB. The presence of these matter fluctuations has also led to the excess formation of galaxies on scales corresponding to the furthest possible physical distance sound waves were able to travel in the age of the Universe at this time - the sound horizon $r_s(z_d)$ (z_d refers to the *baryon-drag* epoch - this is a little after recombination since some high energy photons remained coupled to the baryons for a little while longer, allowing sound waves to continue). The sound horizon is impacted by the sound speed at the time $c_s = \sqrt{\frac{dP}{d\rho}}$ which depends on the equation of state w_i and density ρ_i for Ω_b (baryons, the matter density with the cold dark matter density subtracted off) and Ω_γ (photons) that made up the hot plasma. The expansion rate H_0 , the matter density Ω_m and dark energy density Ω_Λ also impacted the physical scale on which we see the temperature fluctuations today. The relative amplitude of temperature fluctuations on different length scales, also depends on the ratio of matter and radiation in the Universe (Ω_m and Ω_r).

The fluctuations in the CMB temperature map have been measured most precisely by the Planck Collaboration in 2018 (Aghanim et al., 2020). The analysis of the CMB data prefers a flat Universe corresponding to $\Omega_k \sim 0$, which suggests $\Omega_m + \Omega_\Lambda \approx 1$. In combination with constraints from other probes that suggest $\Omega_m \sim 0.3$, this is strong evidence for the presence of dark energy. The CMB also provides strong evidence for the presence of *dark matter* (alongside evidence from galaxy rotation curves and clusters of galaxies) from measurements of the higher-order temperature fluctuations; if all of the matter density was baryonic, the CMB fluctuations would be drastically different since dark matter does not interact with photons (and thus would impact the propagation of sound waves very differently prior to recombination). Such strong evidence for the presence of matter that does not interact like matter in the Standard Model or at all with photons has led to the CDM (cold dark matter) included in ‘ Λ CDM’. It is called ‘cold’ meaning that it is likely made up of heavy, slow-moving particles.

2.3.3 Galaxy surveys

Galaxy surveys, in which we simply use galaxies to trace how matter is distributed through space, can be useful for testing cosmology. This falls under the umbrella of *large-scale structure cosmology*; we will address this much more in the next chapter, but a brief summary will be provided here. In galaxy surveys three-dimensional maps of the galaxy distribution are created from the angular position of galaxies in the sky in addition to a measurement of their cosmological redshift. The redshift may come from either spectroscopy, or multi-bin photometry. In spectroscopy, the wavelengths of light from the galaxy are separated into fine bins using a spectrograph which involves a diffraction grating. Alternatively, photometry refers to making measurements of the galaxies brightness in various filters, which allows one to estimate the redshift, but with a higher level of uncertainty; it is effectively spectroscopy with very coarse wavelength binning. The redshift gives information on the distance to the galaxy from a cosmological model via equation 2.13.




Figure 2.3: The matter power spectrum (left) and correlation function (right) which are related by a Fourier transform and are closely related to the galaxy power spectrum and galaxy correlation function. These have been generated using the CLASS Boltzmann solver (Blas et al., 2011), with $\Omega_b = 0.048$, $\Omega_m = 0.31$, $H_0 = 67 \text{ km s}^{-1} \text{ Mpc}^{-1}$, and $\Omega_\Lambda = 0.69$.

There are various ways in which one can analyse this data for testing cosmology. Typically, one constructs a correlation function of the separation between galaxies as a function of scale. This function allows one to see the excess probability of finding galaxies as a function of their separation. Since the galaxies trace matter in the Universe, this tells us something about the underlying matter distribution in space. One can also choose to look at the data in Fourier space, in which case the data product is referred to as the *power spectrum*. In either case, one can expect to see distinct features that encode information about the expansion history. The power spectrum has a large ‘turn-over’ which can be seen in Figure 2.3. The length scale corresponding to the position of this peak represents the epoch at which matter began to dominate the expansion rate of the Universe, after the period in which radiation dominated the expansion; the epoch is called matter-radiation equality, and depends on the density for each component. The oscillatory feature (called the wiggles) has a wavelength given by the sound horizon scale at the baryon-drag epoch $r_s(z_d)$; this is the point at which the last few high energy photons became completely decoupled from the plasma in the BAOs. In real space, there is a peak in the correlation function that corresponds to the oscillatory wiggles. This feature was detected for the first time by Cole et al. (2005); Eisenstein et al. (2005) but has recently been measured with the best precision by the Dark Energy Spectroscopic Instrument (DESI) collaboration (Adame et al., 2025b).

To test cosmology, historically many have used just the wiggles (peak) from the power spectrum (correlation function) as a *standard ruler*; because the feature corresponds to a characteristic scale, it represents a consistent length that can be measured at different epochs by identifying the feature at different redshifts. As such, it allows us to measure the expansion rate and similarly to the case with Type Ia SNe this contains information about Ω_m and Ω_Λ , but also $r_s(z_d)H_0$; with additional information from other sources, it is possible to individually constrain $r_s(z_d)$ and H_0 .

With increasingly precise data it is possible to extract more cosmological information than just

what is available in the BAO feature; one can compare models for the full shape of the matter power spectrum (or instead the correlation function) to fit for cosmological parameters, extending to those such as n_s , the spectral index which controls the ‘tilt’ of the spectrum and A_s , the power spectrum amplitude (these parameters can both also be fit with CMB data). These parameters are related to the spectrum of primordial fluctuations - the quantum fluctuations that grew during a period called inflation, in which it is believed the Universe expanding by many e-folds very rapidly prior to the ‘hot’ Big Bang. These fluctuations sourced the initial over and under-densities of matter in the Universe by growing in size with inflation. It is also possible to extract information on various other parameters beyond the ‘vanilla’ Λ CDM model which will be discussed later.

2.3.4 Peculiar velocities

Previously we mentioned galaxies recede from each other due to the expanding space which results in a cosmological redshift in the spectrum of each galaxy. However, galaxies are also subject to gravitational interactions and their redshifts are thus impacted by the Doppler effect; the component of a galaxy’s motion due to these interactions are called peculiar motions, or peculiar velocities. Like a galaxy redshift survey, a galaxy peculiar velocity survey traces information about the underlying matter distribution of galaxies and thus provides complementary information.

There are various ways in which peculiar velocities can be used. One approach is to compare a measured galaxy velocity field from a survey to a galaxy density field; these are closely related fields but the observed galaxy density field is impacted by the effect of peculiar velocities on the redshifts of galaxies, called *redshift-space distortions*. This will be covered in much greater detail in the next chapter. The peculiar velocity field $u(x)$ is related to the matter density field $\delta(x)$ in linear theory as

$$u(x, a) = \frac{f(a)}{4\pi} aH(a) \int d^3\mathbf{y} \delta(\mathbf{y}, a) \frac{\mathbf{y} - \mathbf{x}}{|\mathbf{y} - \mathbf{x}|^3}. \quad (2.15)$$

By integrating over the matter density field and comparing to the velocity field one can measure the growth rate of structure $f(a) \approx \Omega_m^{0.55}$. The growth rate of structure represents physically the speed at which galaxies are falling in to each other, but mathematically is defined as $f(a) = \frac{d \ln D(a)}{d \ln(a)}$ where $D(a)$ is the growth factor and represents the relative size of a density perturbation at a time corresponding to a . Additionally, alternative approaches can be taken to measure the growth rate by combining information from galaxy and peculiar velocity surveys (Howlett et al., 2017a) and it has also been shown peculiar velocities can be used to improve constraints on other cosmological parameters or in combination with other probes in other works (Burkey and Taylor, 2004; Koda et al., 2014; Castro et al., 2016; Howlett et al., 2017b; Agrawal et al., 2019; Amendola and Quartin, 2021; Whitford et al., 2022; Quartin et al., 2022; Giani et al., 2023). Lastly, due to the nature of peculiar velocity measurements as having, individually, very large uncertainties, it is common to average them and measure the *bulk flow*. A bulk flow measurement gives the average motion of peculiar velocities of galaxies in a small volume of space. One can compare the amplitude of this measurement to a prediction from linear theory since the amplitude of the motion depends on the local matter distribution; this will be discussed more in Chapter 3 and in great detail in Chapter 4.

2.3.5 Other probes

I will briefly mention here other methods of testing cosmology that are relevant in later sections and are often used to provide complementary information or even independent constraints on cosmological parameters; however this is not a comprehensive overview.

Big Bang Nucleosynthesis (BBN)

Big Bang Nucleosynthesis refers to a period of time in the early Universe, in which various chemical processes took place resulting in the abundance of hydrogen, helium and traces of other light chemical elements that makes up the Universe. Observations of the abundance of these elements we see today allow us to constrain the energy density in the Universe due to baryons, Ω_b . This is because the greater the baryon density, the greater the ratio of baryons to photons, which impacts the abundance of light elements that form. Additionally, BBN can help constrain the *effective number of neutrino species* N_{eff} ; this parameter will be revisited briefly at the end of this chapter and in greater detail in Chapter 5.

Lyman- α forest

The Lyman- α forest refers to the Lyman- α absorption lines seen in the spectra of distant galaxies (such as quasars). The absorption lines occur at various different wavelengths in the spectra, as a result of neutral hydrogen along the line-of-sight between the observer and galaxies absorbing the radiation. As such this helps to probe the matter distribution in the Universe. It is also possible to measure the BAO signal seen as excess Lyman- α absorption in the spectra of galaxies as a function of their separation.

Weak lensing

Weak gravitational lensing refers to the distortions to the elliptical shapes of galaxies we see projected on the sky due to the lensing of their light from intervening matter. The lensing is sourced by objects in the foreground whose gravitational field alters the path light travels, which is a prediction of General Relativity. Galaxy surveys are able to probe this weak lensing signal and obtain constraints on Ω_m since the lensing probes the underlying matter density field in addition to the amplitude of matter fluctuations on different scales.

2.4 The status of modern cosmology

The Λ CDM model

Throughout this chapter I have introduced the following parameters; H_0 , Ω_m , Ω_Λ , n_s , A_s . Additionally, τ_{re} is a cosmological parameter that can be measured using the CMB photons. It represents the opacity of electrons to CMB photons and is called the *optical depth to reionization*; free electrons - which were present after the first stars in the Universe produced photons that ionized atoms in the galactic and intergalactic medium - can scatter the CMB photons, impacting our measurements of the CMB.

These six parameters make up the vanilla Λ CDM model. In Λ CDM it is also generally assumed that $\Omega_k = 0$. Observations from the CMB, BAOs, Type Ia SNe and BBN have provided strong evidence for this model. Despite its success, there is much for cosmologists to try to learn about the nature of dark matter and dark energy, which are poorly understood. Currently, various particle candidates for dark matter exist, but no evidence for these has been confirmed experimentally. Dark matter has only been observed indirectly through the gravitational impact of dark matter on star, galaxies and the paths of photons. Likewise, our observational evidence for dark energy cannot be explained by quantum mechanics or the Standard Model.

Cosmological tensions

Despite the success of Λ CDM, it suffers from tensions in fits to different cosmological datasets. The most well-known of these is the H_0 tension, in which datasets that probe ‘late-Universe’ physics such as Type Ia SNe, prefer larger values of $H_0 \sim 73 \text{kms}^{-1} \text{Mpc}^{-1}$, and datasets that probe ‘early-Universe’ physics, such as the CMB, systematically prefer lower values of $H_0 \sim 67 \text{kms}^{-1} \text{Mpc}^{-1}$ (Di Valentino et al., 2021). Another well-known tension is the σ_8 tension; σ_8 refers to the variance of matter fluctuations over spheres of radius $8 \text{Mpc} h^{-1}$ (where the radius $\sim 8 \text{Mpc} h^{-1}$ has been chosen mainly for historical reasons).⁸ Similarly, there is a systematic difference in the preferred value for σ_8 from probes of late and early-Universe physics (Abdalla et al., 2022). In this thesis I will also discuss the bulk flow tension, which refers to the less severe but nonetheless persistent preference for bulk flows with amplitudes much larger than one might expect to see in the Λ CDM model (Kashlinsky et al., 2008; Watkins et al., 2009; Feldman et al., 2010; Watkins and Feldman, 2015b; Peery et al., 2018; Howlett et al., 2022; Watkins et al., 2023; Whitford et al., 2023); however others have found better agreement with Λ CDM, although more for measurements that probe smaller scales (Nusser and Davis, 2011; Hong et al., 2014; Ma and Pan, 2014; Hoffman et al., 2015; Scrimgeour et al., 2016; Qin et al., 2018, 2021; Hoffman et al., 2024). These measurements are summarised in Figure 2.4. The measured amplitude is generally sufficiently large that uncertainty in Ω_m or σ_8 cannot reconcile the measurements with the Λ CDM model. This will be revisited in detail in Chapter 4.

2.4.1 Beyond Λ CDM

Various extensions to the Λ CDM model can be explored to resolve tensions or simply explore new physics. Here I will discuss some of the most currently interesting and relevant extensions to this thesis.

Time-varying Dark Energy

One extension to the vanilla Λ CDM model is to allow for the dark energy parameter Ω_Λ to have an equation of state with freedom so that w is not necessarily fixed to $w = -1$. One can allow w to vary

⁸It should be noted that the σ_8 tension appears to have been reduced in the latest results by the Kilo-Degree Survey (Wright et al., 2025).

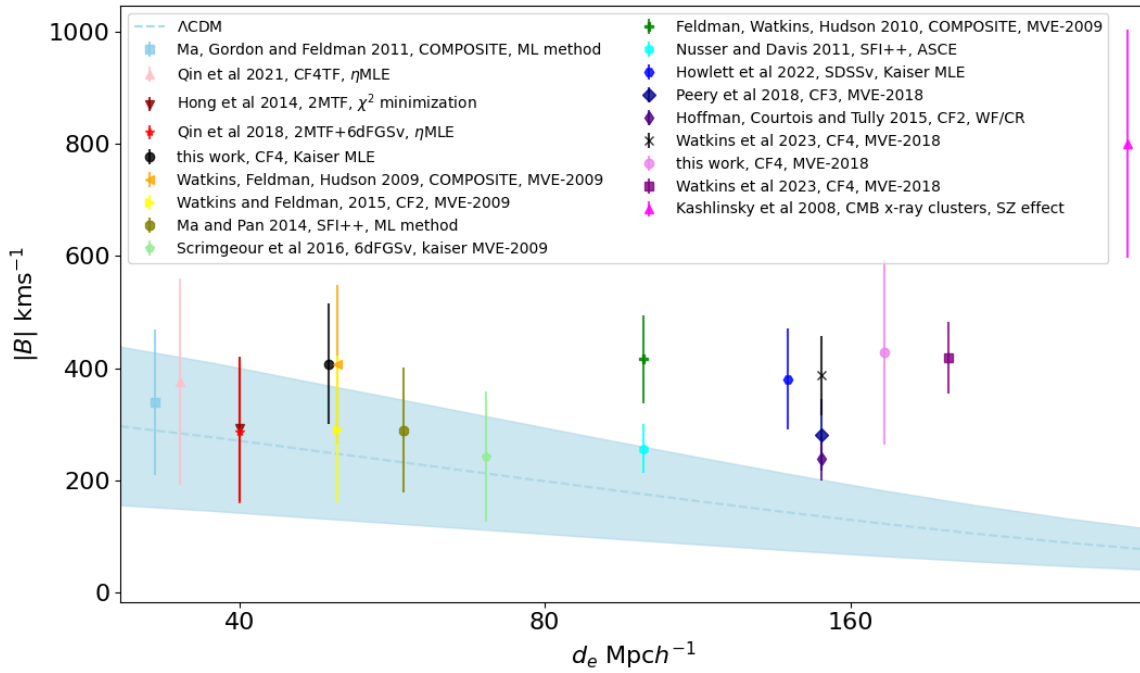


Figure 2.4: Measured bulk flow amplitudes from various authors shown in the legend, where ‘this work’ refers to my measurements made in Whitford et al. (2023) (this figure comes from the conference proceedings Whitford, 2024). The measurements are plotted at varying effective depth d_e , which is essentially the scale probed by the bulk flow measurement (in comoving coordinates), which is a function of the survey geometry. The dashed curve shows the predicted bulk flow amplitude and 1σ statistical uncertainty for the Λ CDM model with $H_0 = 67.74 \text{ km s}^{-1}$, $\Omega_m = 0.31$, for a spherical survey geometry. It should be noted that the surveys used to measure the bulk flows on this plot do not necessarily have a spherical geometry; nonetheless the comparison is useful since we expect the theoretical curve to decrease with the effective depth regardless of geometry.

as a free parameter, which allows for the freedom to have dark energy with a time-varying energy density. It is also possible to go a step further and allow for the equation of state w itself to vary with scalefactor, such as $w = w_0 + w_a(1 - a)$. Until recently, there was not a strong statistical preference for these models over vanilla Λ CDM. However, the Dark Energy Survey (DES) Year 5 SNe Ia results found a moderate preference for models with time-varying dark energy, when combining their data with various datasets (Abbott et al., 2024). Additionally, the DESI DR1 and DR2 BAO results similarly found a preference for the same model (Adame et al., 2025b; Karim et al., 2025), as well as results from the UNION3 supernova compilation and analysis (Rubin et al., 2025). In Figure 2.5, I show a plot of some of the distance measurements obtained from the DESI DR2 BAO analysis as a function of z , that supports the preference for time-varying dark energy.

Neutrino mass constraints

Another interesting extension to Λ CDM is to consider the impact of massive neutrinos on cosmological data; terrestrial particle-physics experiments have not been able to constrain the mass of each kind of neutrino species, which are at least $\sim 10^{-3}$ times smaller than the mass of an electron. However, cosmological probes have constraining power on the sum of neutrino masses, $\sum m_\nu$, which involves summing over the mass of each neutrino specie (or equivalently mass eigenstate). The presence of

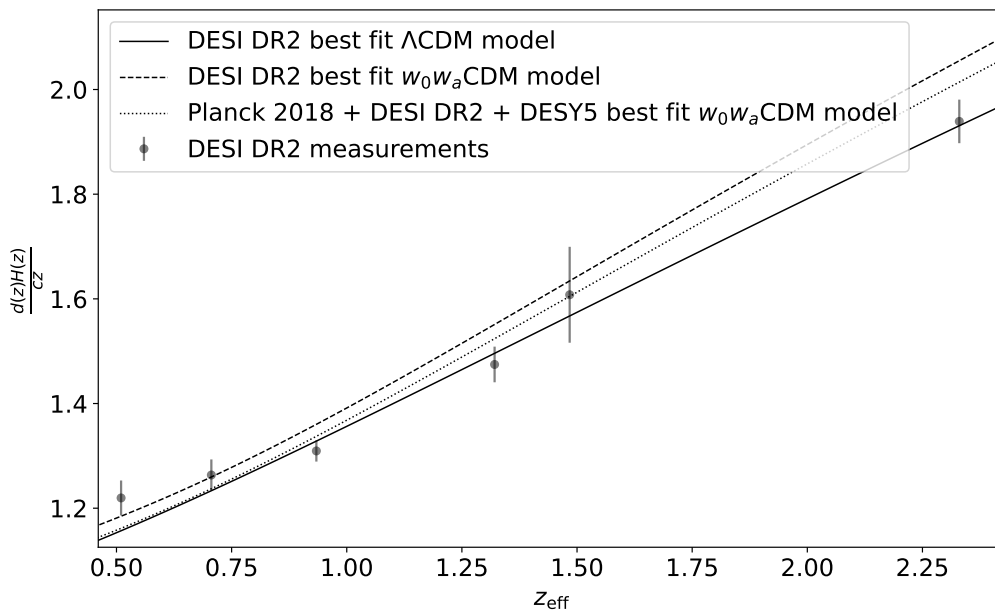


Figure 2.5: Measurements of the ratio $\frac{d(z)H(z)}{cz}$ obtained from DESI DR2 BAO measurements, plotted as a function of the effective redshift z_{eff} , which is essentially a weighted estimate of the redshift for the measurement given the galaxy survey distribution, and uncertainties on the measurements. I have taken the values from Table IV in Karim et al. (2025). For comparison I show lines for the best fit Λ CDM model and best fit w_0w_a CDM model from this data, and the best fit when one includes information from Planck CMB measurements and the DES Year 5 SNe Ia results (best fit cosmological parameters have been taken from Table 5).

massive neutrinos, or particles that freely propagate after ceasing interactions with other Standard-Model particles in the early Universe, impacts the growth of structure (Lesgourgues and Pastor, 2006, 2012). Essentially, the growth of structure on scales below the distance they are able to propagate freely before their velocity's decay is suppressed, proportionally to their contribution to the total matter density in the Universe. The suppression of structure growth is specifically an impact of neutrinos that have non-zero masses that have large thermal velocities at early times. Their presence as a 'radiation-like' particle at early times also impacts the expansion rate. Both of these effects are detectable in the full shape matter power spectrum, and in the CMB photons. Some of the earliest constraints with large-scale structure data are from the 2-degree Field Galaxy Survey (2dFGS) (Elgarøy et al., 2002) and the Sloan Digital Sky Survey (SDSS) combined with CMB data (Tegmark et al., 2004). Interestingly, until recently the strongest constraints on $\sum m_\nu$ came from the combination of CMB data from Planck (Aghanim et al., 2020) with BAO data, with $\sum m_\nu < 0.12$ eV at the 95% confidence level. In principle, $\sum m_\nu$ must be greater than at least 0.058 eV since neutrino oscillation experiments put a lower bound on $\sum m_\nu$ from measurements in the mass splittings, Δm_{ij}^2 , where i and j donate the differences in mass for the i^{th} and j^{th} neutrino mass eigenstate (Hirata et al., 1989; Fukuda et al., 1998). However, recent data from DESI BAOs (Data Release 1) combined with the Planck CMB data find a preference for masses almost below this value, with $\sum m_\nu < 0.072$ eV at the 95% confidence level (Adame et al., 2025b). Analyses by other authors have shown that there is maybe even a preference for masses below zero (Craig et al., 2024; Green and Meyers, 2025). This result suggests our models may not be sufficient to fit the data, prompting us to investigate further beyond Λ CDM. Results from the

DESI BAO Data Release 2 (Karim et al., 2025) with Planck data similarly find a very strong upper bound on $\sum m_\nu$, although in Elbers et al. (2025) it is found that these constraints are relaxed and a normal neutrino mass hierarchy is supported in a more flexible background cosmological model, such as with time-varying dark energy. I show constraints in a plot taken directly from Elbers et al. (2025) for the ‘effective’ neutrino mass in Figure 2.6a.

The effective number of neutrino species

The effective number of neutrino species, N_{eff} , can be included as an additional free parameter to those of Λ CDM. N_{eff} denotes the effective contribution of neutrinos to the radiation density in the early Universe, in terms of the number of species. It is defined as

$$\Omega_r = \left(1 + N_{\text{eff}} \frac{7}{8} \left(\frac{4}{11} \right)^{4/3} \right) \Omega_\gamma, \quad (2.16)$$

where Ω_γ denotes the contribution of photons to the radiation density. Expanding the terms gives the contribution separately from photons and neutrinos. In the Standard Model, there are just three neutrino species. However, in the Standard Model N_{eff} does not need to be an integer, and is expected to be $N_{\text{eff}} = 3.044$. This is because neutrinos gain some extra temperature in the early Universe which raises their energy; it is well approximated as raising the number of neutrinos and as such is absorbed into the definition of N_{eff} (Riemer-Sørensen et al., 2013). Since N_{eff} impacts the total radiation density alterations to N_{eff} impact the full shape matter power spectrum and the CMB photons as well as BBN. N_{eff} has been measured very precisely by Planck CMB data (Aghanim et al., 2020) in combination with other datasets to be $N_{\text{eff}} = 2.99 \pm 0.17$, and a historical review on constraints can be seen in (Riemer-Sørensen et al., 2013). It is also possible to measure a subtle shift in the phase of the ‘wiggles’ of the BAOs (Baumann et al., 2016, 2019; Whitford et al., 2025) or almost equivalently, in the CMB radiation (Follin et al., 2015), due to the presence of neutrinos in the early Universe. The phase shift arises because the neutrinos create a temporal shift in the gravitational potential of the BAOs as they free stream (Bashinsky and Seljak, 2004). The amplitude of the shift is able to constrain N_{eff} in an almost model-independent manner. This will be explored in Chapter 5 and a phase shift for neutrinos beyond the standard model will be explored in Chapter 6.

Neutrinos with self-interactions

The final extension I will consider in this chapter is for neutrino self-interactions. Unlike $\sum m_\nu$ or N_{eff} , allowing for the freedom to have self-interactions among neutrinos goes beyond predictions of the Standard Model. In a very simple model, one may consider a new mediator ϕ (force-carrier particle) that allows for the self-interactions, with a coupling strength g_ν . In the simplistic model I study in Chapter 6 of this thesis, the interaction strength g_ν is universal (the same between each species). Generally, the effect of the self-interactions on cosmological data will depend on both the mass m_ϕ of the mediator and g_ν . However, in the model studied in this work it is possible to capture the effect of both parameters in an effective coupling constant $G_{\text{eff}} = |g_\nu|^2 / m_\phi^2$ (this model is described by Kreisch

et al., 2020).⁹ While experimental data (Lyu et al., 2021; Blinov et al., 2019) strongly constrains models of universal neutrino self-interactions, the parameter space for neutrino self-interactions has not entirely been ruled out. Cosmological data that probes the early Universe allows for one to explore particles in conditions that involve high pressure and temperature and thus provide unique insight to the behaviour of particles that may not be explored in a laboratory. Given that the impact of universal self-interactions on cosmological probes is able to capture the same (or similar) impacts of more novel interactions, studying these simpler models are still useful and can complement terrestrial experiments. At present, neutrinos are poorly understood and there is significant motivation to explore beyond the Standard model. In cosmological data, the self-interactions leave a peak-like imprint in the matter power spectrum and alter the amplitude of the CMB (Kreisch et al., 2020). They also leave a phase shift in both the CMB and BAO oscillations. I discuss the details of the cosmological impacts more in Chapter 3 and Chapter 6.

There are multiple motivations to study self-interactions. Regarding cosmological data, it is motivated as a potential solution to the Hubble tension by altering the sound horizon scale r_s , which allows h to shift (Berryman et al., 2023). Novel interactions could also potentially provide an explanation for very tight upper bounds on $\sum m_\nu$ through changes to neutrino free-streaming which are discussed by Craig et al. (2024). However, it should be noted the delays to free-streaming that are caused by self-interactions which I study in this thesis would not impact $\sum m_\nu$ constraints because a more significant delay is required than what this model allows. Signatures of self-interactions found in cosmological data also provide complementary bounds to laboratory probes, that look for novel interactions to try explain how neutrinos gain mass or to explain anomalous experimental results (Crivellin and Mellado, 2024).

At present, constraints on G_{eff} from cosmological data including the CMB and matter power spectrum suffer from bimodality; this can be seen in Kreisch et al. (2020, 2024); Camarena et al. (2023a); Camarena and Cyr-Racine (2025); Poudou et al. (2025). I show an example of this, taking a plot with constraints directly from Kreisch et al. (2024), in Figure 2.6b. Usually, the two modes found in the constraints are described as the moderate (MIV) interaction mode and the strong (SIV) interaction mode (the peak for G_{eff}). However, various experiments, BBN measurements and CMB measurements can rule out some of the parameter space for interactions for masses $m_\phi \geq 10^{-1}$ (Berryman et al., 2023). However, part of the parameter space has not yet been ruled out.

Summary

In summary, this chapter has introduced the fundamental theory that underlies modern cosmology, current observations supporting the Λ CDM model, and highlighted some relevant issues at the time of writing. In the next chapter, we will discuss in detail how one can perform tests of cosmology with data that probes the large-scale structure of the Universe.

⁹This involves an approximation that assumes that m_ϕ is much larger than the typical energy of a scattering event, and is valid for the energy scales probed for epochs prior to the recombination.

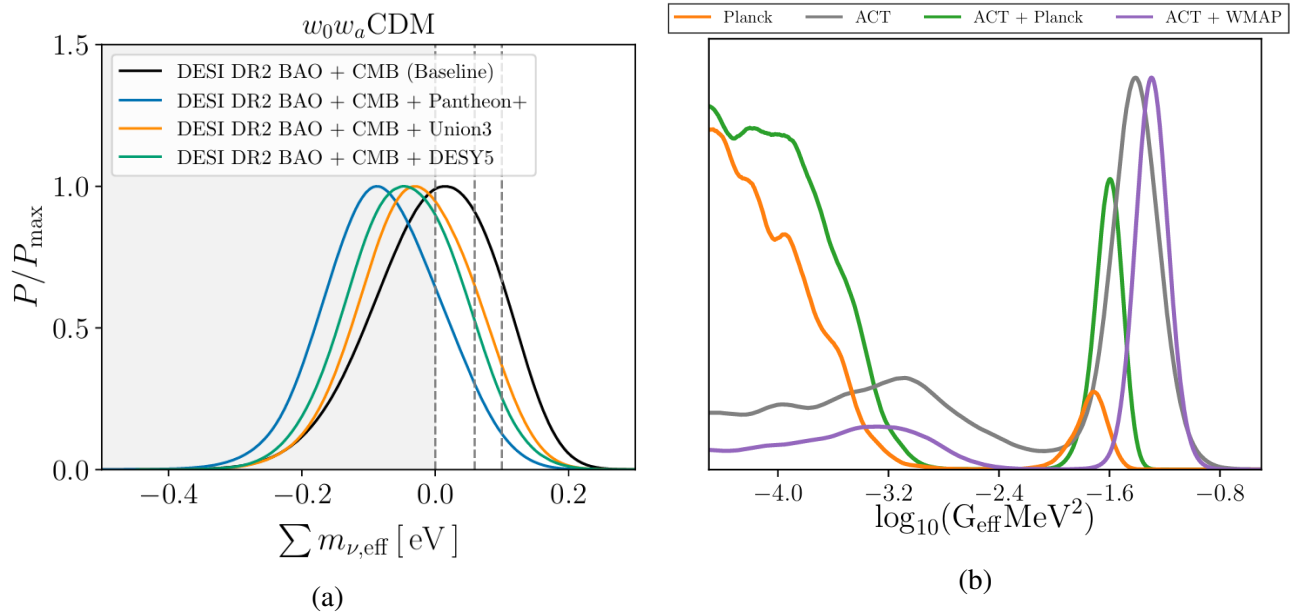


Figure 2.6: Left panel: Constraints on the ‘effective’ sum of neutrino masses, taken directly from Elbers et al. (2025). In this model, the neutrino mass may take negative values, with the purpose being to test what is preferred by the data. The dashed lines show lines corresponding to zero neutrino mass, and also the lower bounds set by oscillation experiments in the two neutrino mass hierarchies. In a $w_0w_a\text{CDM}$ model, it is possible to have the peak of the posterior distribution at positive values although it is noted that the posterior peaks at negative values for combinations of DESI and CMB data with SNe Ia datasets. Right panel: Bimodal constraints on the strength of universal neutrino self-interactions in data from various CMB experiments, taken directly from Kreisch et al. (2024). The datasets correspond to that of Planck 2018, data from the Atacama Cosmology Telescope (ACT), and data from the Wilkinson microwave anisotropy probe (WMAP).

Chapter 3

Theory of large-scale structure cosmology with galaxy and peculiar velocity surveys

In Chapter 2 I introduced galaxy surveys, as a way of tracing the distribution of matter throughout the Universe. Over the last twenty to thirty years, galaxy surveys have grown considerably. While the 2-degree Field Galaxy Redshift survey (2dFGRS¹) conducted in Australia from 1997-2002 mapped over $\sim 230,000$ galaxies, the Dark Energy Spectroscopic Instrument survey (DESI) collected observations for ~ 6.5 million galaxies in the first three years of the survey (Adame et al., 2025b).² Additionally, the CosmicFlows-4 galaxy peculiar velocity catalogue has compiled the largest map of galaxy motions to date with a compilation of $\sim 50,000$ measurements (Tully et al., 2023). This compares to ~ 3500 galaxies in the Mark III catalogue from the 1990s (Willick et al., 1997). As our ability to collect this data grows to larger volumes we are able to conduct more precise tests of cosmological models and map out the Universe with unprecedented detail. In this chapter I will describe the necessary theory for the analysis of data from redshift and peculiar velocity surveys, giving the required background for Chapters 4, 5 and 6. I will also discuss how large-scale structure (LSS) contains information on neutrinos, and has potential to find signatures of particles beyond the Standard Model.

3.1 Galaxy surveys

3.1.1 The matter density field and velocity field

While observations of galaxies are of the visible stars and gas, this traces an underlying field of invisible dark matter throughout the Universe. The matter density field $\delta(x)$, defined as

$$\delta(\mathbf{x}) = \frac{\rho(\mathbf{x}) - \bar{\rho}}{\bar{\rho}} \quad (3.1)$$

where $\bar{\rho}$ is the average matter density, refers to all matter including invisible dark matter in addition to visible baryonic matter. Galaxies trace the underlying field because they are more likely to form in

¹<http://www.2dfgrs.net/>

²<https://www.desi.lbl.gov/category/announcements/>

regions where the matter density is larger. However, the field traced by galaxies, $\delta_g(\mathbf{x})$ is not exactly equivalent to $\delta(\mathbf{x})$. In regions where the underlying dark matter density is not sufficiently dense, galaxies have not formed at all as a critical density is required for baryonic matter to collapse and form stars. In linear theory of the growth of structures in the Universe (that is when we used linearized differential equations to model the growth of structure of a density perturbation), the two fields are related by a galaxy bias factor b_g ,

$$\delta_g(\mathbf{x}) = b_g \delta(\mathbf{x}). \quad (3.2)$$

In reality, this equation is a very simple model, because the galaxy bias may depend not only on time (parameterized by redshift z) but may be scale-dependent, or non-local, meaning the bias on a particular scale may be dependent on other scales, i.e. $b_g(x) \propto \delta(x + \Delta x)$. It is expected that the presence of massive neutrinos introduces a scale dependence to b_g (Vagnozzi et al., 2018), aside from other potential effects due to non-linear structure growth. In linear theory, the density field at a time denoted by z_1 can be related to the field at another time z_2 as,

$$\delta(z_1, \mathbf{x})D(z_1) = \delta(z_2, \mathbf{x})D(z_2) \quad (3.3)$$

where $D(z)$ is the linear growth factor, which gives the amplitude of a density fluctuation, mentioned briefly in Chapter 2. This comes from the mathematical description of the growth of the density field via

$$\delta(z, \mathbf{x}) = D(z)\delta_0(\mathbf{x}). \quad (3.4)$$

$\delta_0(\mathbf{x})$ denotes the field at some time such that $D(z_0) = 1$.

In Font-Ribera et al. (2014); Aghamousa et al. (2016a), the change in galaxy bias itself with z is parameterized by a rescaling of the growth factor, such that $b_g(z) = b_g(z=0)/D(z)$. For theoretical calculations, this allows us to relate the observed galaxy field to the matter density field we are interested in probing for measuring cosmological parameters. I use this in Chapter 6. Furthermore, as we look to more distant galaxies at higher redshifts, we have a manner in which to describe how the galaxy density field we observe will evolve over time (and thus z).

The growth of matter density perturbations

It is also important to understand how the density perturbations δ grow over time. For the next few paragraphs I have relied on the description in Schneider (2006) and lecture notes by Huterer (2020) to describe the growth of perturbations in linear theory. In practice, cosmologists use codes such as CAMB³ or CLASS (Blas et al., 2011) (which I will refer to in the coming chapters) that provide a theoretical estimate of the matter power spectrum (or correlation function) solve for the growth of perturbations numerically. Two equations are required; firstly the continuity equation, which describes the behaviour of a fluid with density ρ and ensures mass conservation,⁴

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}(\mathbf{x})) = 0. \quad (3.5)$$

³<https://camb.info/>

⁴One can expand this here using $\nabla \cdot (\rho \mathbf{v}(\mathbf{x})) = \nabla \rho \cdot \mathbf{v}(\mathbf{x}) + \rho \nabla \cdot \mathbf{v}(\mathbf{x})$.

The matter is assumed to behave as a pressureless fluid in the above form (which we would additionally assume when writing the full Friedmann equation to derive the Universe's expansion dynamics). It should be noted here \mathbf{v} is the velocity field which accounts for motions due to expanding spacetime and due to gravitational interactions (peculiar motions). This equation can be expressed in comoving coordinates as

$$\frac{\partial \rho}{\partial t} + 3\frac{\dot{a}}{a}\rho + \frac{1}{a}\nabla \cdot (\rho \mathbf{u}(\mathbf{y})) = 0 \quad (3.6)$$

where \mathbf{u} is the peculiar velocity field (expansion is factored out); we change to coordinate \mathbf{y} to indicate we have comoving rather than proper spatial coordinates. Secondly, the Euler equation ensures the conservation of momentum and describes the influence of gravitational forces on matter. Assuming pressureless matter allows us to write

$$\frac{\partial \mathbf{v}(\mathbf{x})}{\partial t} + (\mathbf{v}(\mathbf{x}) \cdot \nabla) \mathbf{v}(\mathbf{x}) = -\nabla \Phi(\mathbf{x}, t). \quad (3.7)$$

Here Φ is the gravitational potential. This can be expressed in comoving coordinates as

$$\frac{\partial \mathbf{u}(\mathbf{y})}{\partial t} + \frac{(\mathbf{u}(\mathbf{y}) \cdot \nabla) \mathbf{u}(\mathbf{y})}{a} + \frac{\dot{a}}{a} \mathbf{u}(\mathbf{y}) = -\nabla \phi(\mathbf{y}, t). \quad (3.8)$$

ϕ is the gravitational potential in comoving coordinates. Then, in the case of only small perturbations from a perfectly homogeneous Universe (linear theory), applying the Poisson equation $\nabla^2 \phi = 4\pi G a^2 \bar{\rho} \delta(\mathbf{y})$, and dropping all non-linear terms that arise, it is possible to use the continuity equation and Euler equation to determine another equation for the growth of the density field δ , using $\rho = \bar{\rho}(1 + \delta)$, as

$$\frac{\partial^2 \delta}{\partial t^2} + \frac{2\dot{a}}{a} \frac{\partial \delta}{\partial t} = 4\pi G \bar{\rho} \delta. \quad (3.9)$$

The two solutions to this for δ involve a growing and a decaying solution; it also allows one to write down a solution for the growth factor $D(z)$ assuming the form in equation 3.4 for the growing solution. The linear theory description of the growth of matter perturbations is useful to understand how perturbations evolved during different epochs, such as when radiation, then matter and now dark energy has dominated the expansion history of the Universe. The differential equation above can be solved for limiting cases, and one finds during radiation domination the growth of structure is very slow, then proportional to $a(t)$ during matter domination, before ceasing to grow during a period of dark energy domination. These equations and description of course are not accurate for very large perturbations, or when the assumption of pressureless matter or linear structure growth begins to break down. A useful quantity that can be derived from these solutions is the growth rate of structure $f(a)$, which I mention here since it appears in various equations in this chapter and those following. $f(a)$ is defined as

$$f(a) = \frac{d \ln D(a)}{d \ln a}. \quad (3.10)$$

The growth rate which describes how quickly perturbations are growing under the influence of gravity, can also be thought of as the speed with which galaxies fall towards each other. $f(a)$ can be measured from galaxy surveys and serves as a test of General Relativity and Λ CDM. The growth rate of structure as a function of time (scalefactor) is shown in Figure 3.1.




Figure 3.1: The growth rate of structure $f(a)$, shown as a function of scalefactor a . The growth rate has been computed using the CLASS Boltzmann-solver, for a cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.31$, $\Omega_\Lambda = 0.69$. The broad and flat peak of this function coincides with the matter-domination epoch, discussed in Chapter 2.

The peculiar velocity field

Recall that \mathbf{v} is the velocity field due to both expanding space (the Hubble flow) and due to gravitational interactions, known as peculiar velocities. Thus \mathbf{v} at comoving coordinate \mathbf{r} can be expressed as a field $\mathbf{v}_{\text{total}}$ having two components,

$$\begin{aligned} \mathbf{v}_{\text{total}}(\mathbf{r}) &= \frac{d}{dt}(a\mathbf{r}) \\ &= \frac{da}{dt}\mathbf{r} + a\frac{d\mathbf{r}}{dt} \\ &= H(t)a(t)\mathbf{r} + \mathbf{v}(\mathbf{r}). \end{aligned} \quad (3.11)$$

The first component represents the Hubble flow, and the second the peculiar velocity field. When we observe a galaxy, its motion is due to both its recession velocity that contains information about the expanding Universe and a velocity that arises from gravitational interactions, informing us about gravity and the large-scale structure. Both components of motion impact the redshift we measure for a galaxy (a peculiar motion induces a Doppler shift), so assuming a redshift can be related to a distance given by equation 2.13 leads to *redshift-space distortions*; we will revisit this again soon. In order to use peculiar velocities to probe the matter density field, note in linear theory, the three-dimensional peculiar velocity field can be related to the matter field δ defined previously via

$$\nabla \cdot \mathbf{v} = -aH(a)f(a)\delta(\mathbf{r}). \quad (3.12)$$

For now on we will take \mathbf{v} to only refer to the peculiar velocity field. We can also see now how the growth rate $f(a)$ relates to the speed at which galaxies fall towards each other. It is useful also to write down the velocity divergence field $\theta(\mathbf{r})$ as

$$\theta(\mathbf{r}) = \frac{-\nabla \cdot \mathbf{v}}{aH(a)f(a)}, \quad (3.13)$$

which is equivalent to δ in linear theory. Unlike the galaxy density field δ_g which is a biased tracer of the underlying matter density field, the motions of galaxies are considered an unbiased tracer of

the underlying matter density field. With these measurements (such as in the CosmicFlows-4 sample; Tully et al., 2023) we can test Λ CDM, for example by measuring the bulk flow, as the Λ CDM model allows us to predict the expected distribution of these flows on different scales. However, we will see later in this chapter that the signal-to-noise ratio is much lower for measurements of the galaxy peculiar velocity field than the galaxy density field, but they provide useful complementary information and tests, as well as allowing us to disentangle how z for a galaxy is impacted by expanding space vs gravity. Peculiar velocities are measured usually using a scaling relation, such as the Fundamental Plane (Djorgovski and Davis, 1987; Dressler et al., 1987) or the Tully-Fisher relation (Tully and Fisher, 1977) which are methods of determining a measurement of the distance to a galaxy. The measurement of a log-distance ratio from these relations is largely independent of a cosmological model.⁵ Distances can also be measured using Type Ia SNe (Phillips, 1993).

3.1.2 Galaxy correlation functions

The galaxy correlation function, mentioned briefly in Chapter 2 and its Fourier-transform, the matter power spectrum, contain a wealth of information and are extremely useful for understanding the impact of cosmological parameters on the large-scale structure. For this reason, and since we will revisit these in Chapter 5, I will describe them mathematically here. The correlation function of galaxies is usually denoted $\xi_g(r)$, and expresses the excess probability of finding galaxies as a function of their separation in space. The two-point correlation function is given mathematically as

$$\xi_g(r) = \langle \delta_g(\mathbf{x}) \delta_g(\mathbf{x} + \mathbf{r}) \rangle = \langle \delta_g(x) \delta_g(x + r) \rangle. \quad (3.14)$$

This can be estimated in various ways, but typically it involves comparing the distribution of galaxies in a survey to a completely random distribution of galaxies, such as from a mock catalogue. Due to the assumption of homogeneity and isotropy, there is no need to write $x + r$ as a vector; the two-point correlation function depends only on the separation r between galaxies and not the direction. A correlation function can be defined in an analogous way for the matter density field $\delta(\mathbf{x})$; then we can write in linear theory $\xi_g(r) = b_g^2 \xi(r)$. A correlation function can also be written more generally for *any* field, and thus we can write down a correlation function for the galaxy velocity divergence field $\theta(x)$ as $\langle \theta(x) \theta(x + r) \rangle$.

3.1.3 The galaxy power spectrum

The galaxy density field in a comoving volume V can be represented in Fourier space as

$$\delta_g(k) = \frac{1}{\sqrt{V}} \int e^{ikx} \delta(x) d^3x. \quad (3.15)$$

Likewise for the velocity divergence field we can write

$$\theta(k) = \frac{1}{\sqrt{V}} \int e^{ikx} \theta(x) d^3x. \quad (3.16)$$

⁵The log-distance ratio gives the log of the ratio of the true distance to the galaxy to that predicted by a cosmological model given the observed redshift.

We can write for the correlation function in Fourier space, for the matter field, using $y = x + r$,

$$\begin{aligned}
\langle \delta_g(k) \delta_g^*(k) \rangle &= \frac{1}{V} \int_V \int_V \langle e^{i\mathbf{k}\cdot\mathbf{x} - i\mathbf{k}\cdot\mathbf{y}} \delta_g(\mathbf{x}) \delta_g^*(\mathbf{y}) \rangle d^3\mathbf{x} d^3\mathbf{y} \\
&= \frac{1}{V} \int_V \int_V \langle e^{-i\mathbf{k}\cdot\mathbf{r}} \delta_g(x) \delta_g(x+r) \rangle d^3x d^3(x+r) \\
&= \frac{1}{V} \int_V d^3x \int_V e^{-i\mathbf{k}\cdot\mathbf{r}} \xi_g(r) d^3(x+r) \\
&= \int_V e^{-i\mathbf{k}\cdot\mathbf{r}} \xi_g(r) d^3r = P_{gg}(k) = b_g^2 P_{mm}(k),
\end{aligned} \tag{3.17}$$

and similarly for $\theta(x)$

$$\langle \theta(k) \theta^*(k) \rangle = P_{\theta\theta}(k) = P_{mm}(k). \tag{3.18}$$

The Fourier transform of the galaxy density (galaxy velocity divergence) correlation function gives the definition of the galaxy density (galaxy velocity divergence) power spectrum, $P(k)$. If we are able to exploit spherical symmetry, we can write $d^3r = \sin(\theta) d\theta d\phi r^2 dr$, let $\mathbf{k} \cdot \mathbf{r} = k r \cos(\theta) = k r \mu$, giving $d^3r = -d\mu d\phi d^3r$ and simplify this as (effectively perform a Hankel transform),

$$\begin{aligned}
P(k) &= \int_V e^{-i\mathbf{k}\cdot\mathbf{r}} \xi(r) d^3r = \int_\phi \int_r \int_\mu e^{-i\mathbf{k}\cdot\mathbf{r}} \xi(r) d\phi (-d\mu) r^2 dr \\
&= 2\pi \int_r \int_\mu e^{-i\mathbf{k}\cdot\mathbf{r}} \xi(r) (-d\mu) r^2 dr \\
&= 4\pi \int_r \frac{\sin(kr)}{k} \xi(r) r dr.
\end{aligned} \tag{3.19}$$

With data it is possible to use either the power spectrum or correlation function in order to test cosmological models. However, using peculiar velocities rather than the density field, we do not exactly measure $P_{\theta\theta}(k) = P_{mm}(k)$, but instead from the motions of galaxies we construct a quantity $P_{uu}(k)$ which represents the Fourier transform of the correlation function of the peculiar velocity field along the line-of-sight, $\langle u(x)u(x+r) \rangle$. This is related to $P_{\theta\theta}(k)$ via

$$P_{uu}(k) = \left(\frac{aH(a)f(a)\mu}{k} \right)^2 P_{\theta\theta}(k), \tag{3.20}$$

which can be derived from the relation between $\mathbf{v}(x)$ and $\theta(x)$ given in equation 3.13, which gives $u(k) = -i \frac{aH(a)f(a)\theta(k)\mu}{k}$; the factor of μ arises because we can only observe the peculiar velocity field along the line-of-sight, $\mathbf{u} = \mathbf{v} \cdot \hat{\mathbf{r}}$.⁶ As such, surveys measuring the velocity power spectrum and comparing this to the galaxy redshift power spectrum can measure the growth rate $f(a)$.

Since the matter power spectrum $P_{mm}(k)$ is useful for understanding the impacts of cosmological parameters, we will discuss the impact of the Λ CDM parameters here. I show the effect of varying Ω_c (the cold dark matter density only), Ω_b , n_s , H_0 , $\sum m_\nu$ and N_{eff} in Figure 3.2; we show the effect at both $z = 0$ corresponding to the present day and at $z = 1$ corresponding to the time at which the Universe had only a fraction ($\sim 1/8$) of its present-day volume. Varying τ_{re} impacts the cosmic microwave background (CMB) data, but has no impact on the matter power spectrum. For A_s , we do not show the effect since it is only a constant rescaling of the power spectrum amplitude for all scales.

⁶One also needs to use $\mathcal{F}[\nabla \cdot \mathbf{u}(r)] = ik \cdot \mathbf{u}(r)$ for the field \mathbf{u} which approaches zero as r approaches infinity.

The impact of n_s shown in the first panel (from left to right, top to bottom) is simply to tilt the power spectrum. The impact of H_0 , which for a larger value represents a larger expansion rate of the Universe, is to shift the position of the BAO wiggles and turnover scale; additionally it decreases the impact of photons diffusing in a process called *Silk Damping* resulting in a change in power on smaller scales; we will revisit this later in the chapter. Increased H_0 additionally reduces the time matter fluctuations have to grow. The impact of altering Ω_m is complex, but a greater matter density changes the epoch at which matter-radiation equality occurs (the power spectrum turnover), alters the expansion rate and positions of the turnover and wiggles, and thus also the amount of time perturbations grow in the radiation domination or matter domination era. Likewise a greater matter density enhances matter perturbations. Larger N_{eff} increases the early Universe expansion rate by the contribution of N_{eff} to the radiation density, which shifts the position of the turnover scale and wiggles also. However, another impact is in opposition to that of H_0 , which is to wash out power on smaller scales because greater N_{eff} enhances the impact of Silk Damping. N_{eff} also induces a phase shift in the BAO feature, which is the subject of Chapter 5.

Increasing the ratio of Ω_b -to- Ω_c enhances the peaks in the BAOs (and likewise in the CMB peaks), since the photons interact with the baryons and not the dark matter during the BAOs. This increases the depth of rarefactions and compressions in this period. Finally, we can see the impact of $\sum m_\nu$ which has two main effects. As mentioned previously, since $\sum m_\nu$ contributes to the radiation density and then later the matter density in the Universe, there is some impact of the presence of massive neutrinos on the epoch of matter-radiation equality. Secondly, massive neutrinos tend to reduce power on smaller scales, due to the fact they free-stream after they decouple from the photon-baryon plasma in the early Universe.

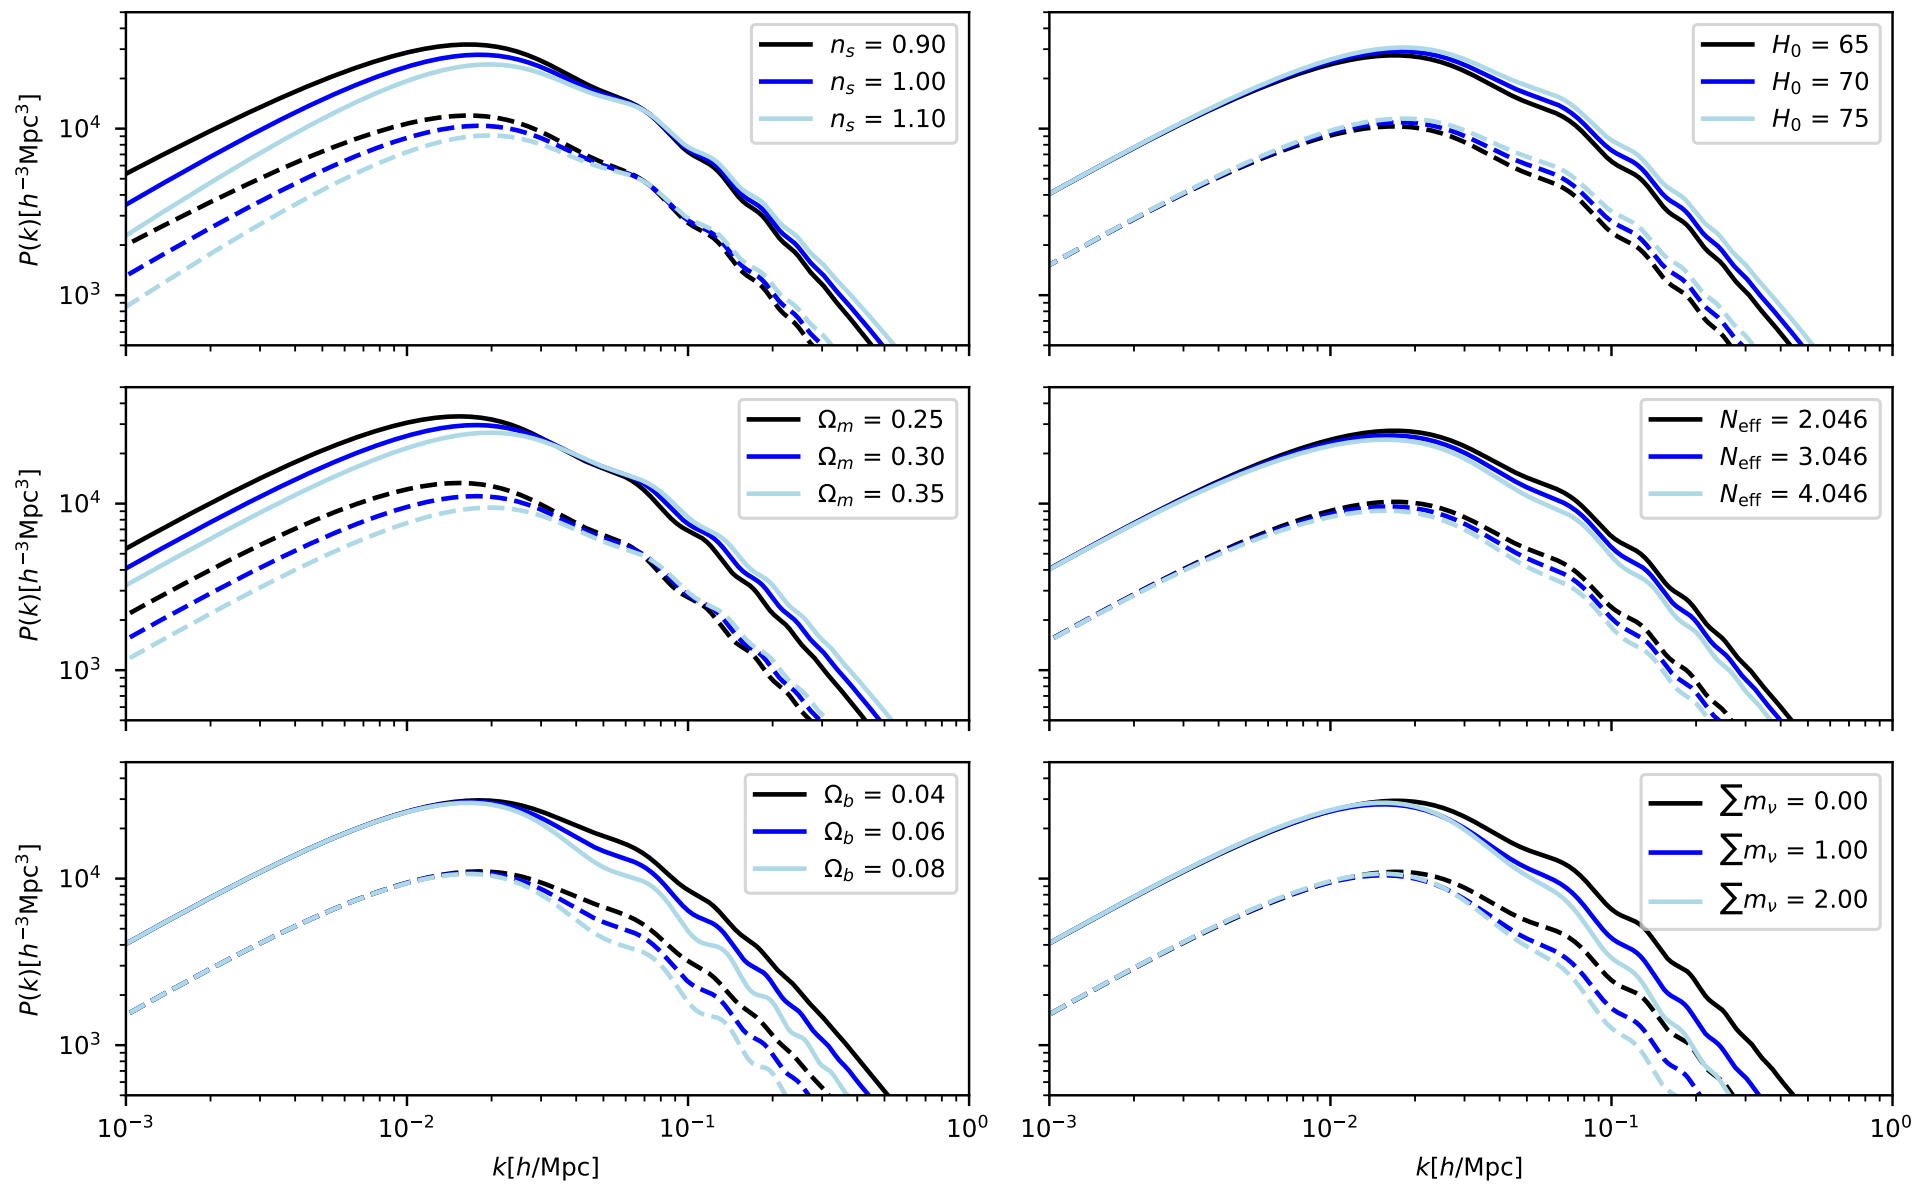


Figure 3.2: The impact of varying (from left to right, top to bottom) n_s , H_0 ($\text{kms}^{-1}\text{Mpc}^{-1}$), Ω_m , N_{eff} , Ω_b (and thus also $\frac{\Omega_b}{\Omega_c}$) and $\sum m_\nu$ (eV). We ensure that Ω_m remains fixed and keep the ratio $\frac{\Omega_b}{\Omega_c}$ fixed while varying $\sum m_\nu$. Dashed lines show the power spectrum at $z = 1$, solid lines at $z = 0$.

Noise in galaxy surveys

When measuring the galaxy power spectrum or the density field or the peculiar velocity field, shot noise arises due to the fact galaxies are effectively ‘randomly’ sampled relative to the total field of galaxies. This introduces an uncertainty into the measurement of the power spectrum, with the noise given as $\frac{1}{n}$ where n is the number density of measurements. Therefore the noise depends on the selection function of the survey (the distribution of observed galaxies in space). For peculiar velocity surveys, one needs to consider additional sources of noise due to the fact most scaling relations (Tully-Fisher, Fundamental Plane) that measure the distances to galaxies have an intrinsic scatter that is usually $\sim 20\%$.

Another source of uncertainty on measurements of the power spectrum (or any cosmological measurement) is known as *cosmic variance*. This is the statistical uncertainty in our measurements that arises due to the fact we have only have a single view of the Universe from our position; it is also generally considered a lower limit on our measurement uncertainty that we can ever possibly obtain. Feldman et al. (1994) shows the best possible approach to measure the galaxy power spectrum to reduce the noise in the measurement involves using weights for the estimation of the matter fluctuation field, known as FKP weights. Following this result, Tegmark (1997) shows an approximation for the theoretical covariance matrix of the i^{th} spatial volume element of the galaxy power spectrum $P(k)$ is written as

$$C_{ij} \approx 2 \frac{P_i(k)P_i(k)}{V_i V_{\text{eff},i}} \delta_{ij}. \quad (3.21)$$

The volume of the spatial element V_i is given as $4\pi k_i^2 dk_i / (2\pi)^3$, $V_{\text{eff},i}$ is the effective volume given as

$$V_{\text{eff},i} = \int_V \left(\frac{nP_i(k)}{1 + nP_i(k)} \right)^2 d^3r. \quad (3.22)$$

Despite the apparent limitations of cosmic variance, it nevertheless is possible at times to make measurements with constraining power below the cosmic variance limit by combining observations with simulations, which is discussed and shown by Sorce et al. (2016) who use observations of peculiar velocities from the CosmicFlows-2 catalogue. Essentially, one can produce simulations for the local Large-scale structure in the Universe and constrain the initial conditions by using the observational data, which exhibit a lower variance in the density field compared to random simulations.

3.1.4 Galaxy surveys in redshift space

So far we have discussed the fields δ and θ that we can probe with galaxy surveys to inform us on structure growth in the Universe, and how these can be used to test models or used to construct the power spectrum or correlation function which contains a wealth of information about cosmological parameters. However, the process of comparing these quantities to a model like Λ CDM is not really quite so simple, for at least two reasons. The following discussion is essential for correctly modelling the observed power spectrum or correlation function prior to comparing to a model, and is thus most important for Chapter 5 in which we do so. Firstly, recall that galaxies have peculiar motions due to

gravitational interactions, and this impacts their observed redshift z_{obs} due to the Doppler effect,

$$z_{\text{obs}} = (1 + z_{\text{rec}})(1 + z_{\text{pec}}) - 1. \quad (3.23)$$

z_{rec} is the motion of the galaxy due to expanding space, which is the redshift component that we require to enter equation 2.13 to compute the distance to a galaxy in a cosmological model. z_{pec} is the redshift that arises due to the Doppler effect due to the galaxy's peculiar motion. Thus the positions of galaxies in redshift space (referring to the angular and redshift coordinates of galaxies) have distortions that need to be accounted for. One can also see how equation 3.23 allows one to extract information about the peculiar velocity of a galaxy, using the relation between z_{pec} and \mathbf{u} for the Doppler effect,⁷

$$u = \left[\frac{(1 + z_{\text{pec}})^2 - 1}{(1 + z_{\text{pec}})^2 + 1} \right]. \quad (3.24)$$

Secondly, there is a somewhat circular aspect to assuming a cosmological model in order to transform galaxy redshifts into distances prior to comparing the constructed correlation function or power spectrum with a cosmological model. However, there exists a clever way to deal with this; the process of transforming redshifts to distances under the assumption of a model that differs from the real cosmology of the Universe results in its own distortions to the distribution of galaxies. A spherical distribution of galaxies around the position of a central observer may transform into an ellipsoid, and the extent of this distortion is related to how much the assumed cosmological model differs from the true cosmology. This is called the AP effect (Alcock and Paczyński, 1979), and will be discussed shortly.

Redshift-space distortions

An observer who measures the redshift of a galaxy from a spectrum will find a component due to the expansion of space z_{rec} and a component due to the galaxy's motion along the line-of-sight z_{pec} ; z_{pec} could indicate a motion towards or away from the observer. If an observer views a cluster of galaxies, those galaxies furthest away from the observer are most likely to have a peculiar motion towards the observer due to the gravitational influence of the cluster, and vice versa for galaxies closest to the observer. This causes the shape of a cluster of galaxies to be 'squished' along the line-of-sight. This is called the Kaiser effect (Kaiser, 1987). Due to the Kaiser effect, the redshift-space power spectrum of the galaxy density field is expressed mathematically as

$$P(k, z, \mu) = (b_g + f\mu^2)^2 P_{\text{mm}}(k, z). \quad (3.25)$$

Here $P(k, z, \mu)$ is the redshift-space galaxy power spectrum, and $P_{\text{mm}}(k, z)$ is the linear matter power spectrum. Here b_g is the galaxy bias introduced earlier and f is the growth rate of structure. μ is the

⁷In reality, extracting the peculiar motion of a galaxy is often not quite so simple as implied here and is mildly dependent on a choice of cosmological model via Equation 2.13. This is because obtaining z_{pec} (or equivalently the difference in the distance from a scaling relation compared to the expected distance to a galaxy given z_{obs}) requires one to make assume a cosmological model to obtain $d(z_{\text{obs}})$. There is mainly a dependence on Ω_m and H_0 that affects this measurement.




Figure 3.3: The impact of redshift space distortions on the correlation function, due to the Kaiser effect and the FoG (Finger-of-God) effect. π gives separations along the line-of-sight, while σ perpendicular to the line-of-sight. The redshift space correlation function was computed via a Hankel transform of the power spectrum moments $P_\ell(k)$ to obtain the configuration space moments $\xi_\ell(s)$, which can be summed to calculate $\xi(\mu, s)$ as $\xi(\mu, s) = \sum_\ell \xi_\ell(s) \mathcal{P}_\ell(\mu)$ (see equation 3.28).

previously introduced $\cos(\theta)$ where θ is the angle relative to the line-of-sight. Perpendicular to the line-of-sight, $\mu = 0$. The Kaiser effect occurs maximally parallel to the line-of-sight, $\mu = 1$ for $\theta = 0$. The Kaiser factor above is not applied to the power spectrum of peculiar velocities since the velocity field is not believed to be biased, so there is no factor b_g ; as such peculiar velocities can help reduce degeneracy between cosmological parameters and b_g . However, as shown in Equation 3.20, there is dependence on angle via μ since the peculiar velocity field is measured along the line-of-sight. The growth rate of structure also enters since this is related to the speed at which galaxies fall towards each other.

Another effect occurs due to the non-linear motions of galaxies. In the center of a galaxy cluster, the galaxies may have virialized and have strong orbital motions around the central mass of the cluster. The components of this orbital motion that are along the line-of-sight results in an apparent ‘squishing’ of the positions of galaxies perpendicular to the line-of-sight, which is the opposite to the Kaiser effect which ‘squishes’ along the line-of-sight. However, it occurs on smaller scales, for those galaxies very close to the center of a cluster. This effect is known as the ‘Finger-of-God’ (FoG) effect (Jackson, 1972). These effects are demonstrated in Figure 3.3.

Accounting for this in the modelling of the redshift-space galaxy power spectrum of the density field gives

$$P(k, z, \mu) = (b_g + f\mu^2)^2 P_{\text{mm}}(k, z) \left[1 + \frac{1}{2}k^2 \mu^2 \sigma_g^2\right]^{-2}, \quad (3.26)$$

with the factor $D_g = \left[1 + \frac{1}{2}k^2 \mu^2 \sigma_g^2\right]^{-2}$ motivated by results from simulations (Ballinger et al., 1996). The velocity power spectrum that we measure in redshift-space is also impacted by the non-linear

motions of galaxies, by a factor $D_u(k) = \frac{\sin(k\sigma_u)}{k\sigma_u}$ (found from simulations by Koda et al., 2014), leading to

$$P_{uu}(k) = \left(\frac{aH(a)f(a)\mu}{k} \right)^2 D_u^2(k) P_{\theta\theta}(k). \quad (3.27)$$

Generally σ_g and σ_u are treated as free parameters in a standard analysis. Often, observed power spectra are decomposed into moments using a Legendre polynomial expansion,

$$P_\ell(k, z) = \frac{2\ell + 1}{2} \int_{-1}^1 d\mu P(k, \mu, z) \mathcal{P}_\ell(\mu), \quad (3.28)$$

where $\mathcal{P}_\ell(\mu)$ are the Legendre polynomials. This allows us to write the following expression for the power spectrum as a summation over the power spectrum moments,

$$P(k, \mu, z) = \sum_{\ell} P_\ell(k, z) \mathcal{P}_\ell(\mu). \quad (3.29)$$

This is useful if the constraining power of the survey allows one to only measure the first moment known as the monopole $P_{\ell=0}(k, z)$, which is the average power spectrum over all lines-of-sight in the sky. Often one measures the monopole, the quadrupole ($\ell = 2$) and the hexadecapole ($\ell = 4$) since the odd moments of the redshift-space power spectra should theoretically be zero due to symmetry about the line-of-sight.

Alcock-Paczynski effect

Our second aforementioned complexity is that the process of transforming galaxy redshifts to distances requires an assumption of a cosmological model. This process distorts an otherwise perfectly spherical distribution of galaxies about an observer to a non-spherical distribution. This is called the Alcock-Paczynski effect (Alcock and Paczyński, 1979). Mathematically, this can be expressed as a distortion to the distances we measure to galaxies along the line-of-sight and perpendicular to the line-of-sight. While assuming a cosmological model may seem problematic when the goal is to test models of cosmology, this effect actually allows us to gain information about the Universe. The level of distortion that occurs contains information about cosmological parameters; measuring the distortions provide a test of cosmology.

Along the line-of-sight, the apparent distances we measure between objects, given by d_{\parallel}^f , are related to the real distance that is measured, by $d_{\parallel}^f = q_{\parallel} d_{\parallel}$. The apparent distances between objects perpendicular to our line-of-sight, d_{\perp}^f are given by $d_{\perp}^f = q_{\perp} d_{\perp}$. The distortion parameters q_{\parallel} , q_{\perp} are

$$\begin{aligned} q_{\parallel} &= \frac{H^f(z)}{H(z)}, \\ q_{\perp} &= \frac{d_A(z)}{d_A^f(z)}. \end{aligned} \quad (3.30)$$

All quantities with the superscript ‘f’ refer to the value of the quantity in the assumed (fiducial) cosmological model. Thus, measuring q_{\parallel} or q_{\perp} can allow us to constrain $H(z)$ or $d_A(z)$. It is straight forward to see how these parameters are derived; distances between objects along the line-of-sight

are given by $\delta d = \frac{c\delta z}{H(z)}$, while distances perpendicular to the line-of-sight can be expressed like $\delta d = \delta\theta d_A(z)$, where the angular diameter distance $d_A(z)$ must be used to express the apparent distance using an angular size on the sky. In the next section we will see how the AP effect can be exploited in the context of measuring baryon acoustic oscillations, which is also specific to Chapter 5.

3.1.5 Analysis with baryon acoustic oscillations

Previously I discussed how cosmological parameters impact the matter power spectrum. One approach to test cosmology is to directly fit the discussed cosmological parameters to the full shape of the galaxy power spectrum. However, another approach is to simply fit the oscillatory feature, the BAO wiggles. This approach has a benefit in that it is semi model-independent,⁸ and because one fits parameters controlling the shape of the oscillatory feature and then only afterwards interprets their constraints on cosmological parameters; thus the constraints can be easily tested against various different models. This approach however does require an algorithm, such as those defined in Hinton et al. (2016); Baumann et al. (2018), to smooth a template power spectrum to isolate the oscillatory feature. The oscillatory feature can be defined (and isolated from a smoothed spectrum) as

$$\mathcal{O}(k) = \frac{P(k)}{P_{\text{nw}}(k)} - 1, \quad (3.31)$$

where $P_{\text{nw}}(k)$ is the smoothed power spectrum. The function $\mathcal{O}(k)$ is plotted in Figure 3.4.




Figure 3.4: The oscillatory BAO feature, computed from a Λ CDM model matter power spectrum with $\Omega_m = 0.31$, $H_0 = 67 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

In a typical BAO analysis, the features of the data are compared to a template that can be modified by parameters that shift the positions of the oscillations in the template along the x -axis. Using the monopole of the power spectrum only, one can analyze the *isotropic* shift required for the positions of the peaks and troughs in a BAO template in order for the template to match the data. This is done using a parameter α which enters the template like

$$\mathcal{O}'(k') = \mathcal{O}(k') = \mathcal{O}(k/\alpha). \quad (3.32)$$

⁸The broadband shape of the power spectrum is fit using a parametric model.

Here, k' are the observed wavenumbers, k are the real wavenumbers. α is closely related to the parameters q_{\parallel} and q_{\perp} mentioned earlier, and is given by

$$\alpha = \frac{D_V(z) r_s^t(z_d)}{D_V^f(z) r_s(z_d)}. \quad (3.33)$$

$D_V(z)$ corresponds to the volume averaged distance,

$$D_V(z) = \left((1+z)^2 d_A^2(z) \frac{cz}{H(z)} \right)^{1/3}, \quad (3.34)$$

and $r_s(z_d)$ is the sound horizon at the baryon-drag epoch z_d , and gives the periodicity of the BAO wiggles. The parameter α is related to q_{\parallel} and q_{\perp} as $\alpha = q_{\parallel}^{1/3} q_{\perp}^{2/3} \frac{r_s^t(z_d)}{r_s(z_d)}$. The superscript ‘t’ refers to the cosmology of the template model, which does not necessarily have to be the same as the fiducial cosmology (‘f’) used when transforming galaxy redshifts to distances (resulting in the AP effect). The ratio $\frac{r_s^t(z_d)}{r_s(z_d)}$ enters because the distance r_s^t in the cosmological model of the template needs to be matched to that of the data r_s . By finding the best choice of α to match the template to the data, we can start to constrain distances by D_V , which can be compared to cosmological models.

In an anisotropic analysis, one measures the BAO wiggles using both the monopole and quadrupole (the quadrupole gives enough information now for us to separate impacts of AP distortions along the line-of-sight and perpendicular to it). We can define the parameters

$$\begin{aligned} \alpha_{\parallel} &= q_{\parallel} \frac{r_s^t(z_d)}{r_s(z_d)}, \\ \alpha_{\perp} &= q_{\perp} \frac{r_s^t(z_d)}{r_s(z_d)}, \end{aligned} \quad (3.35)$$

and thus $\alpha = \alpha_{\parallel}^{1/3} \alpha_{\perp}^{2/3}$. These parameters enter the template like

$$\mathcal{O}'(k') = \mathcal{O}(k'(k, \mu), \mu'((k, \mu))) = \mathcal{O}(k'(\alpha_{\parallel}, \alpha_{\perp}), \mu'(\alpha_{\parallel}, \alpha_{\perp})). \quad (3.36)$$

μ is as defined earlier, the cosine of the angle θ , where θ is the angle relative to the line-of-sight, and μ' is the observed angle. The observed k' , μ' can be related to the real k , μ as

$$\begin{aligned} k' &= \frac{k}{q_{\perp}} \left[1 + \mu^2 \left(\frac{1}{F^2} - 1 \right) \right]^{1/2}, \\ \mu' &= \frac{\mu}{F} \left[1 + \mu^2 \left(\frac{1}{F^2} - 1 \right) \right]^{-1/2}, \end{aligned} \quad (3.37)$$

where $F = q_{\parallel}/q_{\perp}$. These equations can be derived using $k = \left(k_{\parallel}^2 + k_{\perp}^2 \right)^{1/2}$, where $k'_{\parallel} = \frac{1}{q_{\parallel}} k_{\parallel}$ and $k'_{\perp} = \frac{1}{q_{\perp}} k_{\perp}$. The effect of the AP distortions that is used to constrain cosmological parameters is shown in Figure 3.5 for the galaxy correlation function.

With this section, we have thoroughly discussed many of the fundamental technical details essential for modelling the galaxy power spectrum and oscillatory feature in the power spectrum for obtaining fits to cosmological parameters. In Chapter 5, we will discuss fits to the BAO oscillatory feature with α_{\parallel} and α_{\perp} , as well as a more novel parameter β_{ϕ} that captures the impacts of free-streaming neutrinos in the early Universe. I discuss more about the impacts of neutrinos on cosmological data towards the end of this chapter.




Figure 3.5: A demonstration of the Alcock-Paczynski Effect, shown in the galaxy correlation function. Here I have plotted the Fourier transform of the matter power spectrum to give the correlation function, as a function of coordinates π and σ , which are related to r as $r = \sqrt{\sigma^2 + \pi^2}$. π and σ show the parallel and perpendicular components of galaxy separation r , and the contours show the correlation function $r^2 \xi(r)$. The left panel shows the correlation function plotted as a function of true separations r , but in the right panel the separations plotted are the observed separations which are distorted by the AP Effect when the true cosmology (corresponding to $\Omega_m = 0.31$, $\Omega_\Lambda = 0.69$) does not match the fiducial cosmology ($\Omega_m = 0.20$, $\Omega_\Lambda = 0.80$). The black dashed circle is a reference to see the AP distortions.

3.1.6 Galaxy bulk flows

For this section I return back to earlier discussion on directly probing the peculiar velocity field of galaxies for tests of cosmology. Due to the large uncertainties in peculiar velocity measurements, particularly due to non-linear motions on very small scales, a historical approach to gaining information from peculiar velocities is to use a statistic called the *bulk flow*. The bulk flow \mathbf{B} is a measurement of the average of galaxy peculiar velocities in a given survey volume V ,

$$\mathbf{B} = \frac{1}{V} \int_V \mathbf{v}(\mathbf{r}) d^3 r, \quad (3.38)$$

where $\mathbf{v}(\mathbf{r})$ represents the 3-dimensional PV field; note here we have now switched the notation slightly, to use $r(z)$ as the dimensionful comoving distance coordinate for each galaxy when discussing bulk flow estimators. By averaging many velocity measurements we can probe the motion of galaxies on larger scales and reduce the impact of non-linearities and measurement noise. The measured bulk flow can be compared to the predicted distribution of bulk flow measurements from a Maxwell-Boltzmann distribution, with the variance in the velocities that enters this distribution predicted by the Λ CDM model. In the literature, bulk flows are typically measured in one of two ways (Nusser, 2016); the estimator of the bulk flow components along coordinate axes $\alpha = x, y, z$ can be written as

$$\tilde{B}_\alpha = \frac{1}{V} \int u(r) \hat{n}_\alpha w_\alpha(r) d^3 r \quad (3.39)$$

or alternatively as

$$\tilde{B}_\alpha = \frac{1}{V} \int v_\alpha(r) w_\alpha(r) d^3 r. \quad (3.40)$$

In the first approach, the bulk flow component is estimated from the radial component of the PVs of the galaxies which we are able to observe along the line-of-sight (LOS), $u(r)$, with weighting given by $w_\alpha(r)$. The function $\hat{n}_\alpha = \hat{\mathbf{x}}_\alpha \cdot \hat{\mathbf{r}}$, where $\hat{\mathbf{r}}$ is the unit vector for the galaxy position and $\hat{\mathbf{x}}_\alpha$ is the unit vector for the coordinate axis. In the second estimator, the bulk flow component is estimated from the weighted components of the 3-dimensional PVs, v_α . In the first case an optimal weighting scheme is needed for the radial components of the galaxy peculiar motions to take into account that we do not have full 3-dimensional information for each galaxy's peculiar motion. In the second approach, a weighting scheme is found for the galaxy peculiar motions, when we are able to obtain full 3-dimensional information about their motions. This is generally only possible when prior information is assumed from a cosmological model in order to derive an estimate of the 3-dimensional velocity field.

A widely used method in the literature to estimate the bulk flow is the Maximum Likelihood Method (MLE) of Kaiser (1988); this method is an approach of the first kind mentioned above, and investigation into methods of this kind are the focus of Chapter 4. The likelihood function for this estimator for measuring a bulk flow vector \mathbf{B} given a set of N observed radial PVs u_n is written as

$$\mathcal{L}(\mathbf{B}, u_n) = \prod_{n=1}^N \frac{1}{\sqrt{2\pi(\sigma_*^2 + \sigma_n^2)}} e^{-\frac{(u_n - \mathbf{B} \cdot \hat{\mathbf{r}}_n)^2}{2(\sigma_*^2 + \sigma_n^2)}}, \quad (3.41)$$

where σ_n is the uncertainty in each velocity measurement, σ_* is a 1D velocity dispersion to account for random non-linear motions. We want to calculate the bulk flow components B_i as $B_i = \sum_{n=1}^N w_{i,n} u_n$ where $w_{i,n}$ represents the weight for each velocity measurement. The solution for the weights which maximise this likelihood function with respect to the bulk flow is given by

$$w_{i,n} = \sum_j A_{i,j}^{-1} \frac{\hat{m}_{j,n}}{(\sigma_*^2 + \sigma_n^2)}; \quad A_{i,j} = \sum_n \frac{\hat{m}_{i,n} \hat{m}_{j,n}}{(\sigma_*^2 + \sigma_n^2)}. \quad (3.42)$$

Here $\hat{m}_{j,n} = \hat{\mathbf{x}}_j \cdot \hat{\mathbf{r}}_n$. This likelihood function assumes the galaxy PVs are *uncorrelated*, the PVs have gaussian PDFs for their uncertainties, and the bulk flow vector is a *constant* across the survey volume. We will come back to this point in Chapter 4 in which we discuss bulk flows in detail.

Another method of the first approach mentioned previously is the Minimum Variance Method of Watkins et al. (2009); Feldman et al. (2010). In this approach, we attempt to measure the bulk flow from a survey that has an ideal spherical distribution of galaxies on the sky; this is because the geometry of our galaxy survey can significantly impact the measured bulk flow. As such, the goal is to weight the velocity measurements of galaxies appropriately, when the window function of the survey is not necessarily isotropic for each bulk flow component. With the maximum likelihood method just described, we tend to give weight to velocities obtained from small scale flows that are not captured well in our modelling. As such, the bulk flow measurements will be sensitive to scales that are smaller than the depth of the survey. In the Minimum Variance method, the weights are derived

by mathematically minimizing the variance $\langle (b_i - U_i)^2 \rangle$, where U_i represents a measurement of the bulk flow from a survey in which we have a spherical (ideal) distribution of galaxies, and b_i is the bulk flow measurement that we would obtain from our actual galaxy PV sample, that does not have an ideal distribution of galaxies. We can obtain the weights for each galaxy velocity so that $b_i = \sum_m w_{i,m} u_m$. To do this we start with the expression to be minimized,

$$\mathcal{L} = \langle (b_i - U_i)^2 \rangle - \lambda_{ij} \sum_n w_{i,n} g_i(\mathbf{r}_n), \quad (3.43)$$

with respect to the weights $w_{i,n}$, where λ_{ij} is the Lagrange multiplier that ensures the condition $u_n = \sum_i U_i g_i(\mathbf{r}_n)$ is satisfied. $g_i(\mathbf{r}_n)$ are mode functions. Their definition can be understood from an expansion of the LOS peculiar velocity field at some coordinate \mathbf{x} into moments U_i , starting with the bulk flow moments, then higher order modes (such as shear modes),

$$u(\mathbf{x}) = \mathbf{v}(\mathbf{x}) \cdot \hat{\mathbf{r}} = U_i (\hat{\mathbf{x}}_i \cdot \hat{\mathbf{r}}) + U_{ij} x (\hat{\mathbf{x}}_i \cdot \hat{\mathbf{r}}) (\hat{\mathbf{x}}_j \cdot \hat{\mathbf{r}}) + \dots \quad (3.44)$$

$$= \sum_{p=1}^{19} U_p g_p(\mathbf{r}). \quad (3.45)$$

If one counts just the three bulk flow moments U_i , the six shear moments U_{ij} and the ten octopole moments U_{ijk} , there are just 19 moments. In a more simple analysis one can consider just the bulk flow moments. Solving this for a set of weights for our survey gives

$$w_{p,n} = \sum_m G_{nm}^{-1} \left[\langle u_m U_p \rangle - \frac{1}{2} \sum_q \lambda_{pq} g_q(\mathbf{r}_m) \right]. \quad (3.46)$$

We have a solution for the Lagrange multiplier λ_{pq} ,

$$\lambda_{pq} = \sum_l \left[M_{pl}^{-1} \left(\sum_{m,n} G_{mn}^{-1} \langle u_m U_l \rangle g_q(\mathbf{r}_n) - \delta_{lq} \right) \right], \quad (3.47)$$

and

$$M_{pq} = \frac{1}{2} \sum_{n,m} G_{nm}^{-1} g_p(\mathbf{r}_n) g_q(\mathbf{r}_m). \quad (3.48)$$

The matrix, $Q_{i,n} = \langle u_n U_i \rangle$ gives the cross-correlation between the ideal survey (chosen in order to obtain a particular window function as desired) and the real survey as

$$Q_{i,n} = \sum_x w'_{i,n} \langle v'_n v_n \rangle. \quad (3.49)$$

One can calculate the weights for an ideal survey as $w'_{i,n} = 3 \hat{\mathbf{r}}_n \cdot \mathbf{g}_i(\mathbf{r}_n) / N_0$ (N_0 is the size of the ideal survey). $\langle v'_n v_n \rangle$ is the cross-covariance between the galaxy velocities in the real survey and the simulated ideal survey. $G_{mn} = \langle v_m v_n \rangle$ is the covariance matrix of the galaxy peculiar velocity measurements, which can be derived from linear theory. We will also discuss this more in Chapter 4, where we look into the accuracy and precision of these bulk flow estimators. However for the remainder of this theory chapter we will pivot to discuss the impacts of neutrinos in the early Universe on the large-scale structure of the Universe. This final part of the background to my thesis is relevant to Chapter 5 and Chapter 6.

3.2 Neutrinos in cosmology

Neutrinos are one of the most mysterious particles in the Standard Model, due to their nature as a very weakly interacting particle, which was initially believed to be massless. Only due to the discovery of neutrino oscillations between flavour states (Fukuda et al., 1998) are they now known to have a non-zero sum of masses. Oscillation experiments are able to measure the squared differences in mass (called mass splittings) between the different *mass* eigenstates, which are superposition eigenstates of the *flavour* states. However, there are not yet any terrestrial experiments that allow us to measure the absolute masses of any neutrino eigenstate or discern the ordering of the least to the most massive eigenstate (referred to as the mass hierarchy). It is also not well understood how neutrinos gain mass in the Standard Model; they potentially play a role as a ‘portal’ to physics beyond the Standard Model (BSM). Regardless, neutrinos have a role to play in cosmology and the expansion history of the Universe. As mentioned in Chapter 2, it is possible for cosmological probes to constrain $\sum m_\nu$, N_{eff} and signatures of non-standard neutrino interactions.

3.2.1 Impact of neutrinos in the early Universe

During the very early Universe in which baryonic particles formed a hot plasma, various particle species were in equilibrium and interacted with each other. However, these interactions ceased when the interaction rates of the particles dropped below the expansion rate of the Universe. Neutrinos, which are weakly interacting particles, decoupled from this plasma much more quickly than the other particles due to their low interaction cross-section; this is called ‘freeze-out’. Neutrinos began to ‘free-stream’, which refers to them propagating freely through space without further interactions. Their behaviour as relativistic particles (those with kinetic energy greater than their rest mass energy) during the early Universe and then later as non-relativistic particles (when their kinetic energies decreased) also influenced the expansion rate and left a signature on matter perturbations (Lesgourgues and Pastor, 2012). They carried away energy as they free-streamed and their velocities decayed, causing gravitational interactions to be suppressed on small length scales (Lesgourgues and Pastor, 2006, 2012), relative to larger scales on which the neutrinos can cluster and behave like cold dark matter perturbations. The suppression signature which they leave on cosmological probes of matter is dependent on both their masses and the number of different neutrino species. As such, probes of the matter distribution not only to help constrain the masses of neutrinos but also allow us to test for other particle species or test models that invoke similar effects.

3.2.2 Impact of light relics on structure growth

The free-streaming neutrinos from the early Universe, which froze out at a temperature of $\sim 10^{10}$ K, are known as the cosmic neutrino background (CvB). The theoretical number density of neutrinos in the CvB is approximately 336 per cubic centimetre for both neutrinos and anti-neutrinos (more details can be found in Lesgourgues and Pastor, 2006; Vitagliano et al., 2020). Due to their low interaction

cross-section, it is not currently possible for any experiment to directly detect the CνB; it can only be probed indirectly via its effects on the growth of large-scale structure or the CMB.

Since neutrinos were initially in equilibrium with the baryon-photon plasma, if neutrinos and the CMB photons had evolved together without neutrinos decoupling, they would have the same temperature today. The temperature of photons over time has decreased with $(1+z)$ such that $T_\gamma(1+z) = T_{0,\text{CMB}} \approx 2.725$ K (recall this temperature was measured by the COBE satellite, Mather et al., 1994; Fixsen et al., 1996). However the present day temperature of the CνB is actually expected to be ~ 1.95 K (this temperature is calculated theoretically due to a change in temperature that CMB photons receive from a process that injects additional energy into the photon-baryon plasma, see Lesgourgues and Pastor, 2006). The photon-baryon plasma received an energy injection from electrons and positrons annihilating in pairs of photons when they decoupled themselves from the plasma, giving the CMB a boost relative to the CνB.

As mentioned briefly in Chapter 2, the impacts of neutrinos in the CνB on large-scale structure and the CMB is captured by the parameter N_{eff} , the effective number of neutrino species. N_{eff} can be altered by additional sterile neutrinos or if the neutrino temperature is different to what is expected in the Standard Model, via the equation (Renk et al., 2021)

$$N_{\text{eff}} = (\Gamma_\nu)^4 N_\nu + \Delta N_{\text{ur}}. \quad (3.50)$$

N_{ur} allows for contributions from additional ultra-relativistic species, N_ν is the number of standard model neutrinos, and Γ_ν is the ratio of the neutrino temperature to that of the temperature in the Λ CDM model. The power of four can be understood because $\rho_\nu = N_{\text{eff}} \frac{7\pi^2}{120} T_\nu^4$ (Riemer-Sørensen et al., 2013), where T_ν is the neutrino temperature. N_{eff} actually takes the value $N_{\text{eff}} = 3.044$, because neutrinos did not instantly decouple from the plasma of baryons and photons; a fraction of the neutrinos received a tiny energy injection due to annihilation of electrons and positrons, which is well captured by slightly increasing the value of N_{eff} .

The impact of varying N_{eff} in the matter power spectrum can be seen in Figure 3.2. Changing N_{eff} impacts the energy density of neutrinos Ω_ν , which impacts the growth of structure during the radiation-domination epoch via any change this has on the total density of radiation Ω_r . During this epoch, additional radiation retards the growth of structure, and as such length scales inside the causal horizon during the radiation-domination epoch had suppressed growth (at values of k larger than that corresponding to the causal horizon at matter-radiation equality, which is approximately the turnover of the power spectrum). Additionally, altering N_{eff} has an impact on the distance BAOs travelled during the early Universe, given by the sound horizon at the baryon-drag epoch. This is because altering N_{eff} changes the density of radiation which impacts not only the expansion rate but also impacts the fraction of radiation in photons, which impacts the speed of sound at the time, given by $c_s^2 = \frac{\partial p}{\partial \rho}$. Additional or less radiation can also impact *diffusion damping*, which is a process in which photons diffuse and smear the anisotropies in the early Universe. A different scale for diffusion damping impacts the amount of power on small scales in the matter power spectrum, which has less power if diffusion damping is greater.

Finally, a robust impact of free-streaming relics is to induce a phase shift in the power spectrum wiggles (the BAO signal, shown in Figure 3.4). Since free-streaming neutrinos carry energy as they propagate they induce a temporal shift in the gravitational potential of the baryon-photon plasma. This impact leaves a physical phase shift in the BAO signal, with an amplitude that increases with larger N_{eff} . This will be discussed in greater detail in Chapter 5 and Chapter 6.

Existing constraints on light relics

I mentioned in Chapter 2 that tight constraints on N_{eff} can be obtained from the Planck measurements of the CMB, which finds $N_{\text{eff}} = 2.99 \pm 0.17$ (Aghanim et al., 2020). Like in the matter power spectrum, altering N_{eff} impacts the CMB temperature power spectrum through diffusion damping, which can be seen at large values of ℓ (corresponding to small scales) in the spectrum. The left most peak of the CMB, which holds information about the expansion rate of the Universe, is also shifted by altering N_{eff} , in addition to a subtle phase shift in all the peaks of the spectrum due to the free-streaming nature of neutrinos. Figure 3.6 shows the impact of altering N_{eff} on the CMB temperature power spectrum.




Figure 3.6: The CMB temperature power spectrum, which shows the excess power in fluctuations as a function of angular scale. Each $\ell = \frac{180}{\lambda}$ roughly shows the inverse of the angular size of fluctuations, and the C_ℓ 's are the amplitude of each multipole moment of the correlation function of temperature fluctuations $\langle TT \rangle$ from a spherical harmonic expansion (approximately, this is like a Fourier expansion but projected onto a 2-dimensional sphere). The upper panel shows the impact of varying N_{eff} on this spectrum, while the lower panel shows the impact of varying $\sum m_\nu$ (in eV, while keeping the ratio of baryons to dark matter fixed as the total matter density varies).

Additionally, BBN (Big Bang Nucleosynthesis) is able to assist in constraints for N_{eff} . This is because altering N_{eff} impacts the primordial abundance of helium and deuterium; if the expansion

rate is increased by increasing N_{eff} , the amount of time and number of photons that are hot enough to destroy deuterium is decreased, increasing the abundance of deuterium and allowing more helium-4 to form. BBN constraints alone by Mangano and Serpico (2011) determine $\Delta N_{\text{eff}} \leq 1$ at the 95% confidence level (relative to the Standard Model expectation). Results from the Planck CMB measurements combined with measurements of the primordial abundances (Aghanim et al., 2020) find $N_{\text{eff}} = 3.06 \pm 0.22$ (2σ constraints), while a combined constraint from primordial abundances with cosmological data (including Planck data, SNe Ia, BAOs and more) by Stöcker et al. (2021) finds $N_{\text{eff}} = 3.06 \pm 0.24$ ($N_{\text{eff}} = 3.08 \pm 0.24$) for a normal mass hierarchy (inverted mass hierarchy). Riemer-Sørensen et al. (2013) presents a thorough review of constraints on N_{eff} from data including primordial abundances, CMB measurements, and these measurements combined with information from large-scale structure such as BAO data, the matter power spectrum, SNe Ia, and even the Lyman- α forest.

Finally, information in the phase shift in BAO measurements has some ability to constrain N_{eff} . Mainly, constraints by Follin et al. (2015); Baumann et al. (2019); Whitford et al. (2025); Montefalcone et al. (2025c) for the phase shift, and with additional information from CMB measurements by Planck, show the presence of a phase shift caused by free-streaming relics is present in the CMB and BAOs. At present, these measurements cannot provide such constraints on N_{eff} , but do give a robust measurement of a signature of free-streaming relics with little degeneracy with other cosmological parameters.

3.2.3 Massive neutrinos in cosmology

Massive neutrinos have two major effects on large-scale structure probes that can be measured. Firstly, as relativistic particles neutrinos contributed to the total density of radiation in the early Universe, Ω_r , as $\Omega_r = \Omega_\nu + \Omega_\gamma$. As such their presence contributed to the expansion rate of the Universe in the radiation-domination epoch. Their masses dictate the point in time at which they switch from being a relativistic particle that contributes to Ω_r , in which their kinetic energy is on the order of their rest mass energy, to behaving like massive particles. Neutrinos lose kinetic energy as their peculiar velocities decay relative to the expanding Universe. As massive particles, their energy density contributes to the total energy density of matter Ω_m such that $\Omega_m = \Omega_{\text{cdm}} + \Omega_b + \Omega_\nu$, the sum of densities of cold dark matter (cdm), baryons and neutrinos. Consequently neutrinos influence the epoch of *matter-radiation equality*, at which $\Omega_m(t) = \Omega_r(t)$. As mentioned previously this epoch is related to the amount of time that matter perturbations have had to grow in the Universe until today and the length scales that had their growth suppressed in the radiation-domination era.

The second major signature of neutrinos is due to their free-streaming effects. It was discussed earlier that neutrinos interact only weakly and consequently they stopped interacting with other particles in the hot thermal plasma almost immediately after the Big Bang and began to fly freely. Neutrinos initially moved with a velocity very close to the speed of light due to their high thermal energy, but as massive particles later began to slow down relative to the expanding Universe and started clustering with matter. On length scales smaller than the distance over which they could free-stream

before slowing, the growth of matter perturbations is gravitationally suppressed in a way that depends on their masses (Lesgourgues and Pastor, 2014; Riemer-Sørensen et al., 2013). The scale at which neutrinos free-stream (when they still have a peculiar velocity) before clustering with matter at some time corresponding to $z(t)$ is given by (Lesgourgues and Pastor, 2006)

$$k_{\text{FS}} \approx 0.82 \frac{\sqrt{\Omega_\Lambda + \Omega_m(1+z)^3}}{(1+z)^2} h\text{Mpc}^{-1}. \quad (3.51)$$

This equation is valid for the matter-domination epoch, and this scale is derived from a calculation that is analogous to the definition of the Jeans length (Lesgourgues and Pastor, 2006); this involves comparing length scales in which the gravitational potential energy dominates compared to their kinetic energy, given by their thermal velocity. The minimum k (or maximum scale) for which neutrinos can free-stream is determined by the time at which the expansion rate of the Universe surpasses the rate at which the free-streaming scale can continue to grow, and is given by

$$k_{\text{nr}} \approx 0.18 \Omega_m^{1/2} \left(\frac{m}{\text{eV}} \right) h\text{Mpc}^{-1}. \quad (3.52)$$

m is the mass of the neutrino. Figure 3.2 shows the impact of massive neutrinos on the matter power spectrum, which is suppressed on small scales for greater neutrino masses. Figure 3.6 also shows the impact of massive neutrinos on the CMB temperature-temperature power spectrum.

Existing constraints on massive neutrinos

Data from galaxy surveys or the CMB can constrain $\sum m_\nu$. Due to the impact of massive neutrinos on both the expansion rate and damping on small scales of either the matter power spectrum or the power spectrum of temperature fluctuations in the CMB, cosmological probes place an upper bound on $\sum m_\nu$. In principle, three neutrinos of different species and thus different masses could also be constrained as they begin to free-stream on different scales. However, in Lesgourgues and Pastor (2014) it has been pointed out that due to cosmic variance, it will be unlikely for it to be possible for cosmological data to have enough constraining power to constrain each individual neutrino mass. At present, some of the tightest existing constraints, $\sum m_\nu < 0.072$ eV at the 95% confidence level, come from the DESI BAOs (DR1) + Planck measurements of the CMB (Aghanim et al., 2020; Adame et al., 2025b). The most recent constraints from DESI DR2 (Karim et al., 2025) with Planck give $\sum m_\nu < 0.064$ eV at the same confidence level, although this relaxes to $\sum m_\nu < 0.079$ eV with the newer Planck analysis (Tristram et al., 2024). The constraints are much weaker when w_0 and w_a are allowed to vary in time-varying dark energy models (Elbers et al., 2025; Karim et al., 2025). It should be noted here that the BAO information from DESI helps to improve constraints on parameters that determine the expansion rate but the BAOs themselves do not constrain $\sum m_\nu$ alone; this constraint is predominately driven by the Planck data. Figure 3.6 shows the impact of varying $\sum m_\nu$ on the CMB.

Upper bounds on neutrino mass are possible from experiments that measure the neutrino *mass-splittings*. In the Standard Model neutrinos are massless particles. However, neutrino oscillations, in which neutrinos switch between different species as they propagate through space, led to the realization

that neutrinos must have non-zero masses (Fukuda et al., 1998). The mass-splittings refer to the difference in mass Δm_{ij}^2 between each of the neutrino mass eigenstates (these are superposition states of the neutrino species μ , τ , e). The mass-splittings can be measured from oscillation experiments, because the number of neutrinos that oscillate between different flavour types depends on their mass differences as well as the oscillation baseline. Based on data from experiments such as Super-Kamiokande that have measured the mass-splittings from solar, reactor and atmospheric experiments (see references within Stöcker et al., 2021), there are two possible ways in which the measured splittings could be ordered, known as the normal hierarchy in which $\sum m_\nu \geq 0.058$ eV, or the inverted hierarchy in which $\sum m_\nu \geq 0.098$ eV. The two possibilities are due to the fact that it is not generally possible to know the sign of Δm_{ij} from each experiment. We can expect that if cosmological data can reduce the upper bounds on $\sum m_\nu$ sufficiently, it may be possible to determine which hierarchy is correct and determine the individual mass of each neutrino mass-eigenstate.

It is possible to conduct global fits to data from cosmology and neutrino oscillations to constrain neutrino masses, such as the results in Stöcker et al. (2021). Information from cosmological data such as SNe Ia, which constrain the expansion rate, or abundances of primordial elements (Big Bang Nucleosynthesis) which constrain degenerate parameters to $\sum m_\nu$ (such as N_{eff}) are included to improve the constraints. Additionally, data from the Lyman- α forest can improve constraints, by providing further constraining power on the underlying matter distribution; on its own, it can constrain $\sum m_\nu$ at around the ~ 1 eV level (see Yèche et al., 2017, for example). Via the impacts on the matter power spectrum it is possible for redshift space distortions or measurements of the full shape matter power spectrum to add power to these constraints. Particle physics experiments that observe beta decay (in which a neutron decays to a proton, electron and anti-neutrino) or search for a hypothetical neutrino-less double beta decay process (in which two neutrons decay to two protons and two electrons) can also constrain $\sum m_\nu$ (Tanabashi et al., 2018; Zyla et al., 2020).⁹ While these constraints are much less model-dependent than cosmological constraints, they are currently much less competitive.

3.2.4 Neutrino self-interactions in cosmology

In Chapter 2 I discussed some motivation for looking for signatures of neutrino physics beyond the Standard Model, which I review here in more detail. In regards to cosmological data, in the past there has been some motivation to search for this because if neutrinos have non-standard interactions, this can potentially delay neutrino free-streaming and shift the scale r_s of the sound horizon at the epoch of recombination. As mentioned before, neutrino free-streaming introduces a characteristic phase shift in the CMB and BAO oscillations. If free-streaming is delayed due to neutrinos continuing to self-interact, they decouple later than expected for Standard Model neutrinos. This can introduce a scale-dependent change to the free-streaming signal relative to the Standard Model expectation, and shift the positions of the peaks, which constrain r_s . Therefore, neutrino self-interactions allow for flexibility in constraints on H_0 .

⁹Essentially, two beta decay processes occur and there are no neutrinos in the final decay products because they are emitted and absorbed in the process, if neutrinos are their own anti-particles.

Neutrino self-interactions also alter the amplitude of the CMB and introduce a peak-like feature in the matter power spectrum (see this also in Figure 6.1 and Figure 6.2). The description I include for this here is largely inspired by the text in section 4 of Kreisch et al. (2020) and section B of Camarena et al. (2023a). The change in the CMB amplitude, which is increased relative to Λ CDM for stronger neutrino self-interactions, is related to the fact that free-streaming neutrinos in Λ CDM reduce the gravitational potential because they carry away energy. However, if they continue interacting for longer they can contribute to the radiation density and the stress-energy tensor. In the matter power spectrum, this results in a decreased amplitude for structure growth for small scales that do not enter the causal horizon while the neutrinos are still self-interacting, as the relativistic neutrinos are ‘taking away’ from the matter perturbations by increasing the gravitational potential. However, as the neutrinos begin to decouple (corresponding to larger scales), they allow the gravitational potential to decay and boost the amplitude of matter fluctuations relative to Λ CDM, in which the neutrinos have already begun free-streaming. Eventually when the neutrinos begin to free-stream, the behaviour matches that of Λ CDM (largest scales). Hence, this results in a peak in the power spectrum. The changes to the gravitational potential and thus scale-dependent changes in the power spectrum amplitude and CMB amplitude allow for variation in constraints or inferences on σ_8 or Ω_m . Neutrino self-interactions may also be able to relax constraints on models of inflation (Barenboim et al., 2019).

As I also mentioned in Chapter 2, these effects on cosmological data are well captured by studying simplistic models with universal self-interactions. A mediator particle ϕ allows for the neutrino self-interactions, and the effects are captured by a single effective parameter $G_{\text{eff}} = \frac{|g_\nu|^2}{m_\phi^2}$. Studying these impacts in cosmological data may also complement experiments, that search for non-standard neutrino physics. In theories where neutrinos gain mass from a new mediator particle in which lepton-symmetry is spontaneously broken, this allows for neutrinos to also have self-interactions (Berryman et al., 2023). These theories require the neutrinos to be their own anti-particles (majorana). Berryman et al. (2023) also describes how special classes of ‘gauge boson’ (force-carrier) extensions to the Standard Model can induce new interactions between neutrinos. These models can lead to interaction with charged Standard Model fermions, allowing for constraints from more experiments. Self-interactions can also potentially accommodate interactions between Standard Model particles and dark matter candidates, or even a ‘secret’ interaction with a sterile (a non weakly-interacting) neutrino; sterile neutrinos can help to solve anomalies in oscillation experiments or even play as a role as a dark matter candidate (Berryman et al., 2023).

Existing constraints on self-interactions

I described in Chapter 2 how constraints from the CMB and the matter power spectrum suffers from bimodality (with a moderate and strong interaction mode, MIV and SIV respectively). The persistent bimodality in constraints can be seen in various works, however it is possible to determine if the data in the constraints prefers one model to the other. In Kreisch et al. (2020), it is found from the Planck 2015 TT power spectrum + Planck 2015 lensing + combined BAO measurements (eBOSS, 6dFGS, CMASS, SDSS) and a local H_0 measurement (from cepheids) that strong self-interactions are mildly

preferred to Λ CDM in the fit to the data. Allowing N_{eff} and $\sum m_\nu$ to vary in the fit finds larger values of N_{eff} than expected in the Standard Model, and a peak in $\sum m_\nu$ at positive values. In Kreisch et al. (2024), a similar analysis, but applied to data from Planck (2018) and ACT CMB data found that while ACT data also prefers strong interactions, data from Planck 2018 prefers it much less; the polarization (EE) data from ACT drives the preference for the SIV mode. Meanwhile, Camarena et al. (2023a) found a modest preference for the SIV mode in data from BBN (priors for ω_b, N_{eff}) with the full shape of the matter power spectrum measured from the 12th data release of BOSS, suggesting the presence of the SIV independently in LSS and CMB data.¹⁰ However, a combined analysis on Planck 2018 + LSS data by Camarena and Cyr-Racine (2025) disfavors the self-interaction model compared to Λ CDM + N_{eff} + $\sum m_\nu$. They find the failure of the self-interaction model to accommodate both datasets may be driven by polarization data from Planck, that cannot accommodate the lower value required for A_s required to compensate for strong self-interactions. Despite this, there is still the presence of a peak in the profile-likelihood for G_{eff} that corresponds to the SIV mode. Interestingly, an analysis on the new Planck analysis is found by Poudou et al. (2025) to be more accommodating to the SIV mode. The addition of DESI BAO data to the analysis allows for a preference of the SIV model over Λ CDM. They also analyse LSS (power spectrum) data, but find overall the LSS data does not prefer nor disfavour self-interactions over Λ CDM. He et al. (2025) analyses the Planck data with Lyman- α Forest data from eBOSS and finds negligible evidence for self-interactions, and from an analysis of Planck CMB data + DESI BAOs finds a similar result; their result may differ from Poudou et al. (2025) due to the use of a different likelihood for the analysis. In Barenboim et al. (2019), constraints on G_{eff} from an analysis of data from Planck 2015 are converted into constraints on the $m_\phi - g_\nu$ plane, which can be seen in Figure 2 of that work. In the work of Brinckmann et al. (2021), a comprehensive analysis of Planck 2018 CMB data + BAO data (6dFGS, SDSS, BOSS), allowing N_{eff} to vary and allowing for self-interactions, does not find compelling evidence of self-interactions over Λ CDM.

In contrast to the models discussed here, it is possible for cosmological data to provide constraints on models that do not suffer from such bimodality. Noriega et al. (2025) provides bounds on ‘resonant’ neutrino self-interactions, although the bounds are dependent on the mediator mass m_ϕ which are chosen for a range of values in the analysis. They find no evidence for self-interactions in their tests. Additionally, there are constraints from CMB data on scenarios where neutrinos ‘recouple’ to have self-interactions, although Berryman et al. (2023) notes that these constraints are model-dependent in how the self-interaction strength relates to the redshift at which the neutrinos recouple (see references for these constraints within Berryman et al., 2023). For the constraints that have been discussed so far, universal self-interaction models have been used; however the preference for self-interactions is found to be stronger in CMB data when only a single neutrino species has self-interactions (Das and Ghosh, 2021), in line with constraints from laboratory probes that find that self-interactions should not be universal. As a final comment regarding constraints from cosmological data, BBN allows for constraints to be derived on self-interactions via the alterations that such interactions would have on

¹⁰BBN priors can be safely used in the presence of self-interactions, given the strong-self interactions studied in this work only mildly impact BBN predictions for the abundance of light elements.

the primordial abundances of light elements, although they are model dependent (constraints have been derived in Huang et al., 2018). In general, changes to BBN through self-interactions does not significantly change cosmological parameters in a manner that allows for strong constraints (Berryman et al., 2023).

Various particle physics experiments also allow for constraints. Searches for hypothetical double beta decay can allow for constraints because neutrino self-interactions would leave a distinct signatures in the spectrum (Deppisch et al., 2020). Additionally, searches for rare-meson (particle made of two quarks) decays may involve new products that could be observed if self-interactions occur, allowing for bounds on the self-interaction strength. In some classes of models, the particle that mediates self-interactions can be searched for in particle colliders. Self-interactions can also manifest in scattering processes that may result in final state particles with the wrong ‘sign’ for charge or ‘missing’ momentum, allowing for constraints (details of these constraints are also discussed in the comprehensive review of Berryman et al., 2023).

Finally, it is also possible to use astrophysical detections of neutrinos from supernovae or high-energy extragalactic neutrinos to derive constraints. In supernovae, neutrinos can have significant impacts since a large fraction of the energy in a supernova explosion is transferred to neutrinos. Constraints have been derived on neutrino properties using ~ 10 seconds of detected neutrinos from SNe 1987A (Kolb and Turner, 1987). Self-interactions can impact the neutrino energy spectra from supernovae, change the flavour distribution of neutrinos, or even change the SNe explosion mechanism. Additionally, the neutrinos may scatter as they propagate through space and interact with the CνB, leading to further spectral distortions. Neutrino telescopes, such as IceCube,¹¹ allow constraints by also searching for changes to the spectra of high-energy astrophysical neutrinos from self-interactions (Esteban et al., 2021). I show a plot taken directly from Berryman et al. (2023) that summarizes many of the existing constraints on self-interactions in the $m_\phi - g_\nu$ plain, which have been discussed in this chapter for BBN and CMB data, as well as rare meson decays, double beta decay, and scattering measurements from laboratory probes.

Summary

With this chapter, we have thoroughly introduced the relevant theory and current status of large-scale structure cosmology. In addition, the background here gives some introduction to my research on bulk flow measurements in Chapter 4, and measurements of signatures of free-streaming neutrinos in BAO data in Chapter 5 and Chapter 6.

¹¹<https://icecube.wisc.edu/science/icecube/>




Figure 3.7: A summary plot of constraints in the $m_\phi - g_\nu$ plane for neutrino self-interactions, taken directly from Berryman et al. (2023), who themselves have created this figure by reproducing the contours from results of various other works. The dashed and solid lines show projected constraints from future experiments and detectors, while the solid regions show various existing constraints that have been discussed in this chapter.

The following publication has been incorporated as Chapter 4. I have changed the use of ‘we’ in the original publication text to ‘I’ for this chapter.

1. Whitford et al. (2023) **Abbé Whitford**, Cullan Howlett and Tamara M. Davis Evaluating bulk flow estimates for CosmicFlows-4 Measurements, *Monthly Notices of the Royal Astronomical Society* 526, 2, 2023.

Contributor	Statement of contribution	%
Abbé M. Whitford	writing of text	90
	proof-reading	50
	numerical calculations	100
	theoretical derivations	100
	preparation of figures	100
	initial concept	5
Cullan Howlett	writing of text	5
	proof-reading	25
	supervision, guidance	70
	initial concept	95
Tamara M. Davis	writing of text	5
	proof-reading	25
	supervision, guidance	30

This project involving application of bulk flow estimators to data from the CosmicFlows-4 PV catalogue was developed by Cullan. As the project progressed we conducted exploratory analysis that involved ideas and input from myself, involving ideas for tests on simulation data. I created various simplistic mocks (simulations) with my own code, but created more sophisticated mocks with existing codes provided by Cullan. Additionally, the most realistic mock datasets we used for the CosmicFlows-4 data were provided by Cullan and Fei Qin, which I ‘stitched’ together as mocks representative for the entire CosmicFlows-4 catalogue. The mocks of the SDSS data can be found at <https://zenodo.org/record/6640513>. I also used some code shared with me by Anthony Carr for this project. I also thank Eoin Ó Colgáin, Richard Watkins, Hume Feldman and the anonymous reviewer of this publication for their useful comments and thoughts on the first arXiv manuscript. All tests and results from mock data included in this chapter were a result of my own calculations. All measurements of the peculiar velocity data, which were not at all conducted by myself, used in our calculations has been taken from a publicly available compilation that is available at the Extragalactic Distance Database.

Further evidence, plots, code and information for the various tests and results presented and discussed in this chapter can be found here: https://github.com/abbew25/Measuring_bulkflows. This research has made use of NASA’s Astrophysics Data System bibliographic services, the astro-ph pre-print archive at <https://arxiv.org/> and the python libraries MATPLOTLIB, ASTROPY (Hunter, 2007; Collaboration et al., 2013). I made use of the University of Queensland’s Getafix, Tinaroo and Bunya high performance computers to conduct this research.

Chapter 4

Evaluating bulk flow estimators for CosmicFlows-4 measurements

Recall from Chapter 2 our discussion that for over a decade there have been contradictory claims in the literature about whether the local bulk flow motion of galaxies is consistent or in tension with the Λ CDM model. It has become evident that systematics do affect bulk flow measurements, but systematics in the estimators have not been widely investigated. In this chapter, I thoroughly evaluate the performance of four estimator variants, including the Kaiser maximum likelihood estimator (MLE) and the minimum variance estimator (MVE). I find that these estimators are unbiased, however their precision may be strongly correlated with the survey geometry. Small biases in the estimators can be present leading to underestimated bulk flows, which are suspected to be due to the presence of non-linear peculiar velocities. The uncertainty assigned to the bulk flows from these estimators is typically underestimated, which leads to an overestimate of the tension with Λ CDM. I estimate the bulk flow for the CosmicFlows-4 data and use mocks to ensure the uncertainties are appropriately accounted for. Using the MLE we find a bulk flow amplitude of $408 \pm 165 \text{ km s}^{-1}$ at a depth of $49 \text{ Mpc } h^{-1}$, in reasonable agreement with Λ CDM. However, using the MVE which can probe greater effective depths, I find an amplitude of $428 \pm 108 \text{ km s}^{-1}$ at a depth of $173 \text{ Mpc } h^{-1}$, in tension with the model, having only a 0.11% probability of obtaining a larger χ^2 . These measurements appear directed towards the Great Attractor region where more data may be needed to resolve tensions.

4.1 Overview

Local fluctuations in the underlying matter density of the Universe source the gravitational motions of galaxies. These motions create local velocity flows that drag galaxies towards each other. I mentioned in Chapter 3, that the measurement of the average of these motions is called the *bulk flow*. The bulk flow in a particular volume gives us a picture of the direction and amplitude of the overall flow of matter in that region. Because the bulk flow arises due to the large-scale structure (LSS), it is not only a useful tool that can be used to map motions in the local Universe but it also allows us to test models of cosmology. The bulk flow is related to Ω_m , the total matter energy density due to baryons and dark matter, and σ_8 , the variance in matter fluctuations in spheres of radius 8 Mpc h^{-1} . We can also use bulk flow measurements to test theories of gravity, such as General Relativity.

The bulk flow is particularly interesting at the present time due to tensions in bulk flow measurements over the last decade or more. A number of measurements have claimed the measured bulk flows are in some tension with the current concordance model of cosmology, Λ CDM (Kashlinsky et al., 2008; Watkins et al., 2009; Feldman et al., 2010; Watkins and Feldman, 2015b; Peery et al., 2018; Howlett et al., 2022; Watkins et al., 2023). In contrast, a number of measurements have found they are in agreement (Nusser and Davis, 2011; Hong et al., 2014; Ma and Pan, 2014; Hoffman et al., 2015; Scrimgeour et al., 2016; Qin et al., 2018, 2021). Interestingly, many of the measurements that are in tension with the Λ CDM model measure the bulk flow on larger scales than those that appear in agreement. Given these disagreements, the aim here is to make the most robust measurement yet of the bulk flow using the latest CosmicFlows–4 catalogue, while also considering the precision and accuracy of the bulk flow estimators we employ in order to properly determine the consistency of the measurement with the Λ CDM model. In particular, it is interesting to test if additional peculiar velocity data resolves or worsens tensions in bulk flow measurements.

Formally, I mentioned earlier the bulk flow \mathbf{B} is a measurement of the average peculiar velocity of galaxies in a given survey volume V ,

$$\mathbf{B} = \frac{1}{V} \int_V \mathbf{v}(\mathbf{r}) d^3r, \quad (4.1)$$

where as before $\mathbf{v}(\mathbf{r})$ represents the 3–dimensional peculiar velocity (PV) field at coordinate \mathbf{r} , which refers to the motions due to local gravitational interactions rather than due to the expansion of space. I also discussed earlier how only the radial component of the PV of a galaxy can be measured and therefore the components of \mathbf{B} along each coordinate axis in a volume, B_i , may be estimated as a weighted average of radial PV measurements (stated with Equation 3.39). I will restate this here,

$$B_i \approx \frac{1}{V} \int u(\mathbf{r}) \hat{n}_i w_i(\mathbf{r}) d^3r \approx \sum_m u_m \hat{n}_{i,m} w_{i,m} = \tilde{B}_i. \quad (4.2)$$

In the above equation, \tilde{B}_i is the estimator for B_i , where $\mathbf{B} = \sum_i B_i \hat{\mathbf{x}}_i$. The radial PV field we measure is given by $u(r) = \mathbf{v}(\mathbf{r}) \cdot \hat{\mathbf{r}}$ where $\hat{\mathbf{r}}$ is the unit vector for the line-of-sight. Above, u_m represents the observed PV for the m th galaxy in the sample, $w_{i,m}$ represents the weight applied to the m th galaxy for the i –coordinate direction (given some estimator used to determine the optimal weighting scheme) and $\hat{n}_{i,m}$ gives the projection of the radial PV onto the i coordinate direction, $\hat{n}_{i,m} = \hat{\mathbf{r}}_m \cdot \hat{\mathbf{x}}_i$.

Alternatively, one may attempt to reconstruct the full 3–dimensional PV field (which we also discussed relating to Equation 3.40), using a method such as the Wiener Filter (Zaroubi et al., 1995), then sum the 3–dimensional PVs to obtain a bulk flow estimate. An issue that may arise with this method is the need for prior information or model assumptions about the underlying cosmology of the Universe to do the reconstruction. This may dominate the underlying bulk flow signal if the data used to do the reconstruction is noisy. In general however, different methods of measuring the bulk flow may obtain results that are not directly comparable because they define the true bulk flow moment of the data differently (Nusser, 2016). Assumptions in bulk flow estimators needs to be taken carefully into consideration, along with potential systematics, before comparing a bulk flow estimate to theory.

As stated previously the aim is to make a measurement of the bulk flow, primarily using the CosmicFlows–4 catalogue, the largest combined dataset of PVs to date (Tully et al., 2023). This dataset consists of combined data from the Sloan Digital Sky Survey peculiar velocity catalogue (SDSS; Howlett et al. 2022), the CosmicFlows IV Tully–Fisher catalogue (CF4TF; Kourkchi et al. 2020) and the 6-Degree Field Galaxy peculiar velocity survey (6dFGSv; Springob et al. 2014). The dataset also contains PVs measured from other low-redshift objects that can be used to construct a distance ladder and together make a total of 55,877 measurements. This combined sample will also probe more deeply than other combined samples; this is mainly due to the new SDSS sample which contains peculiar velocities for galaxies as deep as $z \sim 0.1$. However, the peculiar ‘top-heavy’ geometry also offers a potential route for systematics to enter. A discussion for the datasets and mocks that I test the methodologies on can be found in section 4.2. I also present results for measurements from the individual SDSS, 6dFGSv and CF4TF datasets, as well as their combination when other low redshift objects are not included.

Prior to this in section 4.3 I discuss systematics in estimators for the bulk flow and the results of our own tests of the performance of popular bulk flow estimator methodologies. I apply these estimators to a range of mocks (ranging from simplistic to sophisticated in nature) for the datasets of interest in order to evaluate their performance thoroughly in section 4.4. The aim here is to closely investigate systematics in the Kaiser Maximum Likelihood Estimator (Kaiser MLE; Kaiser, 1988) and the Minimum Variance Estimator by Watkins et al. (2009) (Watkins MVE). I also investigate the Maximum Likelihood approach suggested in Nusser (2014) (hereon Nusser MLE) and the modifications to the Minimum Variance Estimator in Peery et al. (2018) (hereon Peery MVE). I discuss the pros and cons of these methods, and which methods one can expect are best when applied to survey data. This is an important sanity-check for solving tensions in bulk flow measurements, especially in light of the recent results by Watkins et al. (2023) who claim to make a bulk flow measurement also using CosmicFlows–4, that is in significant tension with the Λ CDM model. Our tests on mock data show that the application of the MVE and the Kaiser MLE to the CosmicFlows–4 data tend to obtain a slightly underestimated measurement of the bulk flow (which increases the tension), however our results also show that the uncertainty on the measurement is also likely underestimated, which overestimates the amount of tension between Λ CDM with the measurement. Furthermore, I explore the effects of the zero–point calibrations applied to the individual datasets that compose CosmicFlows–4 to show how this changes

the measured bulk flow amplitude.

My results from applying the bulk flow measurement techniques to realistic mocks and the various datasets are shown in section 4.5, along with a comparison between the results from the data and mocks with theoretical expectations to quantify the level of tension with Λ CDM. Finally in section 4.6 I conclude with a discussion of recommendations for future work regarding bulk flow measurements and potential ways to improve existing methods of measuring the bulk flow to resolve tensions between datasets and robustly test the Λ CDM model.

4.2 Datasets and mock catalogues

Properties of the CosmicFlows–4 (CF4) data and the largest subsets composing it (SDSS, 6dFGSv and CF4TF) are discussed in this section and summarised in Table 4.1, Figure 4.1 and Figure 4.2. The mocks used for the CF4 data, by combining mocks for SDSS, 6dFGSv and the CF4TF datasets and which capture the majority of the data in CF4 and the survey geometry, are also discussed in this section.

For all the CF4 mocks and data, peculiar velocities are estimated from log-distance ratios using the Watkins and Feldman (2015a) estimator. Furthermore, all galaxy distances calculated for determining weights in any of the estimators or the effective depth of the data use the observed redshifts of galaxies.

4.2.1 The Sloan Digital Sky survey peculiar velocity catalogue

The data

The Sloan Digital Sky Survey (SDSS) peculiar velocity catalogue (Howlett et al., 2022) is a set of 34,059 peculiar velocities on the Fundamental Plane (FP, Djorgovski and Davis, 1987), comprising the largest individual set of peculiar velocities to date. This dataset probes more deeply than the other samples we combine with to a redshift of $z = 0.1$ but only covers only a small region of the sky in the SDSS northern galactic cap. The mean uncertainty on each FP distance measurement is $\sim 23\%$. This dataset contains group redshifts for galaxies from the same underlying dark matter haloes that allow for the galaxy PVs to be measured more accurately. Peculiar velocities are obtained by building on Said et al. (2020) and fitting the Fundamental Plane parameters extracted from the SDSS Data Release 14 (Abolfathi et al., 2018).

The mocks

Howlett et al. (2022) provide a set of 2048 realistic mocks for this dataset. The mocks reproduce the sky mask and selection function of the SDSS PV data and were designed to capture all aspects of the real data, including selection effects, measurement errors and the effects of cosmic variance. More details can be seen in Howlett et al. (2022), and the mocks can be found at <https://zenodo.org/record/6640513>.

4.2.2 The CosmicFlows IV Tully–Fisher peculiar velocity catalogue

The data

The CosmicFlows IV Tully–Fisher catalogue (CF4TF, Kourkchi et al., 2020) consists of 9534 peculiar velocities drawn from the Tully–Fisher plane relation (TF, Tully and Fisher, 1977) that are mostly within a redshift of $z \sim 0.05$ and cover the sky entirely (excluding the galactic plane). Unlike SDSS, this dataset is shallower, but it is apparent in Figure 4.1 that the selection function of this dataset is complimentary to the SDSS dataset by having a high density of objects at low redshift. For this dataset, the H I linewidths and fluxes used for the TF relation parameters come from measurements by the All Digital H I catalogue (ADHI, 78%, Courtois et al., 2009), the Aricebo Fast ALFA Survey (Haynes et al., 2011, 2018), the Springob/Cornell H I catalogue (Springob et al., 2005) and the Pre Digital H I catalogue (Fisher and Tully, 1981; Huchtmeier and Richter, 1989). Photometry data for the galaxies is taken from the SDSS DR12 data release (York et al., 2000). More information can be found within Kourkchi et al. (2020).

The mocks

For this dataset mocks have been produced by Qin et al. (2021), which reproduce the survey geometry and selection function for the CF4TF data. The mock sampling algorithm used here first by Qin et al. (2021) to generate the mocks is the same as that used to produce the SDSS mocks by Howlett et al. (2022) with the corresponding survey properties. As such, the mocks for the CF4TF data are also designed to capture all the aspects of the real data including cosmic variance, measurement errors and selection effects.

4.2.3 The 6-degree Field Galaxy survey PV catalogue

The data

This 6dFGSv dataset (Springob et al., 2014) consists of another 8885 galaxies with peculiar velocities measured from the FP relation. This sample covers the entire Southern sky (excluding the plane of the Milky Way) and is thus complimentary to the SDSS and CF4TF data. This dataset is deeper than CF4TF with more redshifts at $z \sim 0.03 - 0.056$ but is shallower than the SDSS sample and has greater sky coverage. The typical uncertainties on each galaxy distance from this FP are around $\sim 26\%$. The spectroscopic observations of the FP galaxies in this sample were made with the UK Schmidt Telescope as part of the 6dF survey (Jones et al., 2009), and photometric observations from the Two Micron All-Sky Survey (2MASS) Extended Source Catalog (Jarrett et al., 2000). The data used here (and within the full CF4 release) is a reprocessing of the original data including a modified correction for the selection function when extracting distances from the Fundamental Plane. The reprocessing affects only the extraction of distances/velocities, not the publicly available photometric or spectroscopic measurements, and is fully described in Qin et al. (2018) and Tully et al. (2023).

The mocks

Mocks for the 6dFGSv survey are created following Qin et al. (2019), although with slight modifications as presented in Tully et al. (2023) and to align their construction with the SDSS and CF4TF methodology used to produce the mock catalogues in Howlett et al. (2022) and Qin et al. (2021). Again, these reproduce the selection function of the survey and are designed to capture the effects of cosmic variance, measurement errors and selection effects as closely as possible to the real survey.

Survey	No. galaxies	z range (median)	mean error
SDSS	34059	0.0033–0.1 (0.071)	23%
CF4TF	9534	0.0–0.064 (0.018)	22%
6dFGS	8885	0.0–0.056 (0.039)	26%
CF4	55877	0.0–0.1 (0.051)	21.5%

Table 4.1: Properties of each peculiar velocity dataset used in this work, where z refers to the CMB-frame redshift and the mean error column refers to the mean error on each measured galaxy distance.

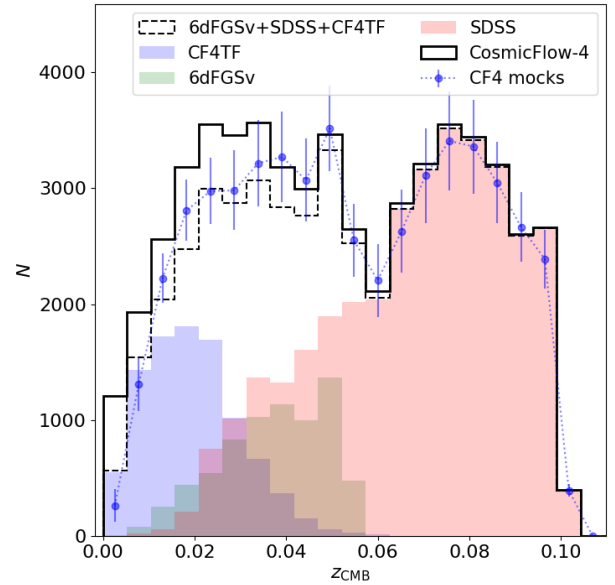


Figure 4.1: Individual selection functions for each dataset, their combination, and CF4. The selection function for CF4 includes low- z objects. z_{CMB} represents the redshift in the CMB frame. Each bin has a width of $\Delta z = 0.005$.

4.2.4 The combined CF4 dataset and mocks

The data

Figure 4.1 shows the redshift selection function for each dataset discussed in the previous section, the redshift selection function for the combined data from these datasets that is included in CF4 (this is not all of the 8885 galaxies for 6dFGSv because some of these are removed in the catalogue due to their classification as spiral interlopers in Tully et al., 2016) and the redshift selection function for the entire CosmicFlows–4 catalogue which includes low-redshift objects used to calibrate the zero-point for the data in Tully et al. (2023).

Figure 4.2 shows the sky coverage for the entire CosmicFlows–4 dataset. The distance measurements in CF4 that are not included in SDSS, CF4TF and 6dFGSv include those from Cepheid variable stars (Leavitt and Pickering, 1912), type Ia SNe (Phillips, 1993), type II SNe (Hamuy and Pinto, 2002), surface brightness fluctuations (SBFs) in elliptical galaxies (Tonry and Schneider, 1988) and tip of the red-giant branch (TRGBs; Lee et al., 1993). Further FP and TF measurements are also used




Figure 4.2: The sky coverage for the entire CosmicFlows–4 dataset shown in galactic coordinates. Each point here represents a group out of the 38008 groups of objects made from the 55877 objects. The color of each point shows the redshift relative to the cosmic microwave background (CMB) reference frame according to the colorbar.

from other samples detailed in Tully et al. (2023). These other samples have overlapping galaxies with the SDSS, CF4TF and 6dFGSv datasets already discussed and independent distance measurements for the same galaxy have been averaged in the combined dataset for all CF4.

Parallax distance measurements and geometric maser distances (Humphreys et al., 2013) help set the absolute distance scale for TRGBs and Cepheid-variable stars. These are used to help provide an absolute distance scale and allow for the zero–point of the datasets to be calibrated using a Bayesian methodology. The zero–point offset of the galaxy peculiar velocities represents the monopole of the field and is degenerate with H_0 . The need for a correction arises during the FP (Fundamental Plane) and TF (Tully–Fisher) fitting procedures, as it is assumed when measuring these relations that the zero–point of the galaxy PVs is null, which causes an unknown constant offset in the true PVs of the galaxies relative to the measurements that requires correction. In Tully et al. (2023), the overlapping objects of various datasets for FP, TF, SNe Ia, SNe II, and TRGB stars are used to calibrate datasets for galaxy distances measured using the same methodology to each other. Then all of the datasets are tied to an absolute scale set by the SNe Ia objects using overlapping galaxy groups in the samples. More detail can be found in Tully et al. (2023). The process for calibrating the various zero–points carries potential for systematic error, which is one of the things I test in this work.

The mocks

The mocks from Qin et al. (2019, 2021); Howlett et al. (2022) discussed in the previous sections for SDSS, 6dFGSv and CF4TF have been all made using the L-PICOLA n -body code (Howlett et al., 2015) with the same settings and initial conditions for each mock, such that all the galaxies are assigned

to haloes from the same dark matter distribution. They have also been placed in the simulation in such a way that the positions of galaxies relative to the observer mimics the relative positions of the true galaxy surveys; this allows the survey mocks to be stacked together in order to make combined mocks for the datasets and also largely capture the geometry of CF4. I present analysis of the bulk flow estimators applied to the mocks in section 4.5. The blue points in Figure 4.1 shows the averaged redshift selection function of the CF4 mocks used, and the error bars show the standard deviation of the number counts in each bin for the mocks.

4.3 Systematics in bulk flow measurements

Various difficulties exist in regards to producing accurate PV measurements and subsequently a bulk flow measurement. Firstly, measurements of PVs derived from distances obtained with empirical scaling relations, such as the FP or the TF relation, are noisy and have an uncertainty on the measured galaxy distances that are typically $\sim 20\%$ (Strauss and Willick, 1995). Type Ia SNe are generally able to obtain more accurate distance measurements with an uncertainty of 5–10% (Fakhouri et al., 2015; Rest et al., 2014). Large uncertainties on the distance measurements that propagate through to PV calculations lead to a noisier measurement of the bulk flow. However, this measurement noise can be combated by increasing the number of PV measurements. PV measurements may also have a non-Gaussian PDF, although this may be overcome in various ways (Watkins and Feldman, 2015a; Qin et al., 2018; Qin, 2021; Hoffman et al., 2021). In general, the galaxy peculiar velocities may have uncertainties that follow a log-normal distribution. Fortunately, the Watkins and Feldman (2015a) estimator allows for PVs with Gaussian error bars to be estimated from Gaussian distributed log-distance ratios,¹ under the assumption that the PV u_m obeys $u_m \ll cz_{m,\text{CMB}}$, where $z_{m,\text{CMB}}$ is the observed redshift of the galaxy from spectroscopy, i.e. the peculiar velocity is much smaller than the recession velocity. Alternatively, the Box Cox transformation also allows for any distribution to be Gaussianised (Qin, 2021).

A second issue which arises is that only the radial (line-of-sight) component of the galaxy’s peculiar motion, $u_m = \mathbf{v} \cdot \hat{\mathbf{r}}_m$ can be measured. This is because the PVs are derived from the galaxy’s redshift, which for non-relativistic velocities is only due to the radial component of the galaxy’s motion away from the observer. Therefore, the weighting scheme applied to the PVs in order to estimate the bulk flow must be derived in such a way that it as closely as possible captures the 3-dimensional bulk flow from what is effectively 1-dimensional PV data for each galaxy. Nusser (2014) shows that under the assumption that the 3-dimensional PV field has no curl component, it is possible to entirely gain the 3-dimensional information about the bulk flow from only the radial projection of the PV field, for a full-sky dataset. Results from tests on mocks (in the next section) show it also is generally possible (on average) when the bulk flow vector can be described as a constant valued vector across the survey

¹The log-distance ratio η of a galaxy is defined as $\eta = \log_{10} \left(\frac{d(z_{\text{CMB}})}{d(z_{\text{rec}})} \right)$, where $d(z_{\text{CMB}})$ is the comoving distance to the galaxy computed from a cosmological model given the observed redshift of the galaxy, and $d(z_{\text{rec}})$ is the true comoving distance to the galaxy. A measurement of η is typically what is obtained from the Fundamental Plane or Tully–Fisher relation in order to calculate a PV.

volume. Alternatively, one can consider scheme's such as the Wiener Filter (Zaroubi et al., 1995) to reconstruct the 3-dimensional PV field from the data and thus estimate the bulk flow. However as mentioned previously, for small or noisy samples of PVs the prior information assumed from a cosmological model when employing this method can dominate the underlying signal. This makes the process of comparing a bulk flow thus derived to a cosmological model somewhat circular.




Figure 4.3: The theoretical bulk flow V_p as a function of the survey radius R for two survey geometries. The shaded regions show the $1 - \sigma$ uncertainty on the theoretical bulk flow due to cosmic variance. For each line plotted, the legend shows the functional form of the number density of objects $n(r)$ for the survey selection function. The predictions here have been calculated using the equations and methods given in Andersen et al. (2016). The predicted bulk flow has been computed using a grid of 500^3 points, with $H_0 = 67.74 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.3089$ and with $\Omega_\Lambda = 1.0 - \Omega_m$ for the standard Λ CDM model.

Andersen et al. (2016) looked into how the survey geometry, in particular the sky mask applied to the survey data, can affect bulk flow measurements. Using simulations, their paper shows that the predicted theoretical bulk flow needs to take into account the geometry of the PV galaxy survey data before making a comparison of a bulk flow measurement to theoretical predictions. They also propose how the bulk flow from theoretical predictions can be more accurately computed to take into account the survey geometry. I demonstrate the effect of the survey geometry on the theoretical bulk flow prediction using their methods in Figure 4.3. This demonstrates how the theoretical bulk flow should change for different survey selection functions as a function of survey depth. Furthermore Andersen et al. (2016) show with simulations that under sampling of the PV field (i.e., a number of measurements < 500) affects the bulk flow amplitude and gives resulting bulk flow measurements that do not necessarily agree with expectation from theory as the variance of the measured bulk flow increases; in the case of small sample sizes, a measured bulk flow needs to be compared to a prediction from simulations or mock catalogues. Overall, it is most important to take care if directly comparing the bulk flow amplitude and cosmic variance uncertainty to a measurement, especially when placing a coordinate for the amplitude of the measurement on a theory plot similar to that shown in Figure 4.3.

Finally, the last issue I will discuss here and which is a main focus of investigation in this work is related to the underlying accuracy and precision of the estimators used to measure the bulk flow. Various popular estimators exist in the literature to derive bulk flow measurements, but incorporate assumptions about the nature of the measured bulk flow, or the data, that can have an affect on their ability to 1) return a realistic precision on the bulk flow from the data and 2) return an accurate bulk flow measurement compared to the moment of the data we expect the estimator to recover.

Here in particular I investigate the precision of the Kaiser (1988) maximum likelihood estimator (Kaiser MLE) which has widely been used through out the literature to measure the bulk flow. Various authors have found the MLE approach does not provide a realistic estimate of the precision on the measured bulk flow moment (the error bars on the measurements are underestimated), according to a χ^2 goodness-of-fit analysis (Qin et al., 2018, 2021; Howlett et al., 2022). Thus this investigation aims to test the performance of the Kaiser MLE on mock data to gain a better understanding of its precision and accuracy. Likewise I apply the same investigation to the Minimum Variance method (MVE) by Watkins et al. (2009). This estimator has been used in the literature and associated with claims of bulk flows measurements in tension with Λ CDM. While Agarwal et al. (2012) has previously shown the estimator to be unbiased and precise for mocks designed for the SFI++, COMPOSITE and DEEP peculiar velocity surveys, I build on this work and further test its performance on current, much larger, datasets. Finally, I also investigate the performance of the MLE method presented in Nusser (2014) and the Minimum Variance Estimator presented in Peery et al. (2018), that are both variants of the Kaiser MLE and Watkins MVE, respectively, and are designed to more accurately capture the bulk flow moment from data, in addition to measure a moment of the data that more closely aligns with the bulk flow as defined in Equation 4.1.

In general, bulk flow estimators such as the Kaiser MLE do not necessarily obtain the same moment as others such as for example, the Watkins MVE (Nusser, 2014, 2016). It is important to distinguish that the moments that are estimated may differ and that one should not assume the bulk flow moment is always consistent with Equation 4.1. The moments from individual estimators should thus carefully be compared to a theoretical bulk flow amplitude or expectation from simulations, as discussed in Andersen et al. (2016).

4.3.1 Discussion of methods

Kaiser Maximum Likelihood method

The Kaiser Maximum Likelihood estimator (Kaiser MLE; Kaiser, 1988) is an analytic weighting scheme to derive bulk flow measurements; this was discussed in Chapter 3 however it is fundamental to this chapter so I repeat the details here. In this method, the likelihood function for measuring a bulk flow vector \mathbf{B} given a set of N observed radial velocities u_m is written as

$$\mathcal{L}(\mathbf{B}, u_m) = \prod_{m=1}^N \frac{1}{\sqrt{2\pi(\sigma_*^2 + \sigma_m^2)}} e^{-\frac{(u_m - \mathbf{B} \cdot \mathbf{f}_m)^2}{2(\sigma_*^2 + \sigma_m^2)}}, \quad (4.3)$$

where σ_m is the uncertainty in each velocity measurement, and σ_* is an extra component of uncertainty to account for random non-linear motions. If a fixed value for σ_* is assumed one can solve for the weights for the galaxy velocities that maximise this likelihood function with respect to the bulk flow by solving $\frac{d\mathcal{L}}{dB_i} = 0$. One estimates the bulk flow components along each coordinate axis as

$$\tilde{B}_i = \sum_{m=1}^N w_{i,m} u_m, \quad (4.4)$$

and the solution for the weights $w_{i,m}$ are given by

$$w_{i,m} = \sum_j A_{i,j}^{-1} \frac{\hat{n}_{j,m}}{(\sigma_*^2 + \sigma_m^2)}, \quad (4.5)$$

where we define

$$A_{i,j} = \sum_m \frac{\hat{n}_{i,m} \hat{n}_{j,m}}{(\sigma_*^2 + \sigma_m^2)}. \quad (4.6)$$

The form of the likelihood function for this method encodes a number of assumptions:

- the galaxy PVs are uncorrelated;
- the galaxy PVs have errors that are drawn from a Gaussian distribution;
- each PV has a component of velocity due to a bulk flow vector that is constant across the entire volume, such that each radial PV can be written as $u_m = B_\alpha \hat{n}_{i,\alpha} + \delta_m$, where δ_m represents a random component of velocity that is not due to the underlying bulk flow (Nusser, 2014);
- that σ_* is a fixed value;
- that the velocity field is well modelled by only the bulk flow modes (higher order modes of the field can be neglected);
- and that the galaxy PV observations are insensitive to small scale flows.

None of these assumptions are generally true. We are interested in exploring how much these may impact the performance of this estimator. This may allow us to understand why, when this estimator is applied to realistic mocks, the analysis of the recovered mock bulk flows compared to the true bulk flow moment of each mock generally results in a reduced $\chi^2 > 1$ for the goodness-of-fit (Qin et al., 2018, 2021; Howlett et al., 2022), which implies the model for this estimator results in underestimated uncertainties even if the results are unbiased.

In the work of Qin et al. (2018), a modified version of the Kaiser estimator is developed, the η -MLE estimator. This estimator searches for the Maximum Likelihood bulk flow using a Likelihood function for the observed log-distance ratios η of the data rather than the PVs in order to avoid the issue of non-Gaussian uncertainties on PV measurements (as an alternative to the other approaches to deal with this issue, discussed previously). Furthermore they take an MCMC approach to search for the best fitting bulk flow modes and also allow σ_* to vary rather than setting a fixed value. Furthermore, in Qin et al. (2021), this approach is used while also allowing the shear moments of the velocity

field to be modelled. In these works the reduced χ^2 is generally still greater than unity. This would imply the assumptions the Kaiser estimator encodes regarding σ_* , the nature of the PV errors, and the higher order moments of the field (which are no longer assumptions for the η -MLE method), might be unrelated to issues regarding the precision of the recovered bulk flows using the Kaiser MLE approach. To explore further I investigate the Kaiser estimator in more detail by testing its performance with mocks in section 4.4.

Nusser Maximum Likelihood Estimator

The Maximum Likelihood method proposed in Nusser (2014) (hereon Nusser MLE) is very similar to the Kaiser MLE approach. Nusser (2014) shows the Kaiser MLE method only obtains the bulk flow of a survey volume as defined in Equation 4.1 if one can reliably assume that the bulk flow is effectively a constant across the survey volume, which is not generally consistent with the definition of the bulk flow as an average of the peculiar velocities in the volume as in Equation 4.1.

In summary, Nusser (2014) shows in the continuous limit and when there is no angular selection function applied to the galaxies, the estimated quantity for each bulk flow mode obtained by maximising the likelihood function given by the Kaiser MLE approach can be written as

$$\tilde{B}_i = \frac{3 \int r'^2 dr' d\Omega \frac{w(r') \bar{n}(r')}{\sigma^2(r')} u(\mathbf{r}') \hat{n}_i(\mathbf{r}')}{4\pi \int r'^2 dr' \frac{w(r') \bar{n}(r')}{\sigma^2(r')}}. \quad (4.7)$$

In this equation Ω is a solid angle, σ is the uncertainty on the galaxy PV (which in this limit is dependent only on r) and \bar{n} is the mean number density of galaxies (which is also assumed to be dependent only on r). For the Kaiser MLE estimator we have $w(r) = 1$. Nusser (2014) shows that this equation is only consistent with the definition of the bulk flow in Equation 4.1 if one sets $w(r) = \frac{\sigma^2(r)}{\bar{n}(r)r^2}$ (\hat{n} is still as before, $\hat{n}_i = \hat{\mathbf{x}}_i \cdot \hat{\mathbf{r}}$). In the Nusser MLE, this modification has been adopted to re-derive the analytical solution to the maximum likelihood weights. In the case that PV field is curl-free and the radial selection function is spherically symmetric, the Nusser (2014) modification results in the following weighting scheme,

$$w_{i,m} = \sum_j A_{i,j}^{-1} \frac{\hat{n}_{j,m}}{(\bar{n}_m r_m^2)}, \quad (4.8)$$

where we define

$$A_{i,j} = \sum_m \frac{\hat{n}_{i,m} \hat{n}_{j,m}}{(\bar{n}_m r_m^2)}. \quad (4.9)$$

In the above equations \bar{n}_m is the number density of galaxies at the comoving distance r_m the galaxy appears to be from the observer. Nusser (2014) shows this weighting scheme obtains a more accurate bulk flow estimate as defined in Equation 4.1 than the Kaiser MLE weighting scheme, using a single simulation. This method has the advantage over the Kaiser method of not requiring the bulk flow to be well-represented as a constant-valued vector across the survey volume, although it has the disadvantage of requiring spherical symmetry in the survey geometry, and otherwise shares the same assumptions that are encoded in the Kaiser method. I investigate its ability to measure the bulk flow accurately and precisely on mock data here.

Watkins, Feldman and Hudson Minimum Variance method

For comparison it is interesting to explore the measured bulk flows obtained from mocks when applying the Minimum Variance estimator (MVE) method by Watkins et al. (2009), also discussed in Chapter 3. This estimator was shown to be unbiased in Agarwal et al. (2012) however we are interested in also exploring the precision and accuracy of this estimator in a similar way to the Kaiser MLE method.

Recall the aim of the MVE method is to obtain a bulk flow measurement from the data that minimizes the variance in the difference between the bulk flow estimate \tilde{B}_i of the data and the measured bulk flow that would be obtained from a survey with an ideal window function, U_i . This in principal allows for the measured bulk flow to be more comparable to other surveys in which this method has been applied even when they have different geometries, as long as the window function of U_i is the same. The approach involves the use of the Lagrange multiplier method in order to also satisfy a constraint so that the measured bulk flow amplitude \tilde{B}_i is correct on average. The Lagrangian function to be minimized with respect to the weights $w_{i,m}$ for each m^{th} galaxy contribution to the i^{th} bulk flow mode is given by

$$\mathcal{L} = \langle (\tilde{B}_i - U_i)^2 \rangle + \lambda_{ij} \sum_m w_{i,m} \hat{n}_{i,m}, \quad (4.10)$$

where λ_{ij} is the Lagrange multiplier and $\hat{n}_{i,m}$ is the same as previously. The resulting weighting scheme can be found after solving for $\frac{d\mathcal{L}}{dw_{i,m}} = 0$. This was described in more detail back in Chapter 3, and other details can be found in Watkins et al. (2009); Feldman et al. (2010); Agarwal et al. (2012); Scrimgeour et al. (2016). Like the Kaiser method, this method also incorporates, although much more weakly, the assumption that the measured radial PVs of galaxies contain a component that is due to a constant bulk flow vector across the entire survey volume; in fact Nusser (2016) shows that these methods are equivalent in the limiting scenarios that the correlations between the ideal survey galaxies and real survey galaxies approaches zero.

We can expect some of the assumptions listed previously which may affect the results from the Kaiser MLE approach (that σ_* can be a fixed value, that the galaxy PV measurements are drawn from a Gaussian PDF) also affect the results from the Watkins MVE in a similar way (although as discussed previously, there are approaches in the literature to Gaussianise PV measurements given in Watkins and Feldman, 2015a; Qin, 2021; Hoffman et al., 2021). However, the Watkins MVE does take into account that there are non-zero linear correlations between galaxy peculiar velocities. It is also less sensitive to small scale flows because of the constraint it enforces on the window function of the measured bulk flow moments. Both this method and the Kaiser method suffer from the potential risk due of leakage of signal due to higher order modes in the velocity field affecting the bulk flow estimate, although in Feldman et al. (2010) this method is extended to include modelling for higher order moments of the field. On the other hand, in Feldman et al. (2010) they find that not including the higher order moments of the field such as the shear and octupole did not contaminate their bulk flow measurements. In our own tests with the Kaiser MLE in which we extended the method to allow modelling for higher order modes and applied it to realistic mocks for the SDSS data, we found that failing to model the higher order modes did not reduce the accuracy of the bulk flow measurements.

This is discussed further in section 4.5.

Peery Minimum Variance Estimator

Finally we discuss the Minimum Variance Estimator that is presented in Peery et al. (2018) (hereon, Peery MVE). This is effectively the same as the MVE method developed by Watkins et al. (2009) discussed previously, but with two modifications. Firstly, the ideal survey that is used to constrain the window function of the estimated bulk flow in the MVE method is chosen to follow a radial distribution with selection function such that the number density of objects $n(r)$ is always proportional to r^{-2} , where r is the radial distance to objects in the ideal survey. Or alternatively, the objects in the ideal survey follow a uniform radial selection function, but are weighted by an additional factor of r^{-2} compared to the ideal weights in the original MVE method. By following either of these approaches this ensures that one obtains an estimate of the bulk flow much more closely aligned to the moment defined in Equation 4.1 without encoding the assumption that the bulk flow is constant across the survey volume, as is true in the Nusser MLE method. The authors arrive at this scheme following and expanding on the derivations shown in Nusser (2014) for application to the MVE scheme.

Secondly, an additional constraint equation is introduced to the Lagrangian function, that ensures $\sum_m cz_m = 0$. This constraint is introduced in order to allow the estimator to be independent on uncertainty in the Hubble constant H_0 . Effectively, we can understand that this constraint equation should make the estimated bulk flow independent of the global zero-point calibration of the dataset. More details regarding this estimator can be found in Peery et al. (2018). We can expect that this variant of the MVE method thus has two advantages over the Watkins MVE approach, and unlike the Nusser MLE approach does not encode any assumptions about the survey geometry. I present results for the application of this estimator to mock data also.

4.4 Evaluation of bulk flow estimators

In this section, I use different sets of mock PV survey data to test the performance of the estimators discussed previously to accurately and precisely recover the bulk flow. For all mocks, galaxies are placed at their observed redshifts. I begin with results from tests on simplistic mocks before moving on to more sophisticated mock data that involves modelling for the growth of structure in the Universe. I conclude by testing these estimators on the fully realistic mocks from numerical simulations that were described in section 4.2 and use these results to comment on the performance of the estimators to recover the bulk flow in real data.

It should be noted, the simplistic mocks and Zeldovich mocks (described in the following sections) were produced for analyzing the performance of the bulk flow estimators specifically when the survey geometry is altered or the statistical properties of the velocity distribution is changed, without the additional complexities that are introduced in realistic mocks for data. As stated previously, the goal has been to thoroughly evaluate the estimators and determine what systematics may affect their performance; the simplistic mocks and Zeldovich mocks are, unlike the realistic mocks for CF4, not

representative of any realistic datasets but are used to gain a general sense of the limitations of the estimators. For reasons that will become apparent, I also do not apply all four estimators for every test or set of mocks with different properties that are constructed.

4.4.1 Tests on simplistic mocks

Generating the mocks

The simplistic mocks were generated as follows:

- I generate points within a sphere with a chosen radial distribution representing galaxies (various are used) and record right ascension (RA), declination (Dec), cosmological redshift (z_{rec}), and the radial comoving distance to the point from the observer in real space ($D(z_{\text{rec}})$). To obtain a radial distribution of objects that is not uniform, I simply use downsampling to obtain the desired distribution of points as a function of comoving distance.
- I generate a radial velocity covariance matrix for the data points using linear theory (see Appendix A.1 for more detail). The covariance matrix depends on a chosen cosmological model (given in the next section) and the coordinates for the data points. I use the linear theory covariance matrix to draw Gaussian random velocities for each object in the mock. I do this using the function `numpy.multivariate_normal()` from the python package NUMPY. Given a covariance matrix and mean vector for a set of parameters (in this case the parameters corresponds to the radial PV for point generated), the function generates random values for each parameter I use to simulate data. The data are drawn from a probability distribution function (PDF) corresponding to a multi-variate Gaussian centered on the mean given from the mean vector. The variance of the PDF and correlations with other measurements are determined by the input covariance matrix. I set the mean vector to $\bar{\mathbf{u}} = \mathbf{0}$.
- I generate a random bulk flow vector with an x , y and z component by drawing from a uniform distribution for each mode. I treat this bulk flow vector as the true bulk flow of the mock data we expect our estimator to recover. The radial component of this bulk flow is added to the radial PVs drawn in the previous step, which are then treated as ‘true’ radial PVs of each galaxy. Each galaxy thus has a velocity component that is due to the same underlying bulk flow \mathbf{B} and their own random component of velocity due to statistical variation generated from linear theory in the previous step, $u_{\text{linear},m}$. Overall for each object labelled by m we have

$$\begin{aligned}
 u_{\text{total},m} &= u_{\text{linear},m} + \mathbf{B} \cdot \hat{\mathbf{r}}_m \\
 &= u_{\text{linear},m} + B_x \cos(\theta_m) \sin(\phi_m) + B_y \sin(\theta_m) \sin(\phi_m) \\
 &\quad + B_z \cos(\phi_m).
 \end{aligned} \tag{4.11}$$

where θ and ϕ are angles in spherical polar coordinates that can be mapped to RA and Dec for each object.

- The radial PVs from the previous step are used to determine a redshift observation in the CMB–frame for each object. I use the equation

$$z_{\text{CMB, obs}} = (1 + z_{\text{rec}})(1 + z_{\text{pec}}) - 1. \quad (4.12)$$

For each object I already have z_{rec} from the first step. z_{pec} (the Doppler redshift due to the radial PV of the galaxy) can be calculated from the radial PVs generated in the previous two steps. I use the special relativistic Doppler shift equation.

- Finally I use $z_{\text{CMB, obs}}$ and z_{rec} to compute truth log–distance ratios η_t for each galaxy as $\eta_t = \log_{10} \left(\frac{d(z_{\text{CMB}})}{d(z_{\text{rec}})} \right)$. To obtain observed log–distance ratios η_o , I draw random η_o for each galaxy from a Gaussian probability distribution function with a mean given by η_t and a standard deviation σ_η given by a constant. This constant is used to represent the uncertainty in the observed η_o . I choose $\sigma_\eta = 0.05$.
- Finally I use the estimator of Watkins and Feldman (2015a) to convert the log–distance ratio observations and their uncertainties to radial PVs and PV uncertainties for each object,²

$$u_m \approx \frac{cz_{\text{mod}}}{1 + z_{\text{mod}}} \eta_o \ln(10) \quad (4.13)$$

and

$$z_{\text{mod}} = z_{\text{CMB}} \left[1 + \frac{1}{2}(1 - q_0)z_{\text{CMB}} - \frac{1}{6}(j_0 - q_0 - 3q_0^2 + 1)z_{\text{CMB}}^2 \right], \quad (4.14)$$

where q_0 and j_0 are the deceleration and jerk parameters, respectively.

Results summary: simplistic mocks

Simplistic mocks are generated each with 5000 galaxy PVs. I choose a Λ CDM cosmological model with $H_0 = 69 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.31$, $\Omega_\Lambda = 0.69$. I generate sets of the mocks with varying geometries by 1) altering the radial selection function by downsampling to a desired number density of objects from an uniform distribution, as discussed previously, and 2) generating objects in a desired sky mask (sky coverage) by limiting the allowed values for θ and ϕ for each object. In summary, I found the Watkins MVE and Kaiser MLE were able to generally obtain an unbiased estimate $\tilde{\mathbf{B}}$ of the bulk flow vector \mathbf{B} added to the mocks, regardless of the radial selection function of the points or the sky mask applied to the data. In general I find the reduced χ^2 goodness–of–fit of the recovered bulk flows to the true underlying bulk flows obtains $\chi^2 \sim 1$. These conclusions are demonstrated in Figure 4.4 and Figure 4.5. It is also apparent from these tests, that when the bulk flow vector can be described as a constant valued vector across the survey volume, then the bulk flow moments obtained from these two different estimator variants are consistent with each other and also the bulk flow moment as defined in Equation 4.1 (Nusser, 2016).

²It should be noted this estimator involves making approximations for the true cosmological redshift-comoving distance relation. A Taylor expansion is used to derive z_{mod} which makes a substitution for the integral $\int \frac{dz}{E(z)}$ that enters the correct expression for comoving distance.




Figure 4.4: The averaged recovered bulk flow $\langle \tilde{B}_i \rangle$ component from simplistic mocks with a cone geometry. Each panel shows $\langle \tilde{B}_i \rangle$ in the x , y and z directions. The value for \mathbf{B} (the underlying constant-valued bulk flow vector) is the same for each mock and the components are shown by the horizontal dashed line in each panel. Each point shows the average recovered bulk flows of 32 mocks with the same opening angle θ , plotted against θ ; $\theta = 180$ corresponds to a spherical survey, $\theta = 90$ corresponds to a hemisphere. For all the mocks the spatial distribution of objects follows a Gaussian radial distribution such that $n(r) \propto e^{-r^2}$ with a standard deviation of $\sim 50 \text{Mpch}^{-1}$. The recovered bulk flows from the Kaiser MLE and the Watkins MVE methods are shown in blue and orange respectively. The error bars on each point represent the standard deviation of $\langle \tilde{B}_i \rangle$.




Figure 4.5: The difference between the recovered bulk flow components $\langle \tilde{B}_i \rangle$ from the components of the constant-valued bulk flow vector \mathbf{B} for 512 mocks using the Kaiser MLE (blue points) and Watkins MVE (red points), plotted against B_i (the components of \mathbf{B}) for simplistic mocks. The panels from top to bottom show the results for the x , y and z directions. The results in the left panel are for a set of mocks with a selection function given by a Gaussian radial distribution, the right column of panels a uniform radial distribution. Each data point is the average recovered bulk flow from 32 mocks with the same vector \mathbf{B} added to the data. The error bars give the standard deviation of $\langle \tilde{B}_i \rangle$.

Figure 4.4 shows the results of applying the Watkins MVE and Kaiser MLE method to mocks with a set bulk flow \mathbf{B} with the shape of a cone with varying opening angles ranging from a narrow cone to a full spherical mock, and with the number density of objects $n(r)$ at a given radius from the observer, r , given such that $n(r) \propto e^{-r^2}$. Figure 4.5 shows the results of applying both estimators to sets of fully spherical mocks with two different radial selection functions specified just above the plot panels and with various bulk flows. In all cases $\tilde{\mathbf{B}}$ from both of the estimators is unbiased compared to \mathbf{B} and $\chi^2 \sim 1$. This tells us that overall the Watkins MVE and Kaiser MLE are good estimators for \mathbf{B} in this simplistic model where the radial peculiar velocities have a component due to $\mathbf{B} \cdot \hat{\mathbf{n}}$ and a Gaussian random component of motion. In this model, the survey geometry does not affect the ability of these estimators to recover \mathbf{B} on average. More supporting results for these mocks not included here for conciseness can be found here. Given the nature of these mocks, we can expect the results will not differ for the Peery MVE estimator and do not show results when applying these estimators; however we will see the Nusser estimator may indeed suffer when a non-spherical sky mask is used.

4.4.2 Tests on mocks with linear structure growth

Generating the mocks

The second set of mocks I test here are more realistic in that they include structure growth in the Universe which changes the spatial distribution of objects compared to the simplistic mocks described previously. However, these have been deliberately made to incorporate only linear structure growth of matter as closely as possible. I refer to these mocks as ‘Zeldovich mocks’. These are generated with the same cosmological model as specified for the simplistic mocks in the previous section.

I use the n -body simulation code L-PICOLA (Howlett et al., 2015) to generate these mocks. The L-PICOLA code uses an approximate analytical solution to calculate the initial velocities and positions of dark matter particles due to linear structure growth, called the Zeldovich approximation. It then evolves the simulation with a numerical algorithm to accurately capture the effects of non-linear structure growth from a specified initial redshift by the user to a final redshift. We expect that non-linear structure growth has an effect on the PVs in the simulation because it is known that the galaxy PVs follow a distribution with higher kurtosis (or extended tails) compared to a Gaussian distribution when non-linear structure growth is present (Sheth and Diaferio, 2001). It was hypothesised that the altered velocity distribution in the mocks when non-linear structure growth is present, in addition to the complexity that is introduced into the spatial distribution by structure growth, may have an effect on how well bulk flow estimators recover the bulk flow. Therefore to isolate the effect of altering the velocity distribution due to non-linear structure growth and test how well the estimators work when it is not present, I evolve only the linear structure growth (which is calculated with the Zeldovich approximation) in the n -body code and factor out the evolution that causes non-linear structure growth. This can be done using L-PICOLA by setting the initial and final redshift in the simulation both to $z = 0$. Since there is only linear structure growth, I do not identify dark matter halos and or galaxy populations within them, but instead assume that each dark matter particle in the simulation can be

approximated as a single galaxy. This is a valid choice for our purposes given it is expected that the PV field is largely insensitive to galaxy bias (Zheng et al., 2015). For each simulation box I create four mocks by specifying four different origins for an observer within the box, and we ensure the distance between the four origins is enough such that data from the mocks is not overlapping and that the measured data will not be correlated.

The realistic mocks in Qin et al. (2018, 2021); Howlett et al. (2022) are created using the L-PICOLA code but *do* include non-linear structure growth. I can compare how well the Kaiser MLE measures the bulk flow using the Zeldovich mocks described here and compare the results to those found by Qin et al. (2018, 2021); Howlett et al. (2022) who apply the Kaiser MLE (or variations of it) to their mocks to determine whether there is any relation between the performance of these estimators and the shape of the distribution of PVs. Therefore, for some of the tests I create SDSS PV survey simulations with the Zeldovich mocks described here by following the exact procedure implemented by Howlett et al. (2022) to create the SDSS mocks, aside from the main difference being that L-PICOLA full-numerical simulations are not used to generate galaxy positions and velocities. For tests with these mocks I focus on just the Kaiser MLE approach since our goal is to purely isolate the effect of having non-linear structure growth.

I apply tests with all four of the estimators described previously to more general Zeldovich mocks. For these more general mocks, which are fully spherical and contain just 5000 galaxies (by downsampling from simulations) with various different choices of radial survey selection functions, the log-distance and PV observations are simulated in the same way as the steps described for the simplistic mocks. I generate observations of η from the true values of η by drawing Gaussian random values with a set value for σ_η for the Gaussian (which becomes the uncertainty on the observation). The estimator by Watkins and Feldman (2015a) is also used to convert the log-distance ratios to PV observations. With these mocks, I am also able to test whether other factors that are present in the mocks, such as the properties of the spatial and velocity distribution or survey geometry affects the precision or accuracy of bulk flow estimators, while having factored out further complexity due to non-linear structure growth.

Results summary: Zeldovich mocks

I begin by noting here that we compare the estimated moment $\tilde{\mathbf{B}}$ from the Peery MVE (that we will hereon label $\tilde{\mathbf{B}}_P$) and the Nusser MLE ($\tilde{\mathbf{B}}_N$) to the moments \mathbf{B}_P , \mathbf{B}_N respectively, defined as the sum of the true 3D PVs of each object in the simulation (essentially Equation 4.1 for both estimators). However, for the Kaiser MLE and the Watkins MVE the estimated moment slightly differs; as mentioned previously it is important to ensure the estimated moment of the data is compared to the correct theoretical expectation (or in this case, the correct moment of the simulations). For the Kaiser MLE, the true 3D PVs of each galaxies are weighted by $(\sigma_n^2 + \sigma_*^2)^{-1}$ to define \mathbf{B}_K and for the Watkins MVE by $e^{-r^2/(2R^2)}$ to define \mathbf{B}_W , where R is the standard deviation of the Gaussian used to define the window function of the ideal survey used in this estimator. This choice is taken because the Watkins MVE attempts to obtain a bulk flow estimate from the (ideal) Gaussian-weighted volume, and the

Kaiser MLE obtains a bulk flow estimate of the most likely constant bulk flow vector across the volume, with the velocities weighted by the uncertainty on each observation.³

I summarise our findings from a number of tests performed on the Zeldovich mocks here, while more evidence and plots to support these results can be found here.

Firstly, I tested the Kaiser MLE on the SDSS–Zeldovich mocks. I found that when applying the Kaiser MLE to these mocks, the results are unbiased, similarly to results found in Howlett et al. (2022). However, I also find that the precision of the recovered bulk flows is under-reported; the reduced χ^2 is much greater than unity, similar to that found by Howlett et al. (2022) with the L–PICOLA mocks. Therefore the conclusion from this result that the effects of non–linear structure growth on the spatial or velocity distribution of galaxies is not likely related to the precision or under–reporting of errors for the Kaiser MLE method in general. I also expect that to find similar results using other estimators given that they do not include any additional modelling for non–linear theory PVs.

I also extended the analysis for these mocks to incorporate modelling for higher order modes of the PV field (see the definition of the shear modes in section 3 of Qin et al., 2019, and references within) by simply substituting the expansion of the PV field with the mode expansion into the likelihood function for the Kaiser MLE approach, in place of the field expansion defined as the bulk flow modes only, $\mathbf{B} \cdot \hat{\mathbf{r}}$. I found that the resulting bulk flows were unbiased but the extra degrees of freedom did not improve the precision of the recovered bulk flows. I thus concluded that aliasing due to a failure to model higher order modes of the velocity field does not influence how precisely the Kaiser MLE can estimate the bulk flow with this set of mocks. I do not extend this test to the Watkins MVE as this result has effectively been shown already to be true by Feldman et al. (2010) for this estimator, although it might be expected the importance of modelling these modes to be more largely dependent on the data or mocks used.

In further tests on more general Zeldovich mocks with 5000 objects and varying choices for the radial selection function of the mock data, I tested all four estimators. In general it was found when applying the Kaiser MLE or Watkins MVE to these mocks, in all choices for the radial selection function of the survey the estimated bulk flows are unbiased compared to the expected moments and both estimators in fact give consistent results to each other for the estimated moments. However, the reduced χ^2 is usually greater than unity. Additionally for both methods, the reduced χ^2 changes depending on the radial selection that is applied to the mock catalogues, in contradiction to our results with the simplistic mocks. I find there is a correlation between the χ^2 goodness of fit and properties related to the radial selection function applied to the data.

In contrast, when applying the Peery MVE and Nusser MLE, which also obtain unbiased estimates of the bulk flow, there is no trend present when applying them to the same sets of mocks, and in general the reduced χ^2 is closer to unity, regardless of the radial selection function. This would lead us to

³One might expect based on how the MVE estimator is designed to weight the galaxies such that the bulk flow window function matches a chosen ideal survey, that we should actually compute the expected bulk flow amplitude as the average of all the velocities in the simulation within the volume defined by the ideal radius. However I found in general that this moment is not well-captured by the estimator and the calculated moment is more accurately and precisely captured by the moment we define for comparison in our tests.

conclude, that an underlying assumption that plays a role is the low-precision in results is that the bulk flow vector cannot be well represented as a constant-valued vector across the survey volume and that the radial selection function of the data is important in the bulk flow measurement; this is one of the differences between the Watkins MVE and Kaiser MLE compared to the Peery MVE and Nusser MLE.

Regarding the observed trend between the radial selection function of the survey and χ^2 goodness-of-fit, I note that when the observed log-distance ratios η_o for the mock data are drawn from Gaussian distributions centered on η_t for each object with a constant σ_η , the distribution of observed PVs that are estimated following the estimator by Watkins and Feldman (2015a) changes in shape and kurtosis depending on the radial selection function of the data, although the distribution of true PVs from the mocks does not. I emphasise here that this is *not* a bias or flaw concerning this PV estimator. However the changing in the PV distribution can be understood because the spread and tails (which I quantify using the excess sample kurtosis of the distribution compared to a Gaussian distribution) of the distribution of *true* galaxy log-distance ratios (with a mean of zero) is expected to be smaller when there are more objects at a further distance from the observer where we might expect the ratio $\frac{d(z_{\text{CMB}})}{d(z_{\text{rec}})}$ to be closer to one.⁴ Thus for a choice of a constant σ_η in our mocks, changing the radial selection function of the mock data will change the observed distribution of log-distance ratios due to the way $\frac{d(z_{\text{CMB}})}{d(z_{\text{rec}})}$ is affected and the same applies to the distribution of observed PVs that are estimated from the log-distance ratios.

As the percentage error on each η_o must be more or less significant relative to their true value η_t when the distribution of η_t changes for a fixed choice σ_η , we might expect to see some correlation between the kurtosis of the PV distribution (which has a shape following the η_o distribution when the PV estimator by Watkins and Feldman (2015a) is used) and the χ^2 goodness-of-fit to the measured bulk flows. This is demonstrated by Figure 4.6. However we would expect that if the uncertainty in the measured bulk flows are well captured by the choice of weighting from the bulk flow estimator and uncertainties given to each data point, this might not be the case. The excess sample kurtosis of the resulting PV distribution appears to be correlated with the reduced χ^2 for a set of mocks that have the same radial selection function, when applying the Kaiser MLE or Watkins MVE. This is demonstrated in Figure 4.7. The correlation is not apparent for the results from the Peery MVE or Nusser MLE.

This is interesting, because the motivation behind the Nusser MLE in its derivation was to solve for an improved weighting scheme for the bulk flow of the volume, without requiring the assumption of a constant valued vector across the survey volume. The Peery MVE estimator attempts to achieve the same. These results suggest that it is this assumption that causes the underestimation of the error bars in the standard Kaiser and Watkins estimators, in a complex way that is also tied to the PV survey geometry. In general, the conclusion is that the Peery MVE and Nusser MLE should be applied for bulk flow estimates, although the Nusser estimator does rely on another assumption regarding spherical symmetry, which the Kaiser MLE approach does not rely on. Unfortunately, as commented on further

⁴At higher z , the impact of z_{pec} on the observed redshift z_{CMB} should be less, so the ratio $\frac{d(z_{\text{CMB}})}{d(z_{\text{rec}})}$ should approach unity and the spread in the distribution should converge if there are more objects at higher redshift.




Figure 4.6: This plot demonstrates the typical difference between an observed value of η_0 and the true value η_t , in the case η_0 is drawn from a gaussian random distribution. Here I have fixed the uncertainty (the width of the gaussian distribution) for all observations in the mock data and fixed the ‘true’ peculiar velocity that impacts the redshift of the observed distance to the object. I draw 500 values of η_0 for each value z (and thus each possible η_t) and plot the average of $\frac{\eta_t - \eta_0}{\eta_t}$. At various distances the average of this quantity shows the observations are unbiased, but the standard deviation grows (error bars). Small offsets in the x -positions separate otherwise overlapping data points.

in the next section we find that this assumption breaks down catastrophically for the CF4 geometry, which makes the Nusser MLE estimator unsuitable.

4.4.3 Tests on realistic mocks for CF4 data

In this section, I test the performance of the bulk flow estimators on mocks which fully capture the same selection effects as the CF4 dataset and non-linear structure growth. The details of the process to generate these mocks were discussed in section 4.2. The cosmological model is the same as for previous mocks. I tested the performance of the Kaiser MLE, Nusser MLE and the Peery MVE in detail, and report primarily on the performance of the first and last of these. This is for the following two reasons:

1. I only test the Peery MVE, rather than also testing the Watkins MVE, due to the conclusions from the previous results but also as we expect it to be more useful given that it incorporates a constraint to ensure the results are independent of systematic errors in the zero-point calibration applied to the CF4 data.
2. I found that the Nusser MLE gives a biased result on the CF4 mocks with extremely large scatter between mock realisations, which we attribute to the lack of spherical symmetry in the geometry of the CF4 data, which is an underlying assumption of the Nusser MLE approach. These results can be seen here.




Figure 4.7: Left panel: observed trend between the reduced χ^2 squared goodness-of-fit of recovered bulk flows to the expected moments of the bulk flows from a set of the Zeldovich mocks with a spherical sky mask, plotted against the kurtosis of the distribution of observed velocities for the set of mocks. Points with the same colour define results for estimators applied to the same set of mocks with a radial selection function that is proportional to a number density of objects specified by the labels on the x -ticks on the top of the plot. Different shaped markers distinguish between results from applying different estimators, as given in the legend. Right panel: density plot of the distributions of observed PVs for mocks with different choices of radial selection function.

For the Kaiser MLE and Peery MVE, I next focus on judging how precisely and accurately these estimators are able to constrain the bulk flow using these methods on the data.

Kaiser MLE method: results for CF4 mocks

I applied the Kaiser MLE to 512 CF4 mocks. I also tested the results when a distance limit in the mock data was introduced at various cut-off scales r_c of between 35 and 345 $\text{Mpc } h^{-1}$, such that objects included involve only those with distances $\leq r_c$. This was because for the real CF4 data, we want to investigate the estimated bulk flow when data is included at different scales and this required us to validate the performance of the estimator on mocks with the same distance limits applied. The ‘truth’ bulk flow moment, \mathbf{B}_K , which I compare the results to, was modified to incorporate only the PVs of objects not excluded by the cut-off scale when running these tests. In general I found for any cut-off scale or when all the data is included, the estimated bulk flow is strongly correlated to the truth bulk flow. On average however, the bulk flow is slightly underestimated compared to the expected moment of the data, and this is apparent from the slopes of the best fit linear regressions in Figure 4.8a which shows the results from the Kaiser MLE applied to the 512 CF4 mocks. In this plot I have used all the data and there is no cut-off scale applied. I found similar results when applying the estimator to the same mocks, but when using the true radial peculiar velocity of each galaxy with zero uncertainty, instead of using observed PVs calculated from observed log-distance ratios. Therefore this bias is due to the estimator and not due to a systematic or error in the observational mock data. The results are shown in plots for the different cut-off scales and for the true radial velocities which can be found here for the interested reader.

Given these results, one can consider correcting the bulk flow that we will obtain from the real data

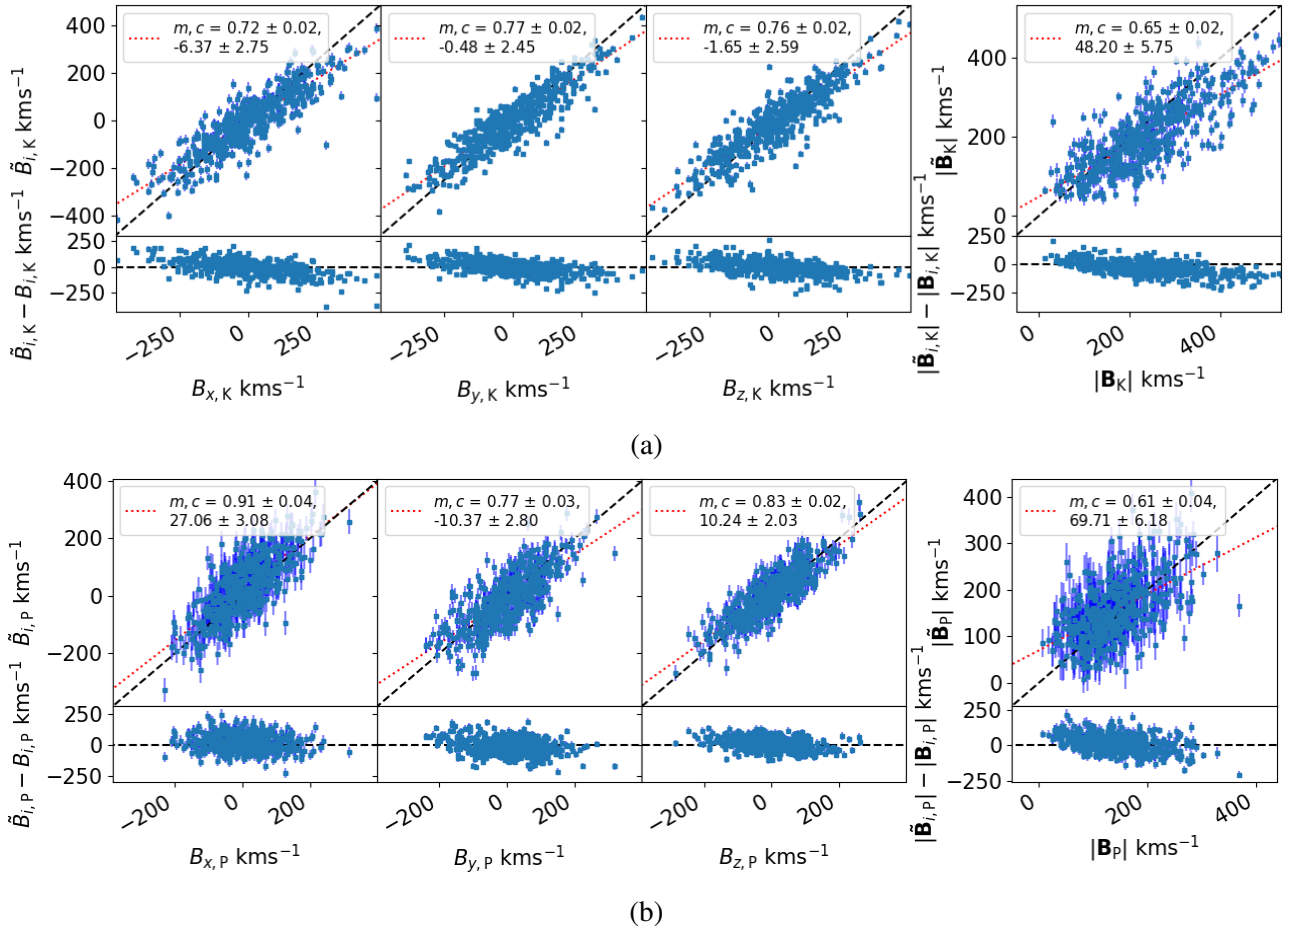


Figure 4.8: a) The recovered bulk flow components $\tilde{B}_{i,K}$ vs the $B_{i,K}$ for the realistic mocks for CF4 when applying the Kaiser MLE method, shown in Supergalactic coordinates. The top panels show $B_{i,K}$ against $\tilde{B}_{i,K}$ directly while the lower panels plot $B_{i,K}$ against the residual $\tilde{B}_{i,K} - B_{i,K}$. In the top panels the black dashed line is simply a 1-to-1 line and in the lower panel the black dashed line is a horizontal line for comparison. The red dashed line in the top panel is the best fit linear regression to the data in each panel, where m gives the gradient and c gives the y-intercept. b) As in a), but when I apply the Peery MVE method.

with the Kaiser MLE method using the mock results, to ensure we can compare the bulk flow that is recovered to a cosmological model. Furthermore, I can use the results from the CF4 mocks where we have used the true radial PVs of each object (as described above) rather than the observational PVs. By averaging the bulk flow from many CF4 mocks and taking the standard deviation of them, this gives us a good representation of what to expect for 1) the bulk flow on a given scale and for a given cosmological model, and 2) for the cosmic variance uncertainty on the bulk flow. By using these for comparison, I can take into account not only how geometry should affect the bulk flow estimate but also any systematics that influence the performance of the Kaiser MLE method, while not being affected by uncertainties in PV observations.

I have also applied the Kaiser MLE to subsets of the CF4 mocks, by applying it individually to the individual SDSS, CF4TF and 6dFGSv mocks that make up the CF4 mocks, with the same cut-off scales applied to the datasets. The results can also be found here. A similar bias with the trend of underestimating the underlying bulk flow also appears to afflict these results. As the bias was

not overtly present in the results for the Zeldovich mocks for SDSS or the more generic zeldovich mocks tested earlier, this might suggest that the presence of the non-linearities in the PV field (for realistic mocks) have some influence on the results after all. This is demonstrated here in Figure 4.9, which shows results when applying the Kaiser MLE to mocks for the SDSS data where non-linear information is included in the peculiar velocity field (top panels) vs when it is not because the Zeldovich approximation is used to generate the mocks only (lower panels).




Figure 4.9: Resulting bulk flow residuals when applying the Kaiser MLE approach to SDSS mock data (256 mocks). The upper panels are mocks that include non-linear structure growth in the PV field, while only Zel’dovich theory for structure growth is included in the PV field for the mocks in the lower panels. The coloured dashed lines show the best fit linear regression in each panel for each coordinate direction, with the legend giving the slope and y-intercept for the regression. The black line gives a horizontal line for comparison. It should be noted here that the mocks do have slightly differently defined projections onto the x , y and z axes for the observer which may affect the scatter seen in the measured bulk flows due to how the coordinate directions point relative to the direction of the SDSS data.

Interestingly, the bias appears to possibly be more significant in the mocks with the CF4TF or 6dFGSv data only, which includes measurements at lower redshifts, where we might expect the systematic error in each measurement of the log-distance ratio to be less significant (due to the way uncertainty increases with more distant objects), thus allowing errors in modelling to dominate.

Furthermore, as SDSS has the largest amount of data more non-linear information can be smoothed out. While I have discussed that the non-linear PVs do not appear to correlate with the χ^2 goodness-of-fit from applying the Kaiser MLE, they are still potentially a source of systematic bias, because the Kaiser MLE method does not account for non-linearities in the PV field. Alternatively, the choice of $\sigma_* = 300\text{kms}^{-1}$ requires variation. This could be tested by applying a MCMC approach or the η MLE approach shown in Qin et al. (2018), which might even have the effect of improving the χ^2 , but I leave this to future work to determine.

Furthermore, the uncertainty due to the data for each estimate appears significantly underestimated as I obtain a reduced $\chi^2 \sim 23$ in general. I also find if a correction is applied to the bulk flows to correct the bias, the reduced χ^2 is ~ 24.8 (where I simply correct $\mu = \mathbf{x} - \mathbf{x}_t$, where \mathbf{x} represents the vector of bulk flows compared to the truth bulk flows \mathbf{x}_t) in the equation for $\chi^2 = \mu \mathbf{C}^{-1} \mu^T$, where \mathbf{C} is the covariance matrix. Therefore I follow the approach of Howlett et al. (2022) and apply a scaling to the covariance matrix \mathbf{C} for $\tilde{\mathbf{B}}_K$ for the uncertainty on the real CF4 data when applying the Kaiser MLE method.

Peery MVE method: results for CF4 mocks

Table 4.2: Results for the reduced χ^2 goodness-of-fit for the bulk flow when applying the Peery MVE for 512 mocks for the CF4 data, compared to the chosen radius R of the ideal survey used for the method.

$R \text{ Mpc } h^{-1}$	69	104	138	173	207	242	276	311	345
χ^2	6.51	5.07	4.23	3.78	3.54	3.43	3.39	3.39	3.39

I also apply the Peery MVE method to 512 mocks. I also wanted to test the results on the mocks when the radius of the ideal survey for this method was varied between scales of 35 and 345 $\text{Mpc } h^{-1}$ (for clarity here, I do not apply any cuts to the data but only change the ideal survey geometry, unlike the cutoff scale introduced when testing the Kaiser method on these mocks). Likewise, this is to look at the performance on mocks before applying the same procedures to the real data. This is also what was done in Watkins et al. (2023) for the CF4 dataset, but with different choices for the ideal survey radii. In general I found the results were strongly correlated to the expected bulk flow moments of the mocks, regardless of the choice of the radius of the ideal survey, although the correlation between the $B_{i,P}$ and $\tilde{B}_{i,P}$ is not as strong as in the application of the Kaiser method. Likewise, there is a similar bias in that the estimated bulk flow is an underestimate of the expected bulk flow moment. However the reduced χ^2 for each set of mocks is much closer to unity compared to the Kaiser result and varies depending on the radius of the ideal survey. The reduced χ^2 is found to obtain closer values to unity at larger radii for the ideal survey; the relationship between these is summarized in Table 4.2. In the case the biases present in the mocks are corrected, I find for the case the ideal survey radius is set to $173\text{Mpc}h^{-1}$, that $\chi^2 \sim 4.48$. Results of applying the Peery MVE to 512 mocks with the ideal survey radius set to $173 \text{ Mpc}h^{-1}$ are shown in Figure 4.8b. The results here (and in the Kaiser approach)

will be used to scale the error bars appropriately when applying this method to the real data, which is described at the beginning of section 4.5.1.

Effects of systematic zero–point offsets

Before finally applying estimators to the real data, I explore the effects of systematic offsets in the zero–point of datasets on the recovered bulk flows from the Kaiser method and the Peery MVE method. I explore both a global zero–point offset, in which case each log–distance ratio in the dataset is offset by a constant amount, σ_η and a zero–point offset between the three main subsamples of CF4; SDSS, 6dFGSv and CF4TF. A global offset in the zero–point is analogous to a change or error in the Hubble constant H_0 . It may also be possible to have an offset due to a calibration error when fixing the zero–points of datasets that have been combined.

Given the constraint equation that is included by Peery et al. (2018) in order to ensure the results of the Peery MVE are independent of the Hubble constant, we can expect that this method should obtain a bulk flow measurement that is unaffected by a global zero–point offset in the data. We found that by applying the Peery MVE to the mocks with a deliberately added or subtracted global zero–point offset of $\sigma_\eta = 0.031$ (corresponding to a change of approximately $5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ on H_0) the estimated bulk flow remains unbiased, validating this expectation.

However, we should have no such expectation when applying the Kaiser MLE estimator and find that the results that are presented (in Supergalactic coordinates, to match the coordinates that the data results are presented in) show significant biases when a global zero–point offset is present; this is shown in Figure 4.10. In particular the Supergalactic y direction, which is the direction of the SDSS cone of data, is affected the most significantly. The Supergalactic x and z directions are also visibly affected due a lesser extent. Due to the fact a constant shift in each η does not translate to a constant shift in each PV, the effect of a positive or negative σ_η is not symmetric. Overall the Bulk Flow amplitude appears to be shifted for the x and y directions by upward of 100 kms^{-1} .

I also explore the effect of a relative zero–point offset between the different datasets that make the CF4 mocks; because the CF4 mocks are composed of individual mocks made from SDSS, CF4TF and 6dFGSv this is straight forward to implement. I choose a systematic offset of $\sigma_\eta = \pm 0.017$ between different datasets; this value is chosen based on the statistical uncertainty found in the calibration procedure of different datasets in Tully et al. (2023). In Figure 16 of Tully et al. (2023) work, the $1-\sigma$ statistical uncertainty on the relative zero–point offset between the FP and TF datasets with the SNe Ia used to calibrate the entire CF4 dataset is five times smaller than this chosen value; I choose a larger value in order to simulate the effect of a significant (5σ) systematic error. I first apply this offset to the CF4TF mock data relative to all other data and then to the 6dFGSv mock data relative to all other data. I find that when applying these offsets there is a visible bias for the recovered bulk flows (as expected) when applying both the Kaiser MLE and the Peery MVE. Interestingly the Peery MVE seems to be more greatly affected by a relative offset between datasets than the Kaiser MLE estimator. Figure 4.11 shows the effect of a relative zero–point offset added to the CF4TF data in the CF4 dataset and how this affects the results from the Peery MVE. More of these results can be found here. Of all individual




Figure 4.10: The recovered bulk flow moments $\tilde{B}_{i,K}$ and the residuals $\tilde{B}_{i,K} - B_{i,K}$ in Supergalactic coordinates, compared to the expected truth bulk flow moments $B_{i,K}$ for the CF4 mocks, when applying the Kaiser MLE. The dashed line shows a 1-to-1 line in the plots showing $\tilde{B}_{i,K}$ vs $B_{i,K}$, and a horizontal line for the residual plots to compare to the results. The red dashed lines show best fit linear regressions for $\tilde{B}_{i,K}$ vs $B_{i,K}$ in each panel. The legend indicates the gradient m and the y-intercept c for each regression. The blue points show the results when there is a global zero-point shift of $\sigma_\eta = -0.031$ applied to each object log-distance ratio in the mocks. For the red points there is a global zero-point shift of $\sigma_\eta = +0.031$ applied to each object. The change in the linear regressions in columns of panels allow us to measure the affect of the change in the zero-point on the results.




Figure 4.11: The recovered bulk flow components $\tilde{B}_{i,P}$ and the residuals $\tilde{B}_{i,P} - B_{i,P}$ in Supergalactic coordinates, compared to the expected truth bulk flow moments $B_{i,P}$ for the CF4 mocks, when applying the Peery MVE. The dashed lines shown are the same as described in Figure 4.10. For the blue points there is a relative zero-point shift of $\sigma_\eta = -0.017$ applied to CF4TF object log-distance ratios in the mocks only. For the red points there is a relative zero-point shift is $\sigma_\eta = +0.017$.

offsets I test with different estimators, the largest bias is apparent in the Supergalactic y direction when a global offset is added to the CF4 mocks and the Kaiser MLE approach is used to recover the bulk flow, shown in Figure 4.10.

4.4.4 Summary: mock performance evaluation

To summarise this section, we found the following;

- The Watkins MVE method and Kaiser MLE are able to recover the expected bulk flow moment of a dataset accurately and with good precision, when the bulk flow is a constant valued vector across the survey volume.
- In general however, these estimators may not measure the bulk flow moment with the same accuracy and precision for realistic mocks, and the measured moment is not in general consistent with Equation 4.1. This issue was first pointed out by Nusser (2014). I demonstrate the precision of the estimators relative to the expected bulk flow moments thoroughly here with mocks.
- I show there is a strong dependence between the survey geometry (in particular the selection function) and the χ^2 goodness-of-fit of these methods.
- Both Nusser (2014) and the Peery et al. (2018) introduce modifications into the Kaiser MLE method and Watkins MVE method to address this issue. The Nusser MLE and the Peery MVE method appear to give a better fit to the expected bulk flow moment, that is consistent with Equation 4.1, with a χ^2 closer to unity for different survey geometries we have tested and also reduce the correlation between survey geometry and the χ^2 . However the Nusser MLE modification depends on the assumption of spherical symmetry, which was found to be unsuitable for the CF4 dataset.
- For mocks with highly non-linear PVs the PV distribution changes and this appears to introduce a bias to the recovered bulk flows from any estimator tested. This may be more apparent in the case the systematic errors in the observed PV are smaller, there are less PV observations, and error due to modelling only linear-theory velocities begins to dominate the measurement. Therefore a more complete solution may need to be derived in future work to be able to estimate bulk flows from data that are both comparable to theory and which generally obtain an accurate estimate independent of the survey geometry or the peculiar velocity distribution. This task is left to future work.

When using the estimators tested in this work, it is suggested that they *always* need to be applied to realistic mocks before applying them to estimate the bulk flow of a real dataset and compare the results to theory. This is because this work has shown one needs to quantify how well the calculated error for the bulk flow due to the data is actually capturing the uncertainty that should be assigned to the bulk flow, and test for the presences of biases. If a bias is present, a correction to the bulk flow can

be made before comparing the measurement to theory. For this reason I have applied the Kaiser MLE and Peery MVE to realistic mocks for the CF4 dataset, which will allow us to draw a more accurate conclusion about 1) the uncertainty on the measured bulk flow vector using these methods and 2) the level of tension or agreement the bulk flow has with theoretical expectations. Based on our results, a correction will be applied to the measured bulk flows from these estimators using the CF4 data and the uncertainties on the measurement will be scaled to recover a reduced χ^2 of 1.

For measurements applied to the real data, one may be concerned about the effects of selection bias or Malmquist bias which are well described in Strauss and Willick (1995) and summarised in Appendix A.2. As mentioned previously, galaxies are placed at their observed redshifts in order to determine their distances in all the mocks tested and the same approach is taken to get galaxy distances for the real data. Using the galaxy redshifts as opposed to using a distance indicator to determine the distances may cause the measured bulk flows to suffer from selection bias, and less so from Malmquist bias. However, since the realistic mocks for CosmicFlows–4 include both the selection function and clustering for the real data, conclusions drawn from applying the estimators to the mocks should also be valid for the real data.

4.5 Measured bulk flows from real data

In this next section I apply the Kaiser MLE scheme and the Peery MVE scheme to the CF4 dataset and compare the measured bulk flows to the expectation of the Λ CDM model, which is deduced from bulk flow results of the mocks for the CF4 data where I have applied these schemes using the true radial PVs of each object in the mock with zero–uncertainty. I note here an analysis of the CF4 dataset using the Peery MVE appeared in preprint in Watkins et al. (2023) while I was completing this analysis. However I still present the results here for both the estimators as we expect the conclusions regarding tension with the Λ CDM model may differ. While we can expect to obtain a similar bulk flow amplitude, I argue the error bars presented for the result in Watkins et al. (2023) are underestimated given the results from applying the estimator to mocks, and that their corresponding claim of a tension with Λ CDM is thus slightly overstated.

4.5.1 Results: CosmicFlows 4 dataset

Table 4.3 shows the estimated bulk flow components for the CF4 dataset when applying the Peery MVE and the Kaiser MLE, in Supergalactic coordinates. The error bars for each component due to the uncertainty in the data have been scaled to account for how well the bulk flow estimators perform on mocks. This is done so that the reduced χ^2 will rescale to unity, as $C_{ij} \rightarrow \beta C_{ij}$ where β is the rescaling factor set by the reduced χ^2 from the mocks, in order to inflate the error on the covariance matrix of

Table 4.3: Estimated bulk flow components in Supergalactic coordinates from the Peery MVE and the Kaiser MLE. d_e gives the effective depth of each bulk flow estimate, and $P(> \chi^2)$ gives the probability of obtaining a larger χ^2 goodness-of-fit for the measured bulk flow with respect to the expectation from the Λ CDM model. $\langle |\mathbf{B}| \rangle$ gives the mean bulk flow estimated from 512 CF4 mocks when the true radial peculiar velocities with zero uncertainty have been used for each object in the mock, as a proxy for comparison to the expectation to Λ CDM, and the error bar represents the standard deviation of this measurement. The error bars include both error due to each object’s PV uncertainty and cosmic variance. The columns for results that specify ‘corrected’ indicate the bulk flow modes and amplitude have been corrected for bias based on the simulations for the estimators. For all columns, the numbers in brackets for error bars and probabilities indicate the result if no scaling is applied to the systematic uncertainty based on the χ^2 from the mocks.

	Kaiser MLE	Peery MVE	Kaiser MLE (corrected)	Peery MVE (corrected)
$\tilde{B}_x \text{ kms}^{-1}$	-281 ± 164 (153)	-328 ± 100 (83)	-382 ± 165 (153)	-391 ± 104 (83)
$\tilde{B}_y \text{ kms}^{-1}$	37 ± 148 (142)	-102 ± 89 (75)	48 ± 149 (142)	-119 ± 93 (75)
$\tilde{B}_z \text{ kms}^{-1}$	-105 ± 152 (140)	-94 ± 119 (108)	-135 ± 154 (140)	-126 ± 122 (108)
$ \tilde{\mathbf{B}} \text{ kms}^{-1}$	302 ± 164 (153)	357 ± 104 (87)	408 ± 165 (153)	428 ± 108 (87)
$d_e \text{ Mpc } h^{-1}$	49	173	49	173
$P(> \chi^2)$	32.7% (26.6%)	0.71% (0.06%)	10.1% (6.59%)	0.11% (0.002%)
$\langle \mathbf{B} \rangle \text{ kms}^{-1}$ (Λ CDM)	130 ± 54	128 ± 59	196 ± 82	194 ± 86

the PV uncertainties.⁵ C_{ij} here refers to the covariance matrix component due to uncertainty on the PVs only and not due to cosmic variance uncertainty. Uncertainties in brackets are given to indicate what the uncertainty is without any scaling applied.

It should be noted that the uncertainty due to cosmic variance (CV) may be imprecise because it is calculated based on linear theory, yet we can understand that bias may arise in the results due to a lack of modelling for non-linear structure growth. One may consider to instead estimate the CV covariance of the bulk flows from the (Λ CDM) mocks. However, we expect this to slightly underestimate the true CV uncertainty for the Kaiser method because the mocks do not include a small portion of lower redshift objects such as Ia SNe, SBFs, etc. that are included in the CF4 data,⁶ and thus will obtain a CV uncertainty corresponding to a slightly lower effective depth. However we find that the linear theory CV covariance matrix and the CV covariance of the mocks only differs by a factor of ~ 1.2 for the Kaiser method. Likewise for the Peery MVE, we find that the linear theory CV covariance matrix is also similar to the CV covariance matrix from the mocks, only differing by a factor of ~ 0.7 . Therefore I simply use the linear theory results to calculate the total error on the measured bulk flows from the data and for comparison to the Λ CDM model. Nonetheless, the mocks form a useful basis for comparison, so in table 4.3 I also include $\langle |\mathbf{B}| \rangle$, the mean bulk flow estimated from 512 CF4 mocks

⁵An alternative approach to scaling the statistical uncertainty may be to instead attempt to add a systematic error. However, our analyses generally found little evidence of biases in tests on mocks. Furthermore, I have pointed out earlier that the difference between estimators cannot be compared because they estimate different quantities, and they do not generally probe the same effective depths.

⁶As a reminder, it was mentioned previously that the mocks used for the CosmicFlows-4 dataset are created by stacking mocks that were produced for the SDSS data, CF4TF data and 6dFGSv data that largely capture the survey geometry and data included in the CosmicFlows-4 catalogue.

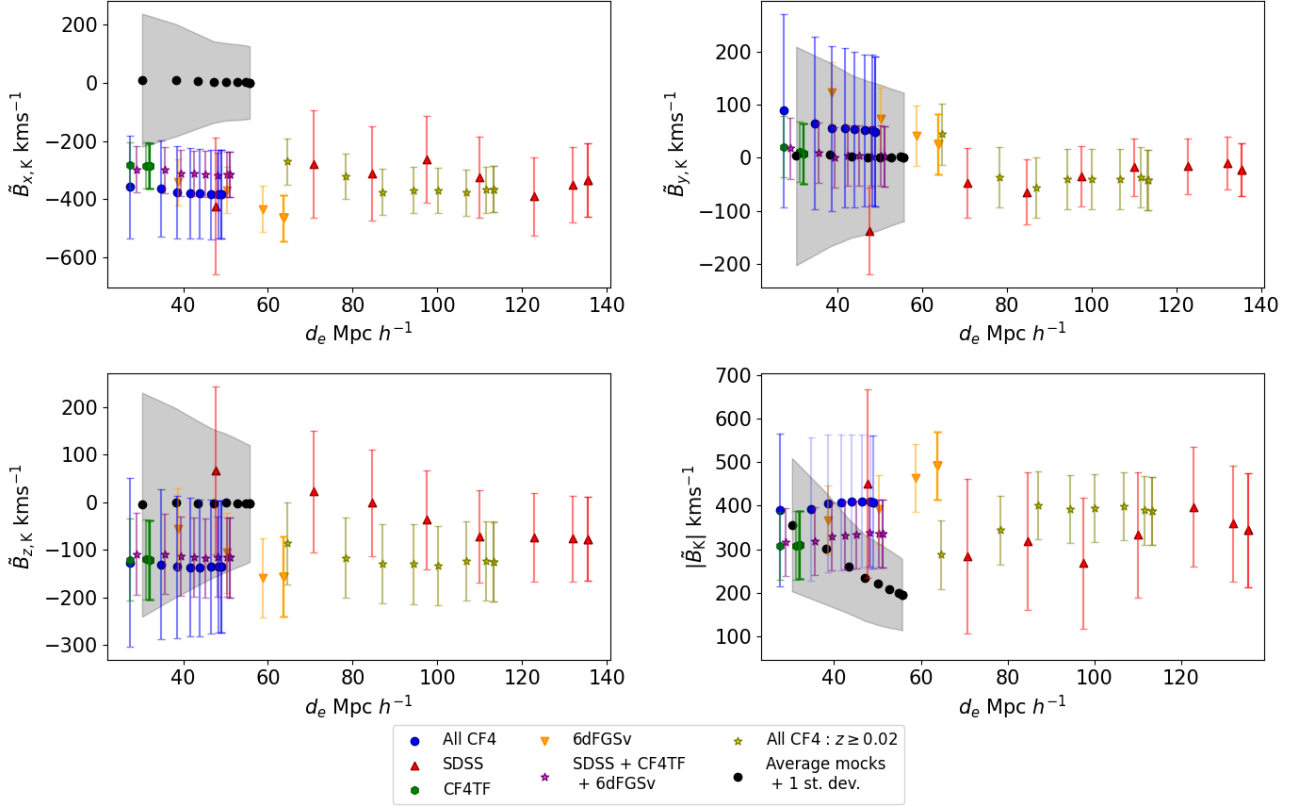


Figure 4.12: Estimated bulk flow components in Supergalactic coordinates and bulk flow amplitude from the CF4 data and subsets of it, using the Kaiser MLE. The bulk flow has been estimated at various different effective depths (the same effective depths are plotted for each panel) by applying cuts to the data at different radii and calculating the effective depth. The error bars do not include cosmic variance for clarity. The different data subsets are shown in the legend. The black points with grey shading show the average recovered bulk flow from 512 mocks $\langle \mathbf{B}_K \rangle$ with the same selection function as SDSS+CF4TF+6dFGSv, and with the true radial PVs of the objects fed as input. The shaded region shows the standard deviation of $|\mathbf{B}_K|$ which gives an estimate of the theoretical cosmic variance uncertainty for the mocks.

when the true radial peculiar velocities with zero uncertainty have been used for each object in the mock.

To compare our measured bulk flows to the Λ CDM model, I take the linear theory covariance and use it to compute the χ^2 of our measurements and have listed the probability of finding a larger χ^2 in Table 4.3. For the effective depths in the table, I use the mean distance to each object weighted by $(\sigma_n^2 + \sigma_*^2)^{-1}$ for the Kaiser MLE, while for the Peery MVE I use the radius of the ideal survey for the estimator, which is set to $\sim 173 \text{Mpc} h^{-1}$ in our results. Furthermore, I have added to the table the expected bulk flow components and amplitude after the bias in the mocks has been corrected in additional columns. I also apply the relevant correction to the Λ CDM expectation calculated from the mocks, $\langle |\mathbf{B}| \rangle$, for comparison.

It is interesting that both estimators give significant bulk flow magnitudes in the negative Supergalactic x direction. While the amplitudes of each measurement from the different approaches are in agreement within statistical uncertainty the Kaiser result would indicate there is no tension with the Λ CDM model. In contrast, the result from the Peery MVE, which is able to probe the bulk flow

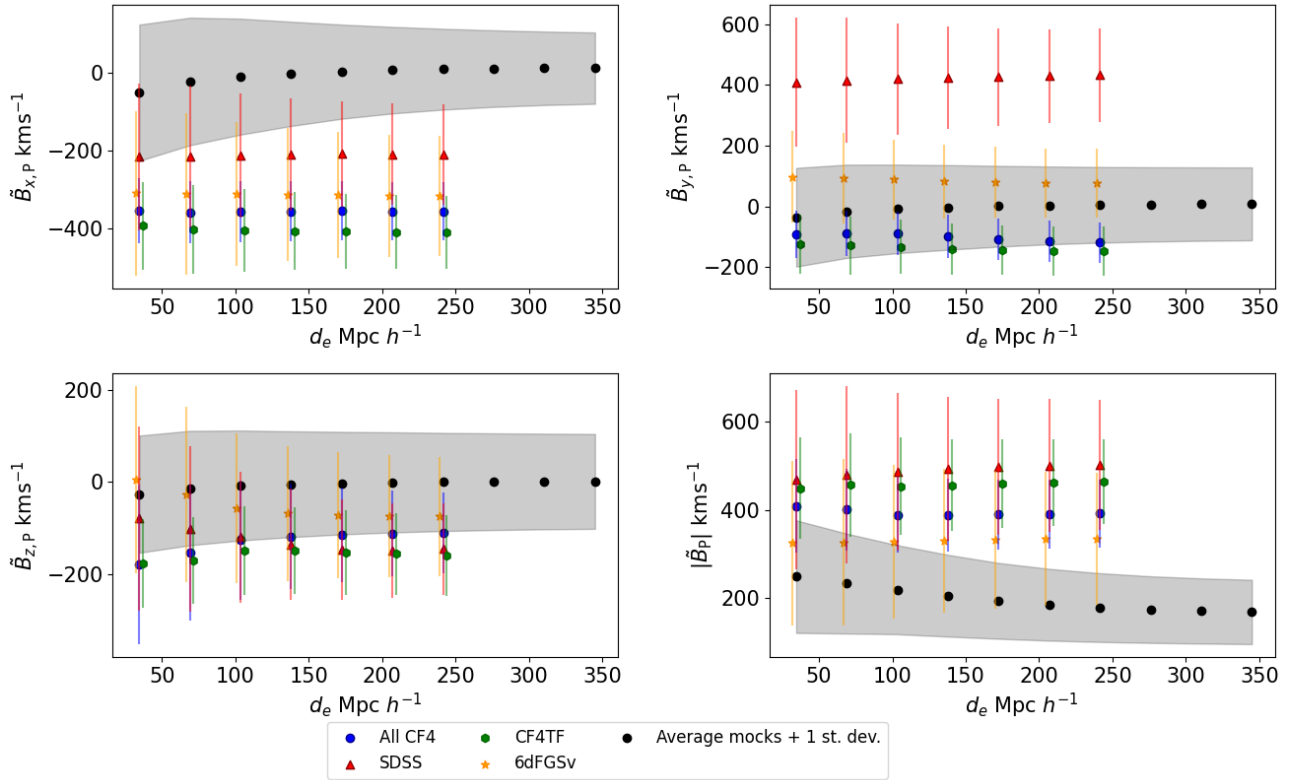


Figure 4.13: As in Figure 4.12, but for the estimated bulk flow components from the Peery MVE and where the different effective depths for estimates arise by changing the radius of the ideal survey for this estimator instead of cutting the data at different scales. A small horizontal offset is added to the CF4TF data points and 6dFGSv data points about their correct value for d_e so that datapoints do not overlap.




Figure 4.14: A map of the CF4 dataset in Supergalactic coordinates, each panel showing a 2D projection of two coordinates. The arrows positioned at the origin show the direction of the bulk flow (with relative lengths given by the amplitude of the bulk flow in each direction, scaled by the effective depth of the measurement) for the Peery MVE and Kaiser MLE. I also show the approximate positions of Shapley (Proust et al., 2006) and the hypothesised Vela Supercluster (Kraan-Korteweg et al., 2017).

at a larger effective depth indicates there is tension with Λ CDM if we set the significance level at 5%. While I find a similar magnitude for the bulk flow to Watkins et al. (2023) the tension found with the Λ CDM prediction is slightly less significant after correcting the error bars and the bias in the results; I find the bulk flow has a 0.11% chance of occurring rather than only the 0.0002% or less they determine.

To try and investigate the origin of this, Figure 4.12 shows the result of applying the Kaiser MLE to the CF4 mocks (and subsets of the CF4 data with the same zero-point calibration) with the same radial cuts as was applied to the mocks. No bias correction has been applied in the results shown, but the values for each CF4 result (blue points) can still be fairly compared to the expected values given from the mocks without bias correction.

One may note the effective depth for the mocks with the same radial cuts (shown in every panel on the x -axis) are generally larger than for the real data; this can be attributed to the fact the mocks do not include data for SNe Ia or other low redshift objects used to calibrate the dataset, as mentioned previously. As these objects typically have PV measurements with smaller errors they carry larger weight for the Kaiser MLE approach. The lower redshift CF4TF data also influences the effective depth of the CF4 data enough for it to sit at lower depths than is obtained for the SDSS data or 6dFGSv data alone; this is seen by comparing the results for all CF4 data to the subset made up of 6dFGSv + SDSS + CF4TF only (purple stars) and then the SDSS subsets (red triangles) and 6dFGSv subsets (orange triangles). The subset of CF4 with only $z \geq 0.02$ objects included has a significantly larger effective depth as much of the contribution of CF4TF catalogue is removed.

Almost all of the data in Figure 4.12 shows unexpected trends for the bulk flow magnitude as a function of depth given the expectation from Λ CDM, which should be similar to what we see in the mock results. Instead of the bulk flow amplitude decreasing with depth, it increases, particularly in the case of the 6dFGSv subset. The clearest feature is that all subsets of data I test prefer a large negative bulk flow in the Supergalactic x -direction at all effective depths, which is the primary driver of the large bulk flow amplitude.

Figure 4.13 shows a similar plot to Figure 4.12 however this time with the Peery MVE. For the CF4 data, the bulk flow amplitude has been evaluated at significantly different scales to the Kaiser MLE but the resulting components and amplitude are not significantly different. Interestingly however, this is with the exception that the Peery MVE obtains a much larger bulk flow amplitude in the positive Supergalactic y direction for the SDSS data, although the direction is negative for the entire CF4 data. We may expect that is related to the fact that SDSS has most of its data in the positive Supergalactic y direction of the sky in a cone-like geometry. In contrast, the Peery MVE seems to give a more reasonable estimate for the bulk flow of the 6dFGSv data alone, compared to the larger bulk flow amplitude of 6dFGSv estimated to be close to $\sim 500\text{kms}^{-1}$ from the Kaiser MLE at the largest scales. It is apparent that at larger depths for the CF4 dataset, at which we expect the estimated bulk flow to be more accurate given the mock results for the Peery MVE, the resulting bulk flows seems to diverge from the expectation of the Λ CDM model (the black points and grey shaded regions). Although, individual datasets present differences in their bulk flows at different depths between the

Kaiser MLE and Peery MVE estimators, the overall bulk flow amplitudes are remarkably consistent in their deviation from Λ CDM.

Figure 4.14 shows a map of the CF4 data and the direction of the estimated bulk flows given in Table 4.3 (a bias correction is applied in the data shown). The results presented in this work would indicate that in general the bulk flow is headed towards the hypothesised Vela Supercluster (Kraan-Korteweg et al., 2017) close to the general region of the ‘Great Attractor’ (Radburn-Smith et al., 2006), and with some pull in the direction of the Shapley supercluster. The former of these hides out of sight in the Zone of Avoidance, where the Milky Way disk inconveniently blocks the view of extragalactic objects, while the latter remains only partially covered by current data in the southern hemisphere.

The strange trends we see in the data could arguably be a symptom of a systematic in the zero-point offsets. However we have noted the bulk flow amplitudes and values of the bulk flow moments measured from the Kaiser MLE and the Peery MVE method are similar. We might expect if a global zero-point offset was present in the data and no relative zero-point offsets were present, the Kaiser MLE and Peery MVE would obtain differing amplitudes because we know the Peery MVE is effectively immune to this systematic via the additional constraint equation it incorporates to minimize uncertainty due to H_0 . However, this would not rule out the possibility of errors due to an incorrect relative zero-point offset between the datasets that make up the CF4 data because they introduce a bias into the recovered bulk flow from each method in different ways. We might expect if a relative zero-point offset error affected the data in such a way that the recovered bulk flow mode value in the negative Supergalactic x direction was overestimated, accounting for this error would make the bulk flow measurements more consistent with the expectation from Λ CDM. A zero-point offset in the CF4TF data of $\sigma_\eta = -0.017$ relative to SDSS and 6dFGSv appears to introduce a bias of $\sim -189\text{kms}^{-1}$ in the Supergalactic x direction for the Peery MVE, which would work in the correct direction to account for the bulk flow in the data, and would also reduce the bulk flow amplitude by roughly the same amount in the SuperGalactic y direction and by a smaller amount in the z direction, and thus would reduce the tension seen in our results with Λ CDM. The same offset would also mean that the bulk flow measured by the Kaiser MLE would require a correction such that the bulk flow amplitude would become lesser in the x and z directions and slightly larger in the y direction, although the required corrections are considerably smaller. However, this scenario does require a zero-point offset between datasets at a level of 5σ compared to the reported calibration uncertainty in CF4. Overall, while we cannot rule out internal zero-point offsets systematically enhancing the bulk flow amplitude, these would have to be very large, unnoticed, and applied in a particularly contrived way to affect both estimators in the right directions by the right amount.

Alternatively, we could also consider the findings in Heinesen (2023) which shows that by neglecting the importance of relativistic effects in bulk flow measurements, the bulk flow amplitude may be overestimated by a factor of $\sim (1 + 1.55z)$, where z here represents the redshift of the effective depth of the bulk flow measurement. This may also explain the trend of increasing bulk flow amplitude with effective depth seen in the results, and would lead to an overestimate of approximately $\sim 10\%$

for the Peery MVE result and $\sim 2.5\%$ for the Kaiser MLE result. However, even after applying this correction, the Peery MVE result is still 389 km/s, which is still in tension with Λ CDM with only a $\sim 0.4\%$ chance of measuring the observed bulk flow.

Finally, one may consider that there is some possibility the tensions seen in bulk flow measurements may be related to the H_0 tension (see references within Di Valentino et al., 2021) and (Aghanim et al., 2020; Heymans et al., 2021) σ_8 tension that currently pervade cosmology. The theoretical bulk flow variance is defined as (Andersen et al., 2016)

$$\sigma_V^2(\mathbf{r}) = \int \frac{dk^3}{(2\pi)^3} P_{vv}(k) |\tilde{W}(\mathbf{k}; \mathbf{r})|^2. \quad (4.15)$$

which are sensitive to H_0 and σ_8 . $P_{vv}(k)$ is the velocity power spectrum while $\tilde{W}(\mathbf{k}; \mathbf{r})$ is the Fourier space window function of the galaxy survey. The most probable bulk flow amplitude can be estimated as $V_p(R) = \sqrt{3/2} \sigma_V(R)$, and thus has dependence on σ_8 and H_0 through the power spectrum. However, altering the power spectrum to vary H_0 and σ_8 within the range seen by tensions only changes the resulting bulk flows in the Λ CDM model by small numerical factors; varying these parameters is not able to solve the tensions seen in bulk flow measurements. Furthermore, we would still expect even after allowing these parameters to vary that the bulk flow amplitude should approach zero at increasingly larger scales.

4.5.2 Comparison to bulk flows from previous literature

Finally I show a comparison of our bulk flow estimates to some of those previously found in the literature. This is shown in Figure 4.15, which includes results taken from Ma and Pan (2014); Hong et al. (2014); Qin et al. (2018, 2021); Watkins et al. (2009); Feldman et al. (2010); Howlett et al. (2022); Watkins et al. (2023); Peery et al. (2018); Scrimgeour et al. (2016); Watkins and Feldman (2015b); Kashlinsky et al. (2008); Ma et al. (2011); Nusser and Davis (2011); Hoffman et al. (2015). The results are shown in order of increasing depth. It is interesting to note that there is a consistently large bulk flow detected in the negative Supergalactic x and z directions, despite the fact different data sets, methodologies and approaches to calibrate the zero-point have been used. However, one should consider that the results between datasets will also be correlated because they all contain some overlap in the data included. Interestingly, the effective depth of each measured bulk flow varies but the bulk flow appears to be consistent in the amplitude and direction. As stated previously, it is interesting to note that bulk flow measurements reported at larger effective depths are those that tend to find more tension with Λ CDM and our measurements of the CF4 bulk flow are consistent with this trend. Although not included in Figure 4.15, Migkas et al. (2021) also finds an apparent spatial variation in H_0 from galaxy cluster scaling relations which could be interpreted as a bulk flow of amplitude $\sim 900 \text{ km s}^{-1}$ at depths of ~ 500 Mpc which would also coincide with similar directions to the other bulk flows in the literature.

The results in this work would indicate and support previous results in the literature that galaxies are headed towards the Greater Attractor or Vela Supercluster region in the negative Supergalactic x

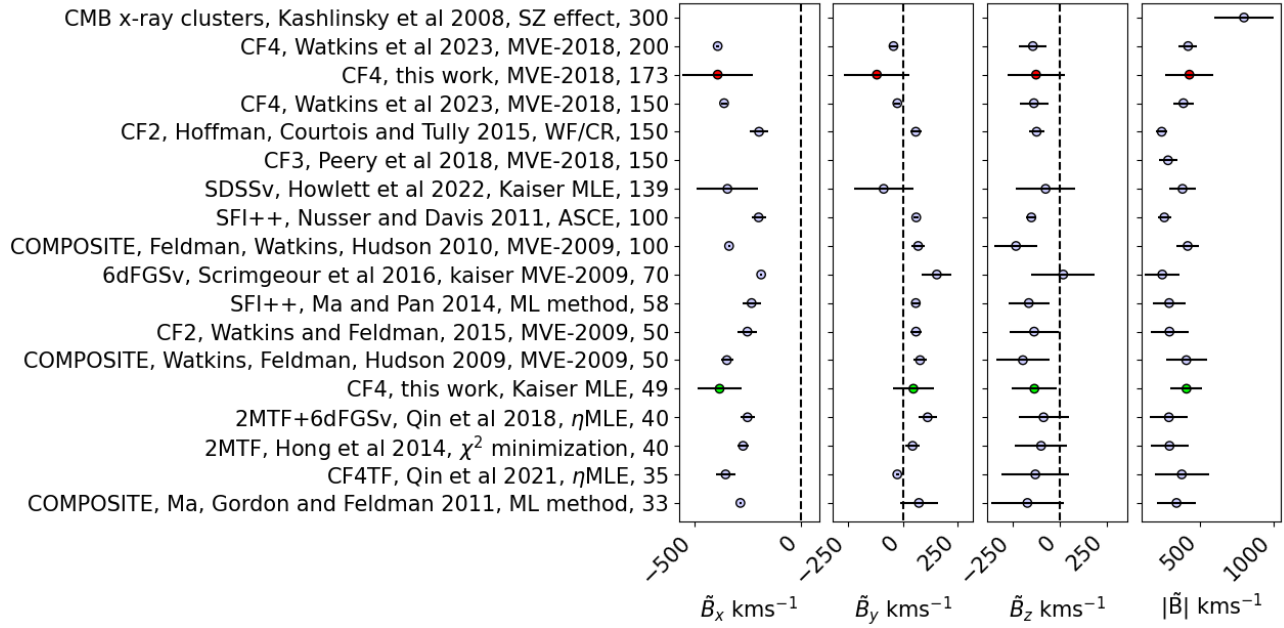


Figure 4.15: Comparison of different bulk flows of datasets shown in the literature, converted to Supergalactic coordinates and listed in order of decreasing effective depth. In cases where data is missing coordinates or bulk flow components were not explicitly listed in writing, only the bulk flow amplitude is shown. Error bars given for $|\tilde{\mathbf{B}}|$ are exact, except for when the uncertainty due to cosmic variance was not given and I have included a rough estimate from theory for a perfect sphere with radius corresponding to the effective depth given. To obtain error bars for each Supergalactic coordinate mode I have assumed the correlation between modes in other coordinate systems was zero, as in general the covariance matrix was not available. For asymmetric error bars I took the mean of the two uncertainties to perform the transformation of the uncertainties to different coordinates. It should be noted that in general authors do not always show uncertainties in a consistent way depending on whether they used mocks to account for underestimated uncertainties by estimators or accounted for the survey geometry when computing the theoretical cosmic variance uncertainty. On the left hand side of the plot I list the dataset used, the authors, the method employed to obtain the measurement (MVE–2018 is short for Peery MVE, MVE–2009 for Watkins MVE, SZ effect for Sunyaev–Zeldovich effect, ML for maximum likelihood estimation, WF/CR for Wiener Filter/Constrained realizations, ACSE for the ‘all space constrained realizations’ method), then d_e in $\text{Mpc } h^{-1}$.

direction. This strongly indicates that future surveys aimed towards the negative Supergalactic x direction (primarily the southern equatorial hemisphere) are needed in order to gain a better understanding of bulk flow measurements and their tension with ΛCDM .

4.6 Conclusions

In this work I have discussed the potential pitfalls on bulk flow estimators, and thoroughly evaluated how well the Kaiser maximum likelihood estimator (Kaiser MLE; Kaiser, 1988), Nusser MLE (Nusser, 2014), minimum variance estimator (MVE) by Watkins et al. (2009) and the Peery MVE (Peery et al., 2018) are able to estimate the bulk flow, using mock data. This work has built on analysis with mock data already conducted for these estimators by others including Agarwal et al. (2012); Nusser (2014); Andersen et al. (2016); Qin et al. (2018, 2019, 2021); Howlett et al. (2022).

Overall, the main significant finding of the tests on mock data is that both the Kaiser MLE and Watkins MVE recover estimates of the bulk flow with a goodness-of-fit typically corresponding to a reduced $\chi^2 > 1$, which is strongly correlated with survey geometry. This implies that these estimators are not able to capture the real systematic uncertainty in the bulk flow. I found that the modified weighting scheme to the Kaiser MLE scheme by Nusser (2014) appears to perform better at estimating the bulk flow (the χ^2 is closer to unity and the correlation between the survey geometry and χ^2 appears to approach zero), but it is not suitable for surveys that do not have spherical symmetry in the selection function, which is the case for most realistic datasets. In contrast the Peery MVE also obtains a χ^2 closer to unity in general and does not rely on the assumption of spherically symmetric survey geometry; arguably, this estimator is the most robust and reliable. This work demonstrates the interesting result that when poorer assumptions about the nature of the bulk flow are encoded in the estimators (as it the case for the Kaiser MLE and Watkins MVE) the χ^2 goodness-of-fit has a strong correlation to the survey geometry. However, tests on the Kaiser MLE, Watkins MVE and Peery MVE show these estimators are unbiased with respect to complex survey geometries generally. Despite the better performance with the Peery MVE and Nusser MLE, all these estimators suffer from some bias due to highly non-linear PVs and some underestimation of the statistical error remains. Both of these need to be corrected (as is done herein) before comparing a bulk flow measurement to a cosmological model.

My recommendations are that either 1) in future work one could develop an improved bulk flow estimator to overcome the issues identified here or 2) that when estimating the bulk flow, ensure the estimator is always tested on realistic mocks to quantify how well it performs given the survey geometry, to ensure the systematic error bars on the bulk flow estimate are accurate. I have taken the latter approach in regards to making an estimate of the bulk flow of the CosmicFlows-4 data using the Kaiser MLE (Kaiser, 1988) and the Peery MVE (Peery et al., 2018).

I analysed CosmicFlows-4 mock data to determine the real systematic uncertainty in the bulk flow and thus scaled the uncertainties up according to the reduced χ^2 fit from the mocks. Using the Kaiser MLE I found a bulk flow with an amplitude of $408 \pm 165 \text{ kms}^{-1}$, at an effect depth of $49 \text{ Mpc}h^{-1}$ that appears to not be in tension with the Λ CDM model. However our bulk flow estimate from the Peery MVE at a depth of $173 \text{ Mpc}h^{-1}$ is in tension with the model with an amplitude of $428 \pm 108 \text{ kms}^{-1}$, compared to the expected $\sim 192 \text{ kms}^{-1}$ based on Λ CDM mocks. The bulk flow amplitude is in excellent agreement with the measurement by Watkins et al. (2023) who used the same method. However, the tension with Λ CDM is very slightly less severe because we derive larger error bars based on the mock results.

Overall, the results indicate (along with previous consistent results in the literature) that the volume of galaxies in our local region is headed towards the Great Attractor (Radburn-Smith et al., 2006) or the Vela Supercluster (Kraan-Korteweg et al., 2017) regions, with some influence from the Shapley supercluster. To gain a better understanding of our bulk flow and test the Λ CDM model more rigorously, it is necessary to gain more peculiar velocity data towards the negative Supergalactic x direction (southern equatorial hemisphere) and further improve bulk flow estimators. Furthermore,

the results with mocks show the large bulk flow amplitude seen in the CosmicFlows-4 data cannot be trivially explained by systematics related to the survey geometry or the peculiar velocity measurement pipelines. This further supports the recommendation that improved bulk flow estimators should be developed, given that statistical uncertainty in peculiar velocity measurements themselves is unlikely to be reduced.

The following publication has been incorporated as Chapter 5. I have altered the text in this chapter, compared to that of the publication, to use ‘I’ rather than ‘we’ in various places. I will use ‘I’ when referring to results produced solely as my own work, and ‘we’ for results in which there was a combined effort with Hugo to produce fits to BARRY and DESILIKE.

1. Whitford et al. (2025) **Abbé Whitford**, Hugo Rivera-Morales, Cullan Howlett, Mariana Vargas-Magaña, Sébastien Fromenteau, Tamara M. Davis, Alejandro Pérez-Fernández, Arnaud de Mattia, Steven Ahlen, Davide Bianchi, David Brooks, Etienne Burtin, Todd Claybaugh, Axel de la Macorra, Peter Doel, Simone Ferraro, Jaime E. Forero-Romero, Enrique Gaztañaga, Satya Gontcho A Gontcho, Gaston Gutierrez, Stephanie Juneau, Robert Kehoe, David Kirkby, Theodore Kisner, Sergey Kopusov, Martin Landriau, Laurent Le Guillou, Aaron Meisner, Ramon Miquel, Francisco Prada, Ignasi Pérez-Ràfols, Graziano Rossi, Eusebio Sanchez, Michael Schubnell, David Sprayberry, Gregory Tarlé, Benjamin Alan Weaver, Pauline Zarrouk, Hu Zou Constraining the phase shift of relativistic species with DESI BAOs, *Monthly Notices of the Royal Astronomical Society*, 538, 3, 2025.

Contributor	Statement of contribution	%
Abbé M. Whitford	writing of text	95
	proof-reading	35
	numerical calculations	60
	preparation of figures	80
	initial concept	5
Hugo Rivera-Morales	proof-reading	5
	numerical calculations	40
	preparation of figures	20
	initial concept	5
Cullan Howlett	writing of text	5
	proof-reading	30
	supervision, guidance	60
	initial concept	80
Mariana Vargas-Magaña	supervision, guidance	25
	initial concept	10
	proof-reading	10
Sébastien Fromenteau	supervision, guidance	10
Tamara M. Davis	proof-reading	20
	supervision, guidance	5

This project involved measuring the baryon acoustic oscillations in data collected by the DESI Collaboration and constraining a parameter representative of the phase shift due to free-streaming neutrinos. I extended upon an existing BAO-fitting code, BARRY, in order to allow it to constrain the phase shift. Hugo conducted calculations with an existing code provided by Mariana, and Arnaud de Mattia helped to extend a BAO-fitting code DESILIKE in order for Hugo to make measurements with this code. I assisted Hugo with his calculations, debugging, and provided codes to create plots. Alejandro Pérez-Fernández assisted in the project by providing codes and mocks to myself and Hugo to use. Remaining authors not yet mentioned or in the table above are builders in the DESI Collaboration who have provided service allowing for the DESI data used in this project to be collected,

post-processed, and for this project to be completed within the DESI Collaboration. Collaboration members included in the acknowledgements of this manuscript have assisted by providing data and feedback on the writing of this document; in particular, Willem Elbers, Ariel Sanchez and Seshadri Nadathur. Willem also provided MCMC chains for fits to the *Planck* 2018 data for extensions to the Λ CDM model, which was used in this work. Additionally, I had interesting discussion regarding this research with David Camarena which was facilitated by Galileo Galilei Institute of Theoretical Physics in Florence, Italy at the *Neutrino Frontiers* workshop in 2024. We also received useful suggestions and data for some of the plots in this work from Benjamin Wallisch. I also received useful feedback from the anonymous referee when this paper was submitted for publication.

This research has made use of NASA's Astrophysics Data System bibliographic services, the astro-ph pre-print archive at <https://arxiv.org/> and the python libraries MATPLOTLIB, GETDIST, CHAINCONSUMER, NUMPY, SCIPY and PANDAS (Hunter, 2007; Lewis, 2019; Hinton, 2016; Harris et al., 2020; Virtanen et al., 2020b; Wes McKinney, 2010). The fits provided by Willem were produced using the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. Furthermore, Hugo, Mariana, and Sébastien were supported by PAPIIT IN108321, PAPIIT IN116024, and PAPIIT IN115424. Finally, this material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science, Innovation and Universities of Spain (MICIU/AEI/10.13039/501100011033), and by the DESI Member Institutions: <https://www.desi.lbl.gov/collaborating-institutions>.

The data to produce the figures in this work can be found at <https://zenodo.org/records/14311742>. The code used to produce fits with BARRY, including modifications for the phase shift amplitude, is available at <https://github.com/abbew25/Barry>. The DESILIKE code to produce fits with the phase shift amplitude can be accessed at <https://github.com/cosmodesi/desilike>.

Chapter 5

Constraining the phase shift of relativistic species in DESI BAOs

For the rest of this thesis, I focus on the ability of large-scale structure probes to measure cosmological parameters, focusing on the impact of free-streaming neutrinos whose presence in the early Universe impacts structure growth today. In Chapter 2 and Chapter 3 I discuss how neutrinos decouple quickly from the primordial plasma and propagate without further interactions. Recall that an impact of free-streaming neutrinos is to create a temporal shift in the gravitational potential that impacts the acoustic waves known as baryon acoustic oscillations (BAOs), resulting in a non-linear spatial shift in the Fourier-space BAO signal. In this work, I make use of and extend upon an existing methodology to measure the phase shift amplitude β_ϕ and apply it to the DESI Data Release 1 (DR1) BAOs with an anisotropic BAO fitting pipeline. Myself and my collaborators validate the fitting methodology by testing the pipeline with two publicly available fitting codes used by the DESI Collaboration, `desilike` and `Barry`. We applied both codes to highly precise cubic box simulations and realistic simulations representative of the DESI DR1 data. I found in my tests that further study towards the methods used in fitting the BAO signal (mainly regarding the broadband power spectrum) will be necessary to ensure accurate constraints on β_ϕ in future DESI data releases. Using DESI DR1, individual measurements of the anisotropic BAO distortion parameters and the β_ϕ for the different tracers are presented, and additionally a combined fit to β_ϕ resulting in $\beta_\phi = 2.7 \pm 1.7$. After including a prior on the distortion parameters from these constraints using *Planck* it is found that $\beta_\phi = 2.7^{+0.60}_{-0.67}$ suggesting $\beta_\phi > 0$ at 4.3σ significance. This result may hint at a phase shift that is not purely sourced from the standard model expectation for N_{eff} or could be an upwards statistical fluctuation in the measured β_ϕ ; this result relaxes in models with additional freedom beyond Λ CDM.

5.1 Overview

As discussed in Chapter 2 and Chapter 3, the present-day distribution of matter perturbations in the Universe has grown from the over and under-densities that froze in place when sound waves in the early Universe ceased to propagate. These sound waves are called the BAOs, and could propagate in the hot and dense plasma at early times due to the coupling between baryons and photons. At the epoch of recombination at $z \sim 1090$, photons were no longer energetic enough to ionize atoms, and photons began decouple from the baryons and stream freely; these photons are measured today as the cosmic microwave background (CMB, Penzias and Wilson, 1965). The BAOs ceased to propagate completely at the baryon drag epoch at which baryons and photons completely decoupled.

In the present day, galaxies throughout the Universe trace the underlying matter distribution; galaxy surveys such as the Dark Energy Spectroscopic Instrument (DESI, Aghamousa et al., 2016a,b; Abareshi et al., 2022; Adame et al., 2025b) aim to probe the large-scale distribution of galaxies to measure the imprint of the BAOs and test models of cosmology. This has already been probed by various surveys (Cole et al., 2005; Eisenstein et al., 2005; Gil-Marín et al., 2016; Abbott et al., 2019) and most recently in DESI data (Moon et al., 2023; Adame et al., 2025a,b) who have measured the ‘peak’ in the two-point auto-correlation function of galaxies or the ‘wiggles’ that can be seen due to this peak in Fourier space (the Fourier-space transform is the power spectrum). These signals are shown in Figure 2.3. Again recall that the peak corresponds to the BAO scale r_s ; the longest distance sound waves could travel up until the baryon drag epoch and also represents a separation scale at which there is an excess probability of finding pairs of galaxies.

Historically, the BAO peak is used to test cosmological models in a partially model-independent way; by measuring a parameter α that represents a ratio of distance scales in a fiducial cosmological model (used to convert the galaxy catalog angular coordinates and redshifts into Cartesian coordinates) to the same distance scales in the true cosmology. This was discussed in section 3.1.5 of Chapter 3 although I will recap the mathematics here. If the fiducial cosmology matches the truth, the parameter α is unity; if not, deviations from unity occur as the physical BAO scale and the BAO scale of the fiducial cosmology do not match. The scaling of the parameter α physically represents an isotropic scaling or distortion of the distance to the BAO peak (Bernal et al., 2020).

I also discussed in section 3.1.5 how the BAO analysis can be extended to account for anisotropic distortions, which are described by the Alcock-Paczynski (AP, Alcock and Paczyński 1979) effect. In essence, the difference between the assumed fiducial cosmology model and the true cosmology distorts the distribution of galaxies in the sky differently in directions transverse or parallel to the line-of-sight (Bernal et al., 2020). The effect arises due to transforming the redshift-space coordinates of objects to real-space coordinates, and the distortions are more significant at higher redshift. In the anisotropic analysis, an additional scaling parameter α_{AP} can be measured, representing the *anisotropic* elongation or compression of the BAO.¹ Depending on the constraining power of the available data, it

¹It is also common to change variables from α and α_{AP} to α_{\parallel} and α_{\perp} , which represent the scaling of the BAO parallel and transverse to the line-of-sight respectively. A simple transformation exists between the two bases, which is presented in Section 5.2.

may only be possible to measure the α parameter using the isotropic fitting methodology. That is, if only the monopole multipole (the first $\ell = 0$ term in the Legendre expansion of the anisotropic power spectrum, corresponding to the spherically isotropic component - see Equation 3.28) of the galaxy power spectrum is constrained by the data. If the quadrupole multipole ($\ell = 2$) is constrained this contains more information about the shape of the power spectrum anisotropy and enough information to measure α and α_{AP} in an anisotropic fit. The fits to the values of α and α_{AP} can eventually be turned into constraints on cosmological parameters or models, or combined with other datasets such as *Planck* (Aghanim et al., 2020) to improve constraints on the matter density Ω_m and the expansion rate H_0 . As various cosmological parameters that change the expansion rate will have degenerate effects on α and α_{AP} , the distortion parameters may require measurements in several redshift bins to obtain useful constraints on cosmological parameters.

The majority of analyses to this date have not taken into account an additional piece of information that the imprint of the BAOs in galaxy surveys may give us about light relics in the early Universe. Bashinsky and Seljak (2004) first showed that free-streaming neutrinos (or any light relics that free-stream) with a speed faster than the sound speed in the early Universe induce a shift in the phase of the BAO wiggles that can be seen in Fourier space (an improved semi-analytic calculation of the phase shift can also be found in Green and Ridgway, 2020). This phase shift can be understood in that as neutrinos propagate after decoupling from the primordial plasma, they alter the gravitational potential that influences it by carrying away energy. This alteration of the potential induced a temporal shift in the BAO sound waves, that consequently led to a shift in the scale at which the sound waves froze in place and seeded the large-scale structure we see today. The phase shift is also discussed in detail in Baumann et al. (2016). Under the assumption that the primordial fluctuations are adiabatic, this phase shift is not degenerate with any other parameters and is a unique signature of neutrinos. Even if there were non-adiabatic primordial fluctuations, the scale dependence of the phase shift may still possibly allow it to be distinguished from non-adiabatic fluctuations (Baumann et al., 2018).

The phase shift was first parameterized and measured in CMB data from the *Planck* satellite by Follin et al. (2015), whose analysis showed a probability $p = 8 \times 10^{-6}$ for the absence of a phase shift due to free-streaming relics, indicating a strong presence of free-streaming relics. More recently, it was parameterized and measured in BAO data from BOSS DR12 by Baumann et al. (2019) who found with 95% confidence a non-zero phase shift by free-streaming relics, and additionally in Montefalcone et al. (2025c) with CMB data. The parameterization by Baumann et al. (2019), like the standard BAO analysis, involves a semi model-independent parameter β that characterizes the amplitude of the phase shift and can be transformed into a constraint on N_{eff} , the effective number of neutrino species. The effect of altering N_{eff} in the matter power spectrum, which shifts the phase in the BAOs, can be seen in Figure 5.1.

The effect in configuration space can be seen in Figure 5.2, where the correlation function has been calculated from the power spectrum using the MCFIT package.² It is apparent from the figures, that the impact of the phase shift from varying N_{eff} on the BAO wiggles is subtle compared to various other

²<https://github.com/eelregit/mcfit>

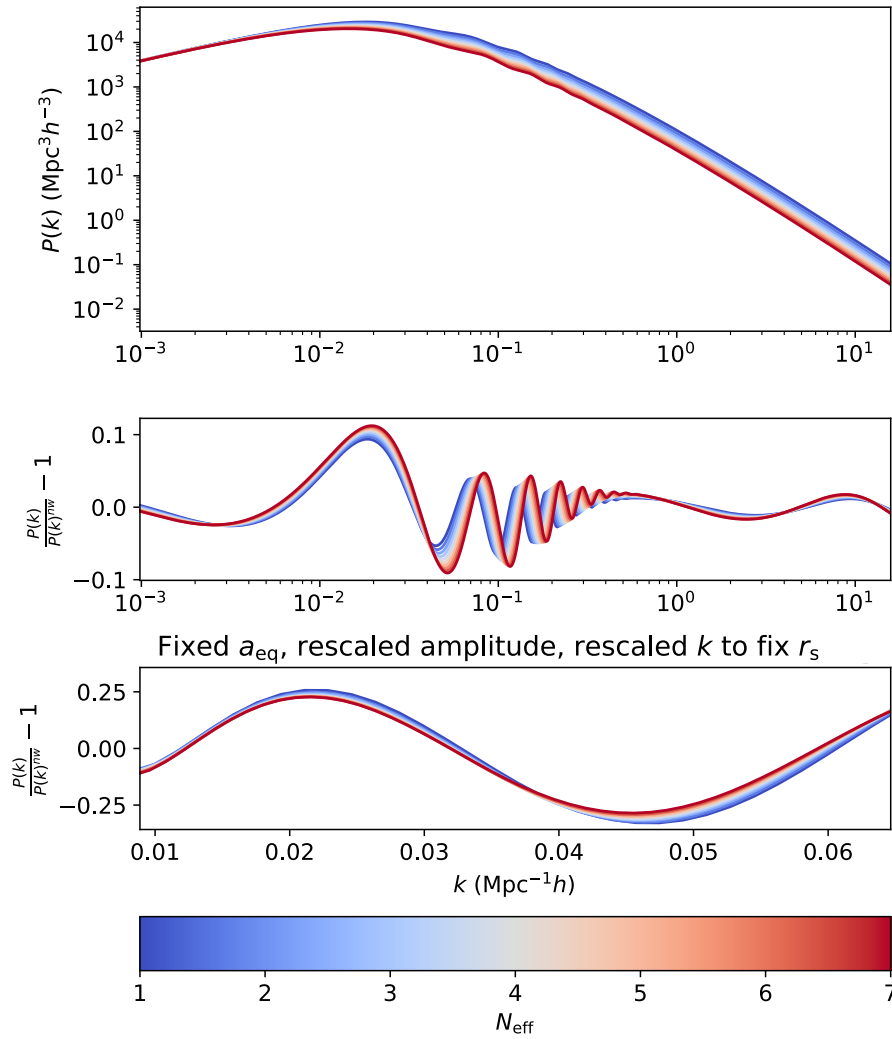


Figure 5.1: This figure has been inspired by Figure 3 in Baumann et al. (2018). The top panel shows the effect of altering N_{eff} on the matter power spectrum, while the middle panel shows the ratio of the power spectrum plotted to a de-wiggled (smooth) power spectrum; all other cosmological parameters are held fixed. The power spectrum has been calculated using the Boltzmann solver CLASS (Blas et al., 2011), and smoothed to isolate the wiggles using the method in Hinton et al. (2016). Altering N_{eff} has several notable effects: 1) it alters the amplitude of the BAO wiggles because the ratio of radiation to baryonic matter is changed and this affects the amplitude of rarefactions and compressions in the BAO sound waves; 2) the amplitude of the spectrum for k -modes that entered the horizon before the epoch of matter-radiation equality are suppressed by increasing N_{eff} because the additional radiation density suppresses structure growth of these modes during the era that radiation dominates the expansion rate; and 3) additional radiation alters the position of the peaks and the large turnover by increasing the expansion rate, which allows the epoch of matter-radiation equality and recombination to occur at an earlier time; the scales are smaller than in a Universe that takes longer to reach these epochs. In the lowest panel these effects have been removed to isolate the effect of N_{eff} on only the phase of the BAO wiggles. This is done by calculating the power spectrum wiggles as in the middle panel, but with various differences. When altering N_{eff} we also alter Ω_{m} such that $a_{\text{eq}} = \frac{\Omega_{\text{r}}}{\Omega_{\text{m}}}$ is held fixed. For each model with different N_{eff} and Ω_{m} , we rescale the k modes plotted on the x -axis by the ratio of r_s^f and the model r_s and (where f is a fiducial model with $N_{\text{eff}} = 3.044$, $\Omega_{\text{m}} = 0.315$) so that the impact of the changing r_s is effectively removed. Finally, the amplitude of the wiggles is rescaled so that there is a fixed maximum height for any of the peaks shown.

impacts N_{eff} has on the power spectrum or correlation function. Nonetheless, it can be detected.

Baumann et al. (2018, 2019) also discuss and provide forecasts for the potential for upcoming galaxy surveys that will probe the large-scale structure to constrain the phase shift and N_{eff} . One such survey is the DESI survey (Aghamousa et al., 2016a,b; Abareshi et al., 2022; Adame et al., 2024), a 5-year galaxy survey that is currently in progress and began in 2020.³ It is using an instrument that consists of 5000 robot fibres, mounted on the focal plane of the Mayall Telescope at Kitt Peak, Arizona, to measure spectra for 35 million galaxies. The DESI instrument has already detected the BAO signal in their DR1 Data (Adame et al., 2025a) with a 9.1σ level detection at 0.86% precision using a sample of luminous red galaxies (LRGs) and emission line galaxies (ELGs). In Baumann et al. (2018, 2019), it was anticipated using Fisher Matrix Forecasts that DESI will be able to measure the phase shift with an uncertainty of $\sigma(\beta) \sim 0.3$ from their 5-year survey; DESI BAOs + *Planck* may achieve $\sigma(N_{\text{eff}}) = 0.15$, and with the full broadband power spectrum from DESI may achieve $\sigma(N_{\text{eff}}) = 0.087$.




Figure 5.2: Top panel: the effect of altering N_{eff} on the correlation function, calculated from a Fourier transform of the power spectrum shown in the top panel of Figure 5.1. Lower panel: the correlation function calculated from the power spectrum, using a Fourier transform, after including a phase shift following the parameterization of the phase shift shown in Baumann et al. (2019). The power spectrum is shifted along the x -axis by shifting the k values by an amount expected by for a given N_{eff} and interpolating the power spectrum. Unlike the top panel, the lower panel shows only the effects due to a phase shift on the correlation function rather than the various impacts that can be seen in the top panel due to N_{eff} , including the phase shift.

³<https://www.desi.lbl.gov/>

In this Chapter, my collaborators and I start down this path by constraining the phase shift with DESI DR1 BAO data. Unlike the constraints shown in Baumann et al. (2019), we have extended the methodology to constrain β with both the isotropic and anisotropic BAO parameters, α and α_{AP} . In Section 5.2, I give an overview of the anisotropic fitting methodology and the parameterization of the phase shift given by Baumann et al. (2019) and the implementation of the method in two BAO fitting codes; Barry (Hinton et al., 2020) and `desilike`.⁴ In Section 5.3, tests by myself and my collaborators show the robustness of the approach and codes applied to mocks produced for the DESI collaboration for datasets including the Bright Galaxy Survey (BGS) (Hahn et al., 2023), LRGs (Zhou et al., 2023), ELGs (Raichoor et al., 2023) and quasars (QSOs) (Chaussidon et al., 2023), and constraints for the combinations of these datasets. The results from applying the methodology to the DESI DR1 BAO data will be shown in Section 5.4 and I conclude in Section 4.6.

5.2 Methodology

In this work we (myself and collaborators) follow the methodology of Baumann et al. (2018, 2019) to constrain the phase shift. As mentioned previously, instead of only constraining the isotropic BAO scaling parameter α and the phase shift, we extend the approach to measure the anisotropic BAO by fitting α_{\parallel} , α_{\perp} and the phase shift amplitude β . Furthermore, the methodology is implemented in two codes to fit these parameters; Barry and `desilike` which can both fit the BAO wiggles using the galaxy power spectrum or the BAO peak in the correlation function. Having these two approaches will allow us to see the robustness of our constraint under different methodologies.

5.2.1 Isotropic BAO analysis

Reminding from Chapter 3, in the standard isotropic BAO analysis, the galaxy power spectrum can be split into a smoothed component (or no-wiggle component/broadband component) and a wiggle component,

$$\begin{aligned} P(k) &= P_{\text{nw}}(k) + P_{\text{w}}(k) \\ &= P_{\text{nw}}(k) \left(\frac{P_{\text{w}}(k)}{P_{\text{nw}}(k)} + 1 \right) \\ &= P_{\text{nw}}(k) (\mathcal{O}(k) + 1). \end{aligned} \tag{5.1}$$

The goal of the analysis is to isolate the BAO wiggles from the observed power spectrum, and fit the distortion parameter α which scales the true positions of the BAO peaks and troughs in Fourier space. The distance scaling of the peaks that occurs is by an amount relative to the expected positions of the BAO peaks in a *fiducial* model that is assumed when the observed redshifts of galaxies are transformed to distances, and also relative to a *template* cosmological model that is assumed when one is fitting the

⁴<https://desilike.readthedocs.io/en/latest/#>

BAOs; often these fiducial and template cosmologies are the same. However in general, the parameter α can be defined as

$$\alpha = \frac{D_V(z) r_s^t}{D_V^f(z) r_s}. \quad (5.2)$$

As before $D_V(z)$ corresponds to the volume averaged distance,

$$D_V(z) = \left((1+z)^2 d_A^2(z) \frac{cz}{H(z)} \right)^{1/3}, \quad (5.3)$$

where $H(z)$ is the Hubble parameter at the redshift z corresponding to the redshift of the BAO measurement, and $d_A(z)$ is the angular diameter distance. r_s corresponds to the sound horizon at the baryon drag epoch (at z_d). The quantities with the ‘f’ superscript refer to the expected value in the fiducial model, while those with the ‘t’ superscript refer those in the template model. When all of these models are the same as the true cosmology, $\alpha = 1$; but deviations from unity allow us to quantify the deviation of the models from the true cosmology. α enters the power spectrum as a rescaling of the measured k -modes, and one writes the measured modes $k' = k/\alpha$, where k denotes the true modes. The power spectrum is also scaled by a ratio of the volume averaged distances cubed due to the fact that the rescaling of distances changes the defined survey volume, which must be accounted for as the power spectrum has units of volume. Typically the redshift-space power spectrum $P(k, \mu)$ is decomposed into moments using a Legendre polynomial expansion, and the isotropic analysis allows us to constrain the monopole moment corresponding to $\ell = 0$. The parameter α enters the wiggle power spectrum template as

$$P'_w(k) = P_w^t(k/\alpha). \quad (5.4)$$

5.2.2 Anisotropic BAO analysis

In this section I will restate the parameters from section 3.1.5 of Chapter 3 for the anisotropic BAO analysis, which can be performed when is possible to constrain both the monopole and the quadrupole corresponding to $\ell = 0, 2$ respectively. The analysis constrains two distortion parameters, α_{\parallel} that scales the BAO peaks along the line-of-sight while α_{\perp} constrains the BAO peaks transverse to the line-of-sight; when the fiducial cosmology and template cosmology agree with the real cosmology, these parameters are unity. Alternatively, the anisotropic analysis can be parameterized in terms of α defined previously and α_{AP} for the distortions, where α_{AP} is also unity when all three cosmologies are the same. The parameters α_{\parallel} and α_{\perp} were defined previously as⁵

$$\alpha_{\parallel} = \frac{H^f(z) r_s^t}{H(z) r_s}, \quad (5.5)$$

$$\alpha_{\perp} = \frac{D_A(z) r_s^t}{D_A^f(z) r_s}. \quad (5.6)$$

⁵Strictly, the geometric AP effect actually is only due to distortions from the ratios $q_{\parallel} = \frac{H^f(z)}{H(z)}$ and $q_{\perp} = \frac{D_A(z)}{D_A^f(z)}$, because the factor of $\frac{r_s^t}{r_s}$ only enters due to the choice of BAO template. However it is common to refer to α_{\parallel} and α_{AP} in the literature when constraining the BAO distortion parameters.

Furthermore, equation 5.5 and equation 5.6 can be related to α and α_{AP} as

$$\begin{aligned}\alpha &= \alpha_{\text{iso}} = \alpha_{\parallel}^{1/3} \alpha_{\perp}^{2/3}, \\ \alpha_{\text{AP}} &= \alpha_{\parallel} / \alpha_{\perp}.\end{aligned}\tag{5.7}$$

If one assumes a flat Λ CDM cosmology, it is possible to directly relate α_{AP} to the matter density Ω_{m} ,

$$\begin{aligned}\alpha_{\text{AP}} &= \alpha_{\text{AP}}(z, \Omega_{\text{m}}), \\ &= (H(z)D_A(z))^{\text{f}} (H(z, \Omega_{\text{m}})D_A(z, \Omega_{\text{m}}))^{-1}, \\ &= (E(z)D_A(z)h)^{\text{f}} (E(z, \Omega_{\text{m}})D_A(z, \Omega_{\text{m}})h)^{-1}.\end{aligned}\tag{5.8}$$

where I have used $H(z) = 100 \text{ km s}^{-1} \text{ Mpc}^{-1} h E(z)$. The term containing the angular diameter distance, $D_A(z, \Omega_{\text{m}})h$, depends on only Ω_{m} and z . α on the other hand relates to both Ω_{m} and the comoving sound horizon $r_s h$,

$$\begin{aligned}\alpha &= \alpha(\Omega_{\text{m}}, r_s h), \\ &= \left[\left(\frac{H(z)}{D_A^2(z)} \right)^{1/3} r_s^{\text{f}} \right] \left[\left(\frac{H(z)}{D_A^2(z)} \right)^{1/3} r_s \right]^{-1}, \\ &= \left[\left(\frac{E(z)}{D_A^2(z)h^2} \right)^{1/3} r_s^{\text{f}} h \right] \left[\left(\frac{E(z)}{D_A^2(z)h^2} \right)^{1/3} r_s h \right]^{-1}.\end{aligned}\tag{5.9}$$

It is therefore convenient to note that, for a flat Λ CDM cosmology, one can uniquely map from a set of BAO parameters to a cosmological model by first computing Ω_{m} given α_{AP} , then using this fixed value alongside α to isolate $r_s h$. The distortion parameters enter the wiggle power spectrum template as

$$\begin{aligned}P_{\text{w}}(k, \mu) &= P^{\text{t}}(k(k', \mu'), \mu(k', \mu')), \\ &= P^{\text{t}}(k(\alpha_{\parallel}, \alpha_{\perp}), \mu(\alpha_{\parallel}, \alpha_{\perp})),\end{aligned}\tag{5.10}$$

where k' and μ' are the observed (distorted) k and μ due to anisotropic distortions. The observed k' , μ' can be related to the real k , μ as

$$\begin{aligned}k' &= \frac{k}{q_{\perp}} \left[1 + \mu^2 \left(\frac{1}{F^2} - 1 \right) \right]^{1/2}, \\ \mu' &= \frac{\mu}{F} \left[1 + \mu^2 \left(\frac{1}{F^2} - 1 \right) \right]^{-1/2},\end{aligned}\tag{5.11}$$

where $F = q_{\parallel} / q_{\perp}$, $q_{\parallel} = H^{\text{f}}(z) / H(z)$ and $q_{\perp} = D_A(z) / D_A^{\text{f}}(z)$.

5.2.3 Anisotropic BAO analysis with β_{ϕ}

Having recapped the earlier theory, now we can look at the extended methodology that accounts for the phase shift β_{ϕ} . Similar to the BAO distortion parameters, the phase shift is constrained by measuring

the difference in the amplitude of a phase shift from a template and the true cosmology. However, one should keep in mind that while the phase shift is a parameter similar to α and α_{AP} in that it is model independent and can be transformed into a constraint on cosmology, it is not a geometric distortion but is a physical change to the true k modes in the power spectrum. The BAO wiggle feature can be expressed as a sinusoidal oscillation with a k -dependent amplitude,

$$O_{\text{lin}}(k) \propto A(k) \sin(kr_s + \phi(k)), \quad (5.12)$$

where ϕ is the phase shift; it is parameterized in Baumann et al. (2018, 2019) as

$$\phi(N_{\text{eff}}, k) = \beta_\phi(N_{\text{eff}})f(k). \quad (5.13)$$

Here β_ϕ is the relative amplitude of the phase shift compared to a template model, and $f(k)$ is a fitting function for the scale dependence of the ϕ . When the template cosmology model chosen and the true cosmology match, β_ϕ is unity. β_ϕ is related to N_{eff} as

$$\beta_\phi = \frac{\varepsilon_\nu}{\varepsilon_\nu^t} = \frac{N_{\text{eff}}/(N_{\text{eff}} + a_\nu)}{N_{\text{eff}}^t/(N_{\text{eff}}^t + a_\nu)}; \quad a_\nu = \frac{8}{7} \left(\frac{11}{4} \right)^{4/3}. \quad (5.14)$$

⁶Given that the phase shift is a physical effect caused by free-streaming relativistic species, rather than an effect that is induced by the choice of the fiducial model used to construct the power spectrum or correlation function, we expect β_ϕ to be independent of the choice of fiducial cosmology and only dependent on the template cosmology. If $N_{\text{eff}} = 0$, we also have $\beta_\phi = 0$. In the limit $N_{\text{eff}} \rightarrow \infty$, $\beta_\phi \approx 2.4$ when $N_{\text{eff}}^t = 3.044$. $f(k)$ is redshift independent and has been approximated by Baumann et al. (2018); it is to first order independent of N_{eff} for a range of values of N_{eff} . This can be seen in Figure 1 of Baumann et al. (2019). The fitting function template $f(k)$ is parameterized as

$$f(k) = \frac{\phi_\infty}{1 + (k_*/k)^\xi}, \quad (5.15)$$

with $\phi_\infty = 0.227$, $k_* = 0.0324h\text{Mpc}^{-1}$ and $\xi = 0.872$. As $k \rightarrow \infty$, $f(k) \rightarrow \phi_\infty$. Finally, the observed wiggle power spectrum can be written in terms of the anisotropic distortion parameters and the phase shift amplitude as

$$\begin{aligned} P_w(k, \mu) &= P^t(k(k'), \mu') + (\beta_\phi - 1)f(k')/r_s, \mu(k', \mu'), \\ &= P^t(k(\alpha_\parallel, \alpha_\perp) + (\beta_\phi - 1)f(k')/r_s, \mu(\alpha_\parallel, \alpha_\perp)). \end{aligned} \quad (5.16)$$

It is also possible to directly map the anisotropic BAO distortion parameters to Ω_m and $r_s h$ in a flat Λ CDM cosmology. I note here that unlike equation 3.4 in Baumann et al. (2019), we include the impact of the distortions in the scale-dependent function $f(k)$ for the phase shift in equation 5.16. However, we found there was no impact on the fits to any parameters by choosing to include the distortion k' rather than simply k in $f(k)$ (which is how the model is implemented in BARRY); however the impact may be more significant in cases one explores exotic models which alter the function $f(k)$.

⁶ a_ν here relates the energy density of neutrinos Ω_ν to N_{eff} , such that β_ϕ in Equation 5.14 gives the ratio of the fractional energy density of neutrinos relative to the total radiation energy density. The factor of $8/7$ accounts for the fact the energy density of the neutrinos follow a Fermi-Dirac rather than Bose-Einstein distribution, and the factor of $(11/4)^{4/3}$ arises from a calculation to relate the temperature of photons to neutrinos after electron-positron annihilation, in which the photon temperature is raised relative to the decoupled neutrinos.

5.2.4 Implementation

The fitting methodology for the BAO scaling parameters and phase shift amplitude has been implemented in both `DESILIKE`⁷ and `BARRY`⁸ (Hinton et al., 2020). These implementations are both *python* packages written for fitting the BAO feature and obtaining constraints on the BAO parameters. While both codes are capable of standard isotropic and anisotropic fits⁹, in this work I have extended `BARRY` in order to additionally constrain the phase shift.¹⁰ `DESILIKE` also contains the implementation via recent updates. While these codes have various available clustering and BAO models to choose from, here I focus on using and implementing the phase shift parameterisation with a model similar to the style of Beutler et al. (2017). This model has been widely used in previous BAO analyses and recently by the DESI collaboration for the analysis of DR1 (the model used is presented and discussed in Chen et al., 2024a). I use the same modelling to be consistent with the DESI Year 1 results (Adame et al., 2025b).

Power spectrum

In `DESILIKE` and `BARRY` the redshift-space power spectrum is generically expressed in the following way for the BAO analysis, using the decomposition of the power spectrum into a broadband (no-wiggles) component and a sinusoidal wiggles component as in equation 5.1 (Chen et al., 2024a),

$$P(k, \mu) = \mathcal{B}(k, \mu)P_{\text{nw}}(k) + \mathcal{C}(k, \mu)P_{\text{w}}(k) + \mathcal{D}(k, \mu). \quad (5.17)$$

\mathcal{B} accounts for peculiar velocity effects such as the Kaiser effect (Kaiser, 1987) and the Finger-of-God effect (FOG, Jackson, 1972), \mathcal{C} is the propagator which accounts for non-linear evolution of the BAO feature, and \mathcal{D} accounts for additional features in the measured broadband component of the power spectrum not captured by the no-wiggle component, for instance due to systematics or non-linear structure growth. This term is discussed in more detail in section 5.2.4. In our model, we express \mathcal{B} as

$$\mathcal{B}(k, \mu) = (b + f\mu^2)^2 \left[1 + \frac{1}{2}k^2\Sigma_s^2 \right]^{-2} \quad (5.18)$$

where b represents the linear galaxy bias parameter, f is the growth rate of structure (distinctly different from the fitting function for the neutrino phase shift, $f(k)$), and where both parameters combine to form the Kaiser factor $(b + f\mu^2)^2$. Σ_s accounts for FoG damping, and all three of these are free parameters in the BAO analysis. The Kaiser factor and the FoG factor may be familiar from Chapter 3. The propagator is expressed as

$$\mathcal{C} = (b + f\mu^2)^2 \exp \left[-\frac{1}{2}k^2(\mu^2\Sigma_{\text{nl},\parallel}^2 + (1 - \mu^2)\Sigma_{\text{nl},\perp}^2) \right] \quad (5.19)$$

⁷<https://github.com/cosmodesi/desilike>

⁸<https://barry.readthedocs.io/en/latest/>

⁹The `BARRY` code has been updated since the work presented in Hinton et al. (2020), see more about the improvements in Chen et al. (2024a)

¹⁰<https://github.com/abbew25/Barry>

where $\Sigma_{\text{nl},\parallel}$, $\Sigma_{\text{nl},\perp}$ are additional free parameters that account for non-linear damping of the BAO due to bulk motions of galaxies (typically on scales $\sim 10h^{-1}$ Mpc) separately for damping along and perpendicular to the line-of-sight respectively. For an isotropic analysis this simply reduces to

$$\mathcal{C} = (b + f\mu^2)^2 \exp\left[-\frac{1}{2}k^2\Sigma_{\text{nl}}^2\right], \quad (5.20)$$

involving only one isotropic damping parameter Σ_{nl} .

Reconstruction

In a BAO analysis it is typical to apply a type of algorithm called *reconstruction* to the data, which enhances the BAO signal allowing for improved signal-to-noise. Data products, such as the power spectrum, where a reconstruction algorithm has been applied are referred to as *post-recon* while those without the application of this algorithm are called *pre-recon*. This algorithm is described in greater detail theoretically in Chen et al. (2024a) and algorithms were tested practically with optimal settings for DESI data found in Paillas et al. (2025); Chen et al. (2024b), however the basic steps are described here. First, it involves smoothing the galaxy density field with a Gaussian filter to filter out more non-linear modes. Then, the Zeldovich approximation is used in order to calculate the displacement of the positions of galaxies in space due to structure growth, and this is used to shift the positions of a random catalogue of galaxies and the galaxies in the data by subtracting the calculated linear theory displacement. The reconstructed density field is then given by the difference between the shifted galaxy density field of the data and the shifted random catalogue. The effect of this is to sharpen the BAO feature in the data, because distortions due to non-linear structure growth and the bulk motions of galaxies are largely removed. Consequently, the parameters $\Sigma_{\text{nl},\parallel}$, $\Sigma_{\text{nl},\perp}$ and Σ_s vary between mocks and data that are pre-recon and post-recon, with reconstruction reducing the values of the damping parameters.

In this work I use post-recon mocks that have been reconstructed using the RecSym convention in DESI, in which case the displacement applied to the random catalogue and data are equivalent and the anisotropy in redshift space is preserved (this in contrast to the RecIso formalism – a more detailed discussion of both can also be seen in Chen et al., 2024a; Paillas et al., 2025).

Polynomial broadband methodology

Historically, the term $\mathcal{D}(k, \mu)$ has been modelled using a polynomial expansion. In the methodology of Beutler et al. (2017), this can be expressed for the monopole and quadrupole for the pre-recon data as

$$\mathcal{D}_\ell(k) = \frac{a_{\ell,1}}{k^3} + \frac{a_{\ell,2}}{k^2} + \frac{a_{\ell,3}}{k} + a_{\ell,4} + a_{\ell,5}k. \quad (5.21)$$

For the post-recon analysis, it is written as

$$\mathcal{D}_\ell(k) = \frac{a_{\ell,1}}{k^3} + \frac{a_{\ell,2}}{k^2} + \frac{a_{\ell,3}}{k} + a_{\ell,4} + a_{\ell,5}k^2. \quad (5.22)$$

The coefficients $a_{\ell,i}$ and number of polynomial terms are chosen based on the best fit to the data. To be consistent with the DESI methodology and formalism shown in Chen et al. (2024a) and Paillas et al. (2025), we fit for all the coefficients in our analysis.

Spline broadband methodology

The polynomial broadband methodology suffers from some limitations (this is described in detail in Chen et al., 2024a) due to the fact it is not always clear what the optimal number or form of polynomial coefficients to use should be, and when transformed to configuration space they may result in a peak (mimicking the BAO peak). Therefore Chen et al. (2024a) present a new approach to modelling additional contributions to the broadband power spectrum, using a summation of cubic splines with bases separated by some choice Δ (set to $0.06 h^{-1}$ Mpc in our work),

$$\mathcal{D}_\ell(k) = \sum_{n=-1}^{\infty} a_{\ell,n} W_3\left(\frac{k}{\Delta} - n\right). \quad (5.23)$$

$W_3(k)$ is a fourth order piece-wise cubic spline that is given in Chen et al. (2024a). This approach importantly removes degeneracy between the broadband component of the power spectrum and the BAO peak and reduces the choice of both the form and number of polynomials to a single choice of Δ (which has a well-defined lower bound — it should be larger than half the BAO wavelength). This choice of Δ alongside the range of scales being fit automatically defines the number of free broadband parameters. Distinct terms can also be set for each multipole moment of the power spectrum. In the analysis our choice of fitting scales ($k = 0.02 - 0.30 \text{ Mpc}^{-1}h$) and Δ give 7 $a_{\ell,n}$ coefficients to be fit for each multipole. In our analysis, we test the performance of both the spline methodology and the polynomial methodology in our fits to the phase shift.

Correlation function

To analyse the BAO in configuration space, the redshift-space correlation function is simply calculated from a Hankel transform of the redshift-space power spectrum to configuration space. The only difference in our analysis of the correlation function compared to the power spectrum is that we set the number of coefficients to only three for analyses using the polynomial broadband methodology. For the spline broadband methodology, the transformation of the cubic spline basis presented above results in two large scale $n = [0, 1]$ terms for the quadrupole which we include, while the remainder are confined to small scales outside our fitting range ($50 - 150 \text{ Mpc}h^{-1}$). In addition, the scale-cuts used when fitting the power spectrum to avoid systematics transform into simple polynomials. As such, the final form of $\mathcal{D}_\ell(k)$ for spline-based correlation function fits is completely analogous to that for the power spectrum and takes the form

$$\begin{aligned} \mathcal{D}_0(k) &= a_{0,0} + a_{0,1} \left(\frac{rk_{\min}}{2\pi}\right)^2, \\ \mathcal{D}_2(k) &= a_{0,0} + a_{0,1} \left(\frac{rk_{\min}}{2\pi}\right)^2 + a_{2,0} B_{2,0}(r\Delta) + a_{2,1} B_{2,1}(r\Delta), \end{aligned} \quad (5.24)$$

where

$$B_{\ell,n}(r\Delta) = i^2 \int \frac{dkk^2}{2\pi^2} W_3 \left(\frac{k}{\Delta} - n \right) j_\ell(kr), \quad (5.25)$$

is the result of transforming the cubic spline functions and k_{\min} refers to the minimum choice of k used in our Fourier-space fits. $j_\ell(kr)$ are the spherical Bessel functions of the first kind. The reconstruction procedure using RecSym convention is also applied in the case of the correlation function analysis.

5.3 Mock validation

In order to validate the methodology we first apply the fitting procedure in BARRY and DESILIKE to various mocks for the DESI data. The DESI survey strategically splits the survey into dark-time and bright-time components, to optimize detection of galaxies when the moon illuminates the sky (bright-time, so that only bright galaxies at low-redshift can be detected) and dark-time for fainter objects. As such, the DR1 data consists of BGS (galaxies detected under the bright-conditions), LRGs, ELGs and QSOs (see more detail in Adame et al., 2025b). Mocks are provided for the individual datasets of galaxies detected in the DR1 data set, split into various redshift bins as required, and we obtain fits for each of these individual bins. We also provide combined fits to the mocks by combining the likelihoods for each of the fits to the individual bins. While we produce fits with the β_ϕ parameter using the anisotropic BAO methodology for the majority of the mocks and corresponding data, to be consistent and comparable to the results provided by the DESI Collaboration for the DR1 data (Adame et al., 2025b) we only fit β_ϕ using an isotropic BAO procedure for QSOs, BGS and for ELGs at $z = 0.8 - 1.1$. Furthermore, we fit a combined measurement of the correlation function and power spectrum for ELGs and LRGs at $z = 0.8 - 1.1$ to account for their cross-correlation. Otherwise, following the DESI Collaboration, we neglect cross-correlations between QSOs with ELGs and LRGs from $z = 0.8 - 1.1$ due to the fact that the number of overlapping tracers in the survey volume is likely to be negligible, and neglect cross-correlations between other tracers which are detected in non-overlapping redshift bins. The details of the survey tracers are summarised in the Table 5.1.

Table 5.1: Summary of DESI DR1 data properties. The LRGs are split across three redshift bins (LRG1, LRG2, LRG3) and ELGs across two (ELG1, ELG2).

Tracer	z	N objects
BGS	0.1 - 0.4	300,017
LRG1	0.4 - 0.6	
LRG2	0.6 - 0.8	2,138,600
LRG3	0.8 - 1.1	
ELG1	0.8 - 1.1	2,432,022
ELG2	1.1 - 1.6	
QSO	0.8 - 2.1	856,652

5.3.1 Fitting DESI first-generation mocks

Description of first-generation mocks

In this paper, DESI first-generation mocks (Adame et al., 2025a,b) refer to a set of single-redshift snapshots of full N-body mocks based on the Abacus Summit simulations (Maksimova et al., 2021).¹¹ The first-generation mocks are cubic boxes with a *Planck* 2018 (Aghanim et al., 2020) cosmology (called c000 cosmology in this paper, with $N_{\text{eff}} = 3.044$), used here to test the BAO fitting methods on mocks with precision far beyond that of DESI DR1. The dark matter haloes are populated in the mocks using a Halo Occupation Distribution algorithm (HOD, Rocher et al., 2023; Yuan et al., 2024). We also make use, where available, of first-generation ‘control variate’ (CV) mocks, which have been processed to reduce the noise in the galaxy clustering signal by subtracting off the intrinsic sample variance in the simulated observables (Hadzhiyska et al., 2023). These allow us to test the methodologies with greater precision, and quantify the effects of systematics when the statistical noise in the mocks is smaller than we expect to see even for the final DESI data. More details about the mocks can also be found in Garcia-Quintero et al. (2025); Mena-Fernandez et al. (2025). The second-generation mocks discussed later capture the selection function of the DR1 data and involve an improved HOD algorithm compared to the first-generation mocks. The first-generation mocks do not include systematics in the DESI pipeline like in second-generation mocks. The covariance matrices for the mocks have been produced using 1000 realisations of EZmocks (Chuang et al., 2015), which are produced with the same selection function as is used for the Abacus Summit mocks. We apply the Hartlap correction to the covariance matrices for these mocks (Hartlap et al., 2007; Percival et al., 2014, 2022).

For all the fits to the first-generation mocks, to be consistent with the modelling in Chen et al. (2024a), a Gaussian prior $\mathcal{N}(\mu, \sigma)$ with mean μ and width σ has been used for Σ_s given by $\mathcal{N}(2, 2)$ (Mpch^{-1}), for $\Sigma_{\text{nl},\parallel}$ given by $\mathcal{N}(5, 2)$ (Mpch^{-1}) and $\Sigma_{\text{nl},\perp}$ given by $\mathcal{N}(2, 1)$ (Mpch^{-1}). We have allowed β_ϕ to vary over $\mathcal{U}(-1, 3)$ (uniformly over the range $[-1, 3]$) as a minimum range. For polynomial fits, three and five terms with power-law indices of $[-2, -1, 0]$ and $[-1, 0, 1, 2, 3]$ are used to fit the broadband signal for the correlation function and power spectrum respectively. For the spline methodology, seven coefficients are used for the cubic splines and two coefficients are used for the spline basis which is Hankel transformed for the correlation function fits (in addition to two more coefficients for polynomial terms – the details can be seen in Section 5.2.4). We let α vary over $\mathcal{U}(0.8, 1.2)$ (uniformly over the range $[0.8, 1.2]$) and the parameter $\varepsilon = \alpha_{\text{AP}}^{1/3} - 1$ varies over $\mathcal{U}(-0.2, 0.2)$ in Barry (α_{AP} varies over $\mathcal{U}(-0.2, 0.2)$ in *desilike*).

Comparison of DESILIKE with Baumann et al pipeline

In order to validate the methodology implemented in DESILIKE (using the EMCEE sampler,¹² Foreman-Mackey et al. 2013) we begin by fitting the post-recon correlation function of 25 CV ELG mocks at $z = 1.1$ using both this code and the isotropic pipeline of Baumann et al. (2019). The results, which

¹¹<https://abacussummit.readthedocs.io/en/latest/abacussummit.html>

¹²See Appendix A.4 for explanation on Markov Chain Monte Carlo (MCMC), nested sampling and importance methods, which are used for the fits in this paper.

can be seen in Appendix A.3, indicate the implemented scheme reasonably fits the BAO and phase shift parameters, and we see the expected correlation between α and β_ϕ . We expect that by including α_{AP} as an additional parameter, this should have little impact on the β_ϕ constraints, as there is no degeneracy between these two parameters. The results mainly show there is consistency in the fits to β_ϕ in the isotropic or anisotropic case using the pipeline of Baumann et al. (2016).

Comparison of DESILIKE with BARRY

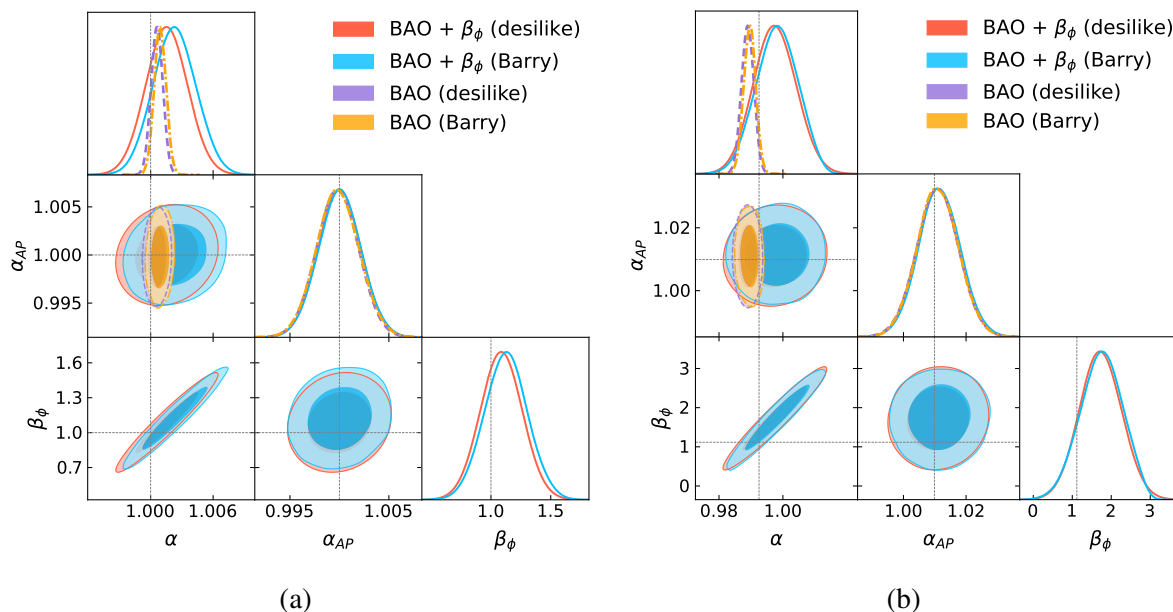


Figure 5.3: Left: fits to the mean of 25 CV mocks using DESILIKE and BARRY for the correlation function. Here we show a comparison of the results both when β_ϕ is fixed or allowed to vary freely. The covariance matrix has been reduced by a factor of 25 compared to a single realization. Right: fits to the mean of 6 power spectrum mocks using BARRY and DESILIKE with the polynomial broadband methodology when the true cosmology of the mocks is c003 and the fiducial and template cosmology are set to c000 (hence the crosshairs deviate from 1-1 positions). This plot also allows one to compare the robustness of the fits to α when β_ϕ is or is not allowed to vary. The covariance matrix has been reduced by a factor of 6 compared to a single realization.

We additionally test the BARRY fitting code with the CV mocks to validate that the fits are consistent using different fitting tools for the anisotropic fitting pipeline. For all our fits with BARRY I use the Nautilus nested sampler (Lange, 2023). In Figure 5.3a we show post-recon fits in comparison to fits with DESILIKE, where β_ϕ is both fixed or allowed to vary. The difference between the peak of the likelihood for α between the two codes when β_ϕ is free vs when it is fixed is $\sim 0.08\%$, and $\sim 0.03\%$ respectively. The difference in the fit to β_ϕ for the two codes is $\sim 4\%$, with variation occurring along the degeneracy direction for these parameters, although they are consistent within 1σ . In both cases, the differences in α and α_{AP} are well within the expected variance from Chen et al. (2024a), and we find no evidence of systematic variation between the two codes.

We also consider the robustness of the fitting methodology in the case the true cosmology of the mocks has N_{eff} set to a value that is not $N_{\text{eff}} = 3.044$, and also in the case the fiducial or template cosmology differs from the cosmology of the mocks. Below we show the result of testing power

spectrum mocks with a true underlying cosmology with $N_{\text{eff}} = 3.7$, which we refer to as c003 (‘CV’ mocks were not available for this cosmology however). The fit to the mean of 6 mocks with a c003 cosmology is shown in Figure 5.3b. We use similar priors for Σ_s , $\Sigma_{\text{nl},\parallel}$, $\Sigma_{\text{nl},\perp}$ in these fits. For these mocks an analytic covariance matrix was available (O. Alves et al., 2024; Forero-Sánchez et al., 2024).

We can expect when using a template with $N_{\text{eff}} = 3.7$ to fit the c003 simulations we should recover $\beta_\phi = 1$, but for a template with $N_{\text{eff}} = 3.044$ we expect that $\beta_\phi \approx 1.117$. The fit shown in Figure 5.3b is statistically consistent with this expectation, although not precise enough to differentiate clearly between the expected value and $\beta_\phi = 1$. However, we still see here that for these mocks there is excellent agreement between the results from our two codes, with no obvious systematic error.

Robustness of modelling choices in BARRY

I additionally compare the results of fitting in BARRY using the polynomial broadband fitting scheme and the spline broadband fitting scheme; these broadband fitting schemes can also be accessed in DESILIKE. A comparison of the broadband methodologies for the fit to the mean of 25 post-recon CV power spectrum mocks and CV correlation function mocks (c000 cosmology snapshot simulations described earlier) can be seen in Figure 5.4. A comparison is shown only for fits when the phase shift is allowed to vary; for fits where the phase shift is fixed (standard BAO analysis) we do not see any significant difference in the fits using either broadband methodology or choice of clustering statistic (this can be seen in Appendix A.5). However, as can be appreciated from Figure 5.4 there is a noticeable difference in the fits to α and β_ϕ between the two different broadband methodologies for the power spectrum, although the variation to the parameters lies along their degeneracy direction and the fits are consistent with the truth within 1σ .




Figure 5.4: Fits to the mean of 25 CV mocks using BARRY, comparing spline and polynomial broadband methodologies. The covariance matrix has been reduced by a factor of 25 compared to a single realization.

In Table 5.2 I show the average difference between the fits to 25 mocks and the difference between the fits to the mocks means for β_ϕ when using the two choices of broadband fitting methodologies,

Table 5.2: The difference in the best fit to β_ϕ , when different broadband methodologies are used. Here we compare the difference in the best fit to the mean of 25 mocks, \bar{P}_w , vs the mean of separate best fits to 25 mocks for the ELGs at $z = 1.1$ after reconstruction. The brackets show the significance of the difference in units of σ (where we use the error on the mean as the uncertainty).

Mock data	$ \langle \beta_{\phi,\text{poly}} - \beta_{\phi,\text{spline}} \rangle _{\bar{P}_w}$ Mock mean	$\frac{1}{25} \sum_i^{25} \beta_{\phi,\text{poly}} - \beta_{\phi,\text{spline}} _i$ 25 mocks
Pk	0.40 (2.00)	0.35 (1.75)
Pk CV	0.30 (2.40)	0.29 (2.32)
Xi	0.09 (0.45)	0.08 (0.40)
Xi CV	0.04 (0.32)	0.05 (0.40)

for all the different cases discussed here. Considering the full set of measurements, the polynomial method applied to the power spectrum appears to be the outlier; I find a systematic difference in the fit to β_ϕ with a significance of 3σ compared to results from the pipeline using the polynomial method with the correlation function or instead using the spline broadband fitting method with either the power spectrum or correlation function.¹³ I find no other significant shifts in the fit to β_ϕ under other comparisons between variations to the pipeline. As such, we treat the correlation function as our default going forward, interchanging between spline and polynomial broadband for various tests as these demonstrate excellent consistency. As such, we do not adopt any systematic error associated with this choice, although caution that this difference should be investigated further for future DESI data releases where we expect tighter constraints. Nonetheless, I note that (given we are fitting the average of 25 mocks) the size of this shift in β_ϕ (shown in Table 5.2) is much smaller than we expect for the DESI DR1 or second-generation mock precision, and there is no clear evidence that this shift is systematic rather than statistical (given the fit still overlaps with the truth at 1σ). As shown later, we also find our results for the second-generation mocks and DR1 data are robust to changes in the broadband methodology.

The effect of altering the number of free terms used in fitting the polynomial broadband terms, the prior ranges used for the nuisance parameters, the choice of the effective redshift for the BAO template or the Boltzmann-solver used to generate the template at a fixed cosmology (which used a *Planck* 2018 cosmology with $z = 0.0$, Aghanim et al., 2020) and the choice of k_{\min} used had little effect on the fits. It should be noted that there can be an impact on the best choice of priors used for nuisance parameters *if* the effective redshift of the BAO template is altered significantly, because some of the nuisance parameters are redshift dependent. However, the optimal choice of priors was thoroughly tested in Chen et al. (2024a) and we used these results to inform our choices of priors used for the fits to the mocks. To test the methodology thoroughly, I investigated the effect of altering the method used to smooth the power spectrum to isolate the wiggle component of the power spectrum $P_w(k)$ from the no-wiggle component $P_{\text{nw}}(k)$. In BARRY there are currently three methods that can be specified to smooth the power spectrum to isolate these, however it was also found to make negligible difference to

¹³I note here this systematic difference is not found to be present in the analysis by Chen et al. (2024a) for the standard BAO analysis, and is thus unlikely to have impacted previous cosmology results.

Table 5.3: A comparison of the absolute difference to the best fits to the BAO and phase shift parameters for the mean of the CV mocks, \bar{P}_w , for ELGs at $z = 1.1$, when the fitting pipeline is varied. The estimate of the significance of the difference relative to the error on the mean is given in the brackets for β_ϕ .

Comparison	$100\Delta\alpha _{\bar{P}_w}$	$\Delta 100\alpha_{\text{AP}} _{\bar{P}_w}$	$\Delta\beta_\phi _{\bar{P}_w}$
BARRY vs DESILIKE (Xi, CV, polynomial)	0.08 (0.53)	0.000 (0.00)	0.04 (0.32)
Pk vs Xi (BARRY, CV, polynomial)	0.47 (3.13)	0.010 (0.01)	0.38 (3.04)
Pk vs Xi (BARRY, CV, spline)	0.05 (0.33)	0.070 (0.05)	0.04 (0.32)
Spline vs polynomial (BARRY, CV, Pk)	0.38 (2.53)	0.040 (0.02)	0.30 (1.71)
Spline vs polynomial (BARRY, CV, Xi)	0.04 (0.05)	0.001 (1.43)	0.04 (0.32)
c000 vs c003 fiducial (BARRY, Pk, polynomial)	—	—	0.18 (1.9)

the fits of the mocks. In Table 5.3 I summarise the most significant differences in the fits to α , α_{AP} and β_ϕ for different modelling choices that we identified. The most significant systematic is the difference we see in the fits for the power spectrum compared to the correlation function in the case we apply fits using the polynomial broadband scheme, which further supports our choice to treat the spline-based method as our default.

It is interesting to test whether we still recover $\beta_\phi = 1$ in the case we alter the fiducial cosmology. This is because in the case the template is set to have $N_{\text{eff}} = 3.044$ such that it is the same as the true cosmology, we expect $\beta_\phi = 1$, regardless of the choice of fiducial cosmology for the clustering (this does not apply for α and α_{AP} , whose truth depends on both the true, fiducial and template cosmology). As mentioned earlier, this is because changing N_{eff} (and thus β_ϕ) in the fiducial cosmology should not alter the fit to β_ϕ since β_ϕ arises as a physical shifting of the k modes in the BAO wiggles rather than a geometric distortion (AP effect) that arises from transforming between redshift-space and real-space coordinates. To test this, I fit the mean of the 25 mocks (non CV mocks) again in Figure 5.5. I use c000 mocks (for the truth and template) where the fiducial cosmology is set to the c003. For brevity I show only the results using the polynomial broadband methodology and I show the absolute difference in the fit to β_ϕ in this scenario in Table 5.3. Note this *differs* from Figure 5.3b, because here I have altered the fiducial cosmology rather than altering the real cosmology. The effect of the fiducial (grid) cosmology in DESI mocks for the standard BAO analysis is studied in detail in Pérez-Fernández et al. (2025) for the more sophisticated ‘second-generation’ mocks.




Figure 5.5: Similar to Figure 5.3b, however only showing the fit to the mean of 25 CV mocks with the polynomial broadband methodology in the case where the template and true cosmology are both set to c000, while the fiducial cosmology is set to c003. In this case, we still expect to recover $\beta_\phi = 1.0$. These fits demonstrate that our method is robust to the choice of fiducial cosmology.

Summary

Overall, our results show a clear degeneracy between α and β_ϕ , which indicates it may be important to include β_ϕ as a free parameter when fitting the BAOs generally (although as I will show later, applying an informative prior to this recovers the published DESI Year 1 results Adame et al. 2025a). Of course, when comparing the systematic differences to the fits in α for the standard case vs the case where we include β_ϕ , the impact is insignificant, due to the fact β_ϕ is weakly constrained. However, for future data we may expect it to be possible the degeneracy between α and β_ϕ could lead to small systematic biases in the case β_ϕ is fixed to one but the data prefers a different value.

From Table 5.3 and Figure 5.4, it is apparent that the fits using the polynomial broadband method for the power spectrum data results in a difference between polynomial and spline fitting methodologies at a significance of 3σ ; for fits where β_ϕ is a free parameter, $\sigma_{\text{sys}}(\beta_\phi) = \pm 0.38$. Given we find no other methodologies result in such systematics, we avoid using the polynomial fitting method for power spectrum fits for the second-generation mocks and data. As mentioned previously, we find all other choices explored for the fitting methodology do not show significant evidence of systematic differences between the best fits to the BAO parameters or β_ϕ relative to the statistical uncertainty in the fits, although we caution this will need to be considered carefully for future DESI data releases.

5.3.2 Fitting DESI second-generation mocks with DESILIKE

In this section, we show the results of fitting realistic second-generation mocks. Unlike the first-generation mocks, these include potential data systematics and the selection functions matching the DESI data (Adame et al., 2025a,b). Data systematics include the impact of the algorithm used to assign

fibers to targets, which can impact the clustering signal due to fiber collisions.¹⁴ We explore fits in which we do not reduce the covariance on the mocks in order to see the expected constraining power for β_ϕ with the realistic uncertainty for a single realisation. For these fits, the covariance matrices for the correlation function are semi-analytic and are created using `RascalC` (Philcox et al., 2020; Rashkovetskyi et al., 2025),¹⁵ while for the power spectrum the covariances are made using `covaPT` (Wadekar and Scoccimarro, 2020; O. Alves et al., 2024).¹⁶

Fits to β_ϕ from individual tracers

Figure 5.6 shows fits to the individual mocks for the correlation function of each tracer produced using `DESILIKE` and `BARRY`; we show fits with the post-recon mocks. We only show the fits using the spline broadband methodology. As with the real data results in section 5.4, our ‘baseline’ results will use the ‘spline’ methodology to be consistent with Adame et al. (2025b,a). However, we expect based on the analysis of systematics from first-generation mocks presented earlier, the second-generation fits will be independent of the choice between a spline or polynomial broadband approach as they are consistent when using the correlation function. For `ELG1` and `LRG3` which overlap in redshift, we use a combined mock for the correlation function (the same is used for the fits to the real data). The priors we use on nuisance parameters for second-generation mocks are specified in Table 5.5. For polynomial fits we set the coefficients for the powers in the fit the broadband shape as $[-2, -1, 0]$ and $[-1, 0, 1, 2, 3]$ for the correlation function and power spectrum respectively. Seven coefficients are used to fit the broadband shape using the spline methodology for power spectrum fits (and as before, the Hankel transformation of the cubic splines to configuration space results in fitting just two coefficients, in addition to two more coefficients for polynomial terms).

It is clear that there is consistency in the fits between `BARRY` and `DESILIKE` tracers. The only significant differences are for fits in cases where neither `BARRY` or `DESILIKE` obtains a strong constraint on the BAO parameters and β_ϕ . In such cases, β_ϕ can hit the end range of the priors used, and we may expect some differences due to differences in how the sampling is performed or due to, subtle differences in the modelling of the Kaiser factor between `desilike` and `Barry` (see discussion in Adame et al. 2025a, although it is expected that these model differences will have negligible effect on the fits to the BAO parameters and β_ϕ and only affect the fits to the nuisance parameters involving the galaxy bias). For the more constraining tracers (i.e., `LRG3+ELG1`), we see excellent agreement between the codes and both `BARRY` and `DESILIKE` obtain unbiased constraints on the phase shift, which validates our methodology for application to the DR1 data. Table 5.4 shows the difference (on average) between the fit to each BAO parameter + β_ϕ between `Barry` and `desilike`, and for comparison the difference between the fit to the mock means. This, like previously for the first-generation mocks, allows us to verify any potential systematic uncertainties on the fits to the parameters due to differences

¹⁴Fiber collisions refer to the effect of two fibers attempting to simultaneously observe two targets close to each other in the sky and colliding; an up-weighting procedure is used to correct for this impact which otherwise leads to systematics in the measured galaxy clustering on small scales (see more detail in Lasker et al., 2024).

¹⁵<https://github.com/misharash/RascalC-scripts/tree/DESI2024/DESI/Y1>

¹⁶<https://github.com/cosmodesi/thecov>

Table 5.4: Average of the difference in the fits to α , α_{AP} and β_ϕ between Barry and desilike. The numbers in brackets also give this uncertainty in the number of standard deviations (taking the weighted standard deviation of the 25 mocks) - thus quantifying the significance of the difference in numbers of σ . We also show for comparison the difference between the fits to the mock means, \bar{P}_w , for each parameter.

Tracer	Δx from the weighted mean of 25 mocks, $\frac{1}{25} \sum_i^{25} x_{\text{Barry}} - x_{\text{desilike}} \left(\frac{\Delta x}{\sigma}\right)$			Δx from fit to data mock mean \bar{P}_w		
	$100\Delta\alpha$	$100\Delta\alpha_{\text{AP}}$	$\Delta\beta_\phi$	$100\Delta\alpha _{\bar{P}_w}$	$100\Delta\alpha_{\text{AP}} _{\bar{P}_w}$	$\Delta\beta_\phi _{\bar{P}_w}$
BGS	1.7 (2.1)	—	1.711 (2.62)	0.03	—	1.950
LRG1	0.7 (0.8)	< 0.1 (< 0.1)	0.570 (0.80)	1.5	0.1	0.947
LRG2	0.2 (0.2)	< 0.1 (< 0.1)	0.550 (0.82)	0.9	< 0.1	0.570
LRG3 + ELG1	0.3 (0.5)	< 0.1 (< 0.1)	0.239 (0.49)	0.8	0.1	0.550
ELG2	0.7 (0.6)	0.3 (0.5)	0.361 (0.44)	2.2	0.7	1.361
QSO	1.6 (1.5)	—	1.13 (1.47)	2.5	—	1.425

Table 5.5: Summary of the choices of priors and ranges used for fitting the second-generation mocks and DR1 data for each DESI tracer. The subscript i denotes different tracer bins (BGS, LRG1, LRG2 etc.) and thus allows us to specify cases where the prior differs for specific cases which are otherwise the same.

Parameter	Gaussian prior: $\mathcal{N}(\mu, \sigma)$		
	Post-recon	Pre-recon	
Σ_S ($\text{Mpc}h^{-1}$)	$\mathcal{N}(2, 2)$	$\mathcal{N}(2, 2)$	
$\Sigma_{\text{nl}, \parallel}$ ($\text{Mpc}h^{-1}$)	$\mathcal{N}(6, 2)$	$\mathcal{N}(x_i, 2)$	$x_{\text{ELG2}} = 8.5$ else $x_i = 9$
Σ_{nl} ($\text{Mpc}h^{-1}$)	$\mathcal{N}(x_i, 2)$	$\mathcal{N}(y_i, 2)$	$x_{\text{BGS}} = 8, x_{\text{QSO}} = 6, y_{\text{BGS}} = 10, y_{\text{QSO}} = 9$
$\Sigma_{\text{nl}, \perp}$ ($\text{Mpc}h^{-1}$)	$\mathcal{N}(3, 1)$	$\mathcal{N}(x_i, 1)$	$x_{\text{BGS}} = 6.5, x_{\text{QSO}} = 3.5$ else $x_i = 4.5$
b	$\mathcal{U}(0.5, 4)$	$\mathcal{U}(0.5, 4)$	
β_ϕ	$\mathcal{U}(-8, 10)$	$\mathcal{U}(-8, 10)$	
α	$\mathcal{U}(0.8, 1.2)$	$\mathcal{U}(0.8, 1.2)$	
$\varepsilon = (\alpha_{\text{AP}}^{1/3} - 1)$ (BARRY)	$\mathcal{U}(-0.2, 0.2)$	$\mathcal{U}(-0.2, 0.2)$	
α_{AP} (DESILIKE)	$\mathcal{U}(0.8, 1.2)$	$\mathcal{U}(0.8, 1.2)$	

between the fitting codes, which can be considered in the error budget for the fits to the data. However, for each of the tracers, the relative difference in the fits to each parameter to the statistical uncertainty on the mean is $\lesssim 2.1\sigma$ for α , $\lesssim 2.6\sigma$ for β_ϕ and $\lesssim 0.5\sigma$ for α_{AP} ; the more weakly constraining tracers tend to show larger differences. This indicates no strong detections of any systematic differences between the codes, and that any systematics are far below the expected precision of the Year 1 data (noting that the statistical error on the data is expected to be 5 times larger than the error on the mean used here).

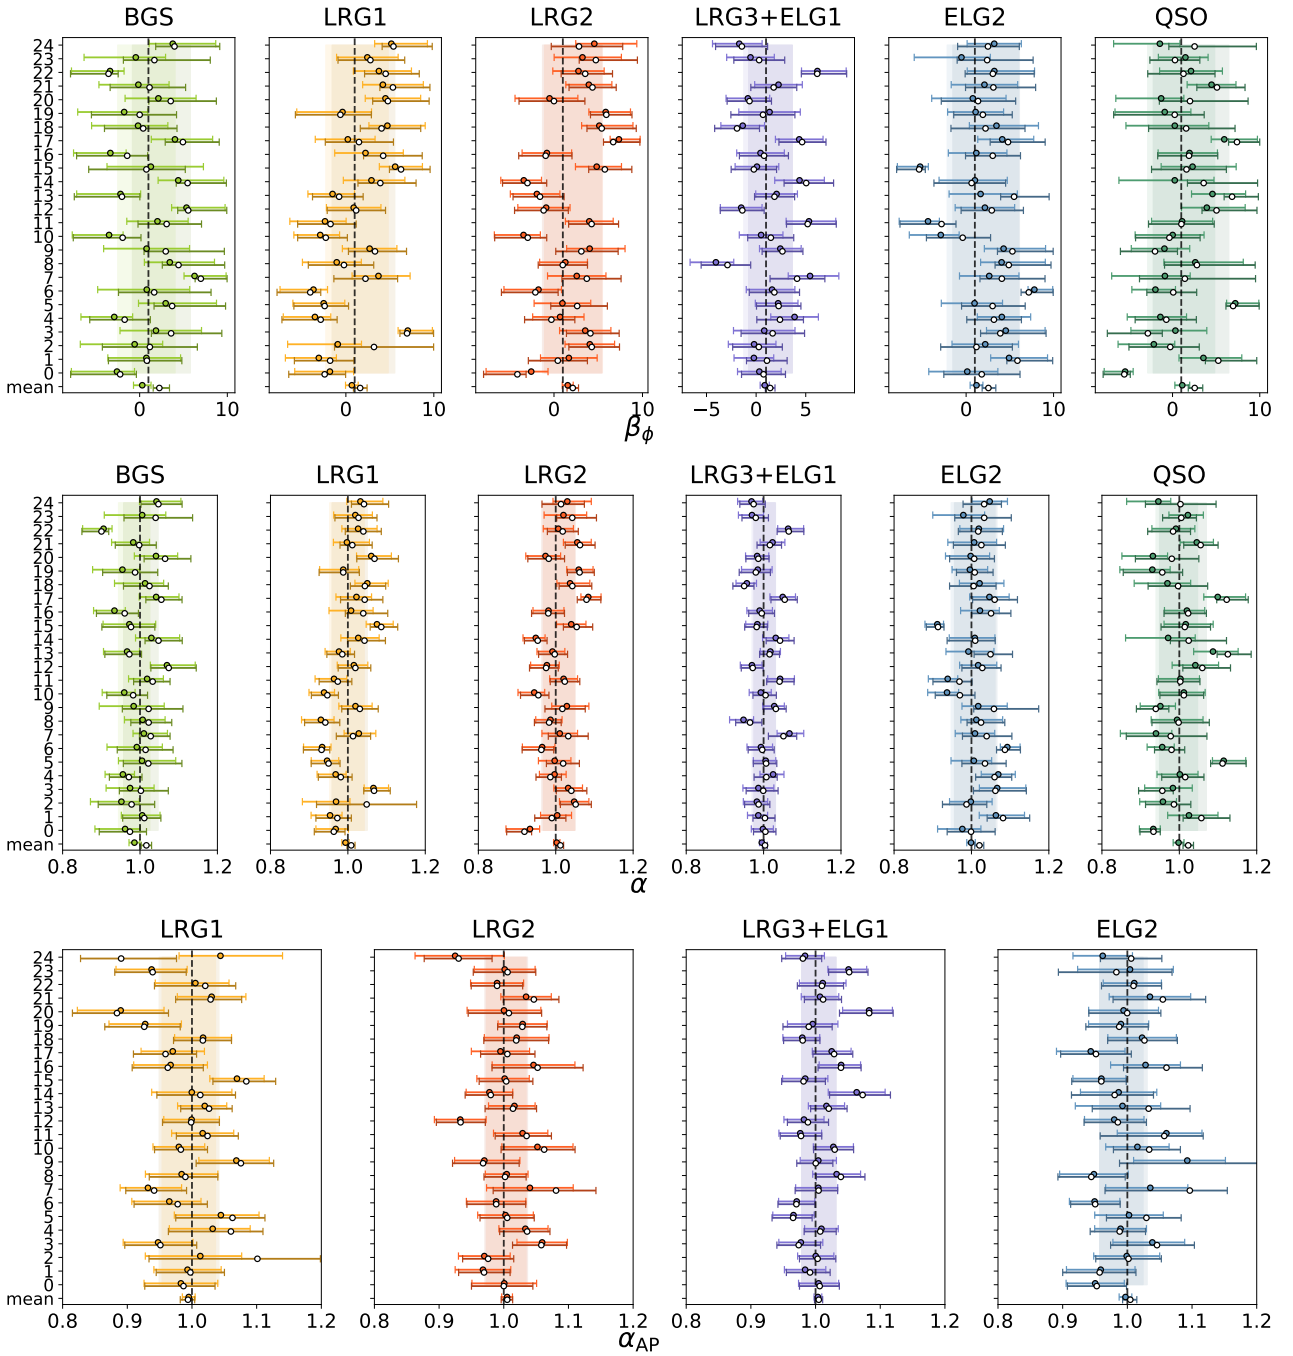


Figure 5.6: Our fits of the 25 mocks and the mean for β_ϕ (top panel), α (middle panel) and α_{AP} (lower panel), for the 6 DESI samples considered here. The coloured points show desilike fits while the white points show the fits using Barry. Shaded regions show one standard deviation about the mean fit for the 25 mocks, where the darker colored shaded regions are desilike and the lighter are Barry. The mean and standard deviation has been weighted by the uncertainty on each mock, so that more poorly constrained fits have less weight in the mean. For the BGS and QSO samples, we only fit an isotropic BAO model, and hence do not report α_{AP} .

Determining a combined fit to β_ϕ from multiple tracers

We are also interested in looking at the combined constraining power for β_ϕ from the tracers in addition to the constraints from individual tracers. In order to do this, I take a mock for each tracer, then importance sample the individual constraints LRG1, LRG2 and ELG2 to weight the fit to β_ϕ for the tracer that obtains the most constraining fits, LRG3+ELG1. In this process, we are effectively computing the posterior distribution of β_ϕ from a combined fit of LRG1, LRG2, LRG3+ELG1 and ELG2,

$$P_{\text{combined}}(\beta_\phi) \propto \prod_i \int d\alpha_i d\alpha_{\text{AP},i} P_i(\alpha_i, \alpha_{\text{AP},i}, \beta_\phi | \xi_i), \quad (5.26)$$

where the index i runs over LRG1, LRG2, LRG3+ELG1 and ELG2 and the individual α and α_{AP} constraints are treated as independent, while β_ϕ is shared. Given that BGS and QSOs obtain weak constraints on β_ϕ , I only use LRGs and ELGs for the combined fits to the mocks and data. This result, which we will refer to as ‘combined fits’ to the mocks can be seen in Figure 5.7. This shows we can obtain reasonable fits to β_ϕ using this methodology and find that combining the multiple tracers improves the uncertainty on β_ϕ by a factor ~ 1.9 compared to the fits to LRG3+ELG1 alone. The fits are also consistent between BARRY and DESILIKE when combining the tracers in this way. The fits here are done using individual mock fits with the ‘baseline’ methodology. The average absolute difference of each mock fit to the combination of the tracers, for fits to the tracers done using `desilike` and Barry is $\Delta\beta_\phi = 0.431$ (or $\frac{\Delta}{\sigma} = 1.282$ where σ is the weighted standard error of the mean of the combined fits). The absolute difference between the fits to the mock mean data (average correlation function of 25 mocks) is $\Delta\beta_\phi = 0.809$.

5.4 Results with DR1 DESI data

Having studied the robustness of our fitting methodology to different BAO fitting software and analysis choices in the previous section, we present fits to the BAO parameters and phase shift β_ϕ with the DR1 data. In all cases, we use the same choices of coefficients for the broadband fitting and choices of priors as done for our fits to second-generation mocks. These choices are the same as those used for the BAO results presented by Adame et al. (2025b,a). For the remainder of this chapter, unless stated otherwise, I will refer to and present our main results using the ‘baseline’ methodology; as before, the ‘baseline’ fits involve post-recon fits to the correlation function using the spline fitting methodology and using the DESILIKE software.

5.4.1 Constraints from individual tracers

In Figure 5.8 we show the individual fits to the BAO distortion parameters α_{\parallel} and α_{\perp} and phase shift β_ϕ for tracers for the DR1 data. The figure also allows a comparison of the baseline results to various individual changes to the fitting pipeline, allowing one to see how robust our results are. One can see how the best fits vary for each tracer under various choices;




Figure 5.7: The fits to β_ϕ from the combinations of the LRGs and ELGs (all tracer bins except QSOs and BGS) with the second-generation mocks, for *desilike* (black points) and *Barry* (white points). The shaded regions show the (weighted) 1-standard deviation about the weighted mean for the fits to individual mocks, the fit to the mean of the mocks is shown at the bottom of the plot.

- when we do not apply reconstruction to the data,
- when we use the polynomial broadband fitting methodology rather than the spline approach,
- when we fit the power spectrum rather than the correlation function,
- or when we produce the fits with *BARRY* rather than *DESILIKE*.

Compared to the fits presented in Adame et al. (2025a) for DESI DR1 (see Figure 12 in their work which is analogous to our Figure 5.8), there is slightly less consistency in our fits presented when the pipeline is varied; the results shown in Adame et al. (2025a) are more statistically robust to changes. This is most noticeable (and expected) for the pre-recon fits, due to the fact that not applying reconstruction significantly weakens the constraints. However, we can expect that for all of the fits, the

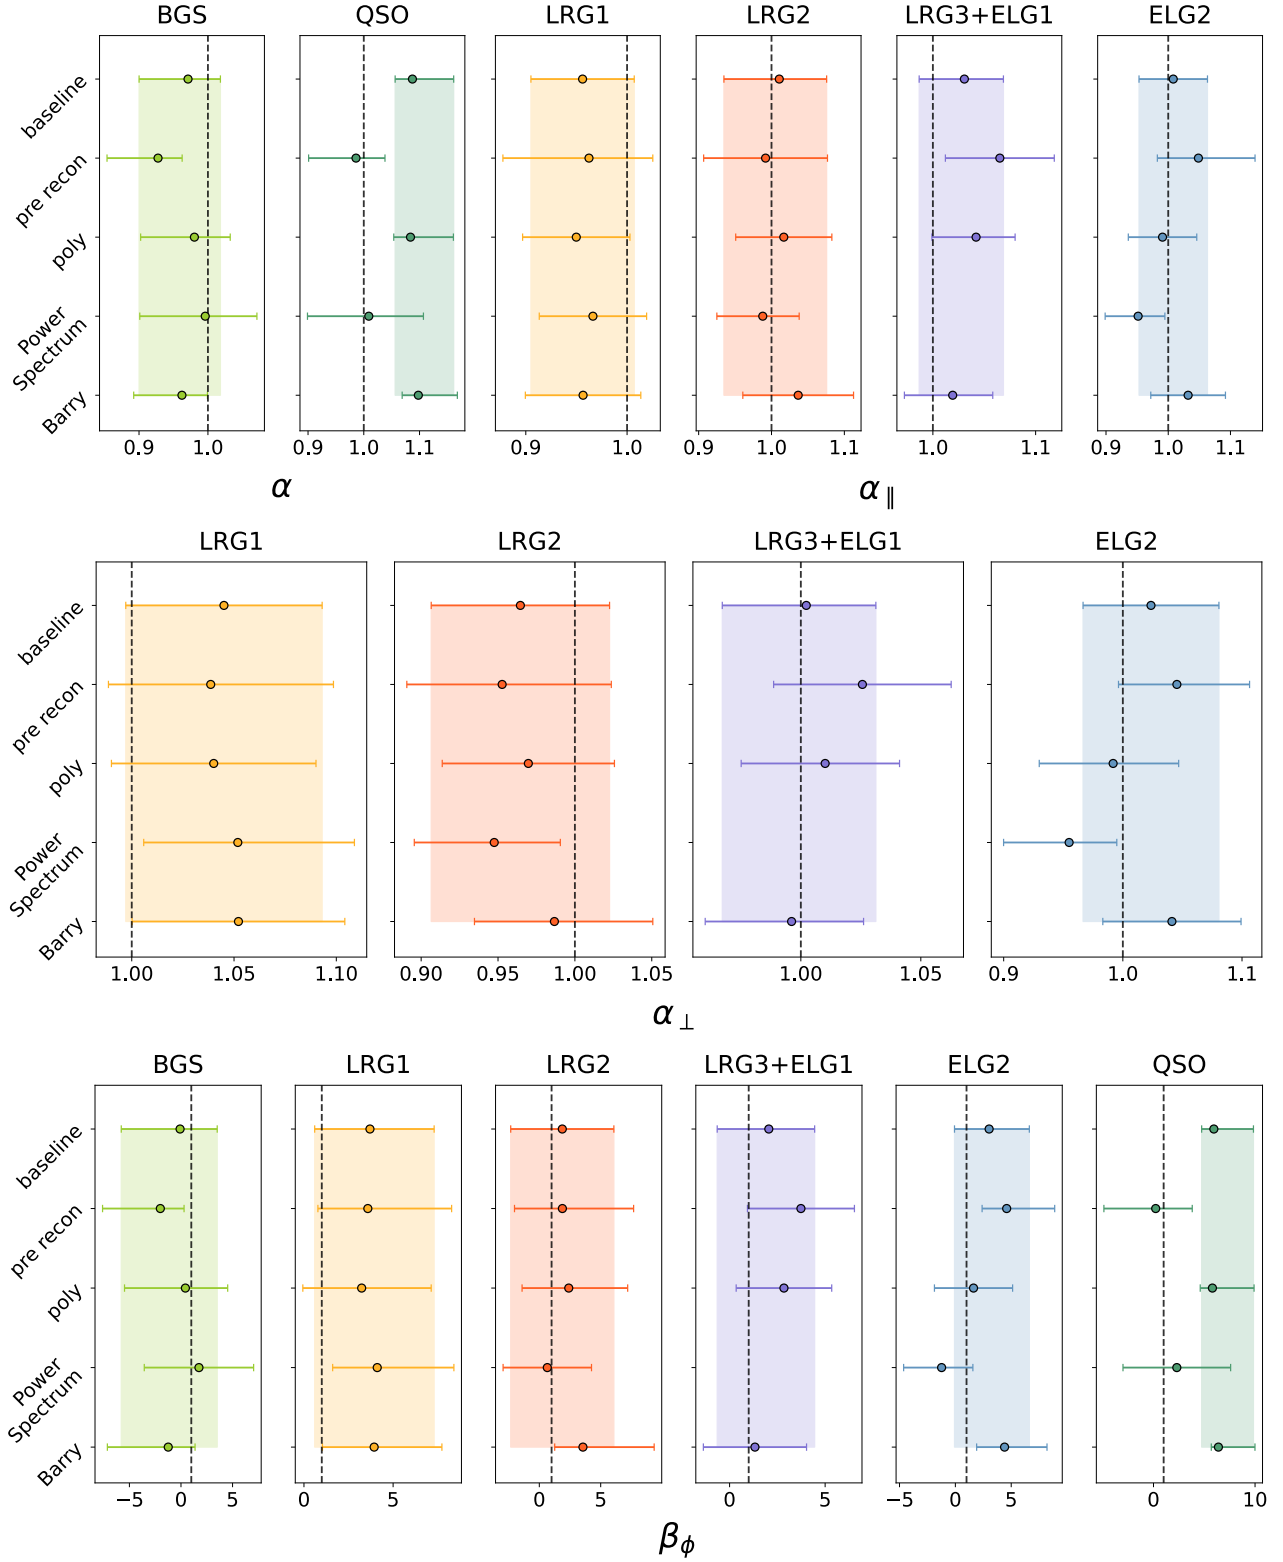


Figure 5.8: Our fits to the DESI tracers in DR1. For ease of interpretation for the distortions to physical distances, we plot α_{\parallel} , α_{\perp} in place of α and α_{AP} and additionally β_{ϕ} . For BGS and QSOs we only study the isotropic scaling; as such we only plot α .

strong degeneracy between α and β_{ϕ} should lead to less constraints that move along this degeneracy direction due to statistical noise. We also do not expect very strong constraints from BGS or QSO tracers given the results with the mocks showed that these tracers are much less constraining when β_{ϕ} is allowed to vary, and this is apparent in our results.

It is also important to validate that our fits to α and α_{AP} are consistent with those of the DESI DR1 results presented in Adame et al. (2025b,a) and that the additional uncertainty in our fits when β_ϕ is allowed to vary does not indicate that the uncertainty in the DESI DR1 results are underestimated. To check this, we recalculated these parameters in a thin slice of the parameter space for β_ϕ by importance sampling the data to impose a prior. This was done by taking the *Planck* 2018 MCMC chain for N_{eff} from fits to the $\Lambda\text{CDM} + N_{\text{eff}}$ model, which gives $N_{\text{eff}} = 2.99 \pm 0.17$. In principle, we expect that this is equivalent to rerunning the fits with a very tight informative prior on β_ϕ from *Planck*. In doing so, we found the best fits to the parameters expressed as α and α_{AP} were fully consistent with the DR1 results presented in Adame et al. (2025a). The fits and comparison to the numbers in Adame et al. (2025a) can be seen in Table 5.6. The 1σ uncertainty on α_{\parallel} varies by approximately 5 – 8% for different tracers, and for α_{\perp} there is a 3 – 15% variation in the 1σ uncertainty. We additionally looked at the constraints in the case when the width of the *Planck* prior for N_{eff} was increased by a factor of 10 times; the fits are still consistent with those presented in Adame et al. (2025a) and only the fits to β_ϕ are weakened.

We also show the correlation function data and best fit models for each tracer for our baseline fits in Figure 5.9. Unlike the results presented in Adame et al. (2025a), the model here shows the fit to the data when β_ϕ is included in the analysis. For comparison, Figures 5 and 7 in Adame et al. (2025a) show the model fits compared to the data for the standard BAO analysis. The χ^2 goodness-of-fit and number of degrees of freedom of each tracer is given in Table 5.7 and in all cases these are comparable to those presented in Table 15 of Adame et al. (2025a) and demonstrate excellent fits.

Figure 5.10 shows the individual fits to each tracer and the posterior distributions for the parameters of interest. The triangle plot on the left allows one to see the posterior distribution for each tracer for α , α_{AP} (excluding BGS and QSOs) and β_ϕ . Each individual bin for the tracers gives relatively consistent fits on β_ϕ (we do not expect consistency for the α , α_{AP} as they are redshift dependent). The BGS and QSO constraints show the largest differences, but as mentioned previously, this is likely because QSOs and the BGS data individually are not able to constrain the β_ϕ parameter in addition to α strongly. Assuming a flat ΛCDM cosmological model, the anisotropic BAO constraints can also be uniquely mapped to a constraint on Ω_m and $r_s h$ (see equation 5.8 and equation 5.9). These are shown in the right-hand panel of Figure 5.10. As is shown in the DESI DR1 results in Figure 2 of Adame et al. (2025b), the contours for $r_s h$ and Ω_m have a rotating degeneracy direction for different tracers (due to the changing redshift of each tracer bin), but the constraints between different tracers are internally consistent and the combination allows for a tighter constraint on all parameters.

5.4.2 Combined constraints and cosmological interpretation

As the different tracers give consistent cosmological results, I can combine them using the methodology described in section 5.3.2 to combine the ELG and LRG mocks. The left panel of Figure 5.11 focuses on just the 1-dimensional posterior for β_ϕ and includes the result for the combined fit to the tracers as a black-dashed line. The right panel of Figure 5.11 shows the contours for the combined fits for Ω_m

Table 5.6: Constraints on β_ϕ , α and α_{AP} from the DR1 data baseline results. We also include the case where a tight prior is included on β_ϕ by importance sampling with the *Planck* 2018 $\Lambda\text{CDM} + N_{\text{eff}}$ MCMC chain. Results quoted from Adame et al. (2025a) are also presented for comparison, demonstrating excellent agreement between our work at previous DESI results when an informative prior on β_ϕ is included.

Data	β_ϕ	α	α_{AP}
<i>(β_ϕ free)</i>			
BGS	$-0.1^{+4.2}_{-5.0}$	$0.971^{+0.05}_{-0.066}$	—
LRG1	3.7 ± 3.2	1.014 ± 0.044	0.915 ± 0.043
LRG2	1.9 ± 4.2	0.979 ± 0.057	1.048 ± 0.048
LRG3+ELG1	2.1 ± 2.6	$1.012^{+0.031}_{-0.035}$	1.028 ± 0.030
ELG2	$3.0^{+3.6}_{-3.2}$	1.018 ± 0.052	0.986 ± 0.045
QSO	$5.9^{+4.0}_{-2.3}$	$1.087^{+0.073}_{-0.045}$	—
<i>(β_ϕ free + <i>Planck</i> prior on N_{eff})</i>			
BGS	0.97 ± 0.04	0.982 ± 0.020	—
LRG1	0.97 ± 0.04	0.979 ± 0.012	0.915 ± 0.041
LRG2	0.98 ± 0.04	0.966 ± 0.013	$1.047^{+0.047}_{-0.052}$
LRG3+ELG1	0.97 ± 0.04	0.998 ± 0.008	1.026 ± 0.029
ELG2	0.97 ± 0.04	0.987 ± 0.015	0.988 ± 0.043
QSO	0.98 ± 0.04	0.998 ± 0.024	—
<i>(Results presented in Adame et al. 2025a)</i>			
BGS	—	0.983 ± 0.019	—
LRG1	—	0.979 ± 0.011	0.915 ± 0.037
LRG2	—	0.966 ± 0.012	1.046 ± 0.043
LRG3+ELG1	—	0.998 ± 0.008	1.026 ± 0.028
ELG2	—	0.988 ± 0.015	0.990 ± 0.047
QSO	—	1.002 ± 0.026	—

Table 5.7: The χ^2 and degrees-of-freedom (dof) for each of the baseline fits to the different DESI tracers when β_ϕ is allowed to vary (corresponding to the data shown in Figure 5.9).

Tracer	χ^2/dof
BGS	15.8/14
LRG1	40.8/36
LRG2	40.6/36
LRG3+ELG1	31.6/36
ELG2	59.8/36
QSO	28.8/14

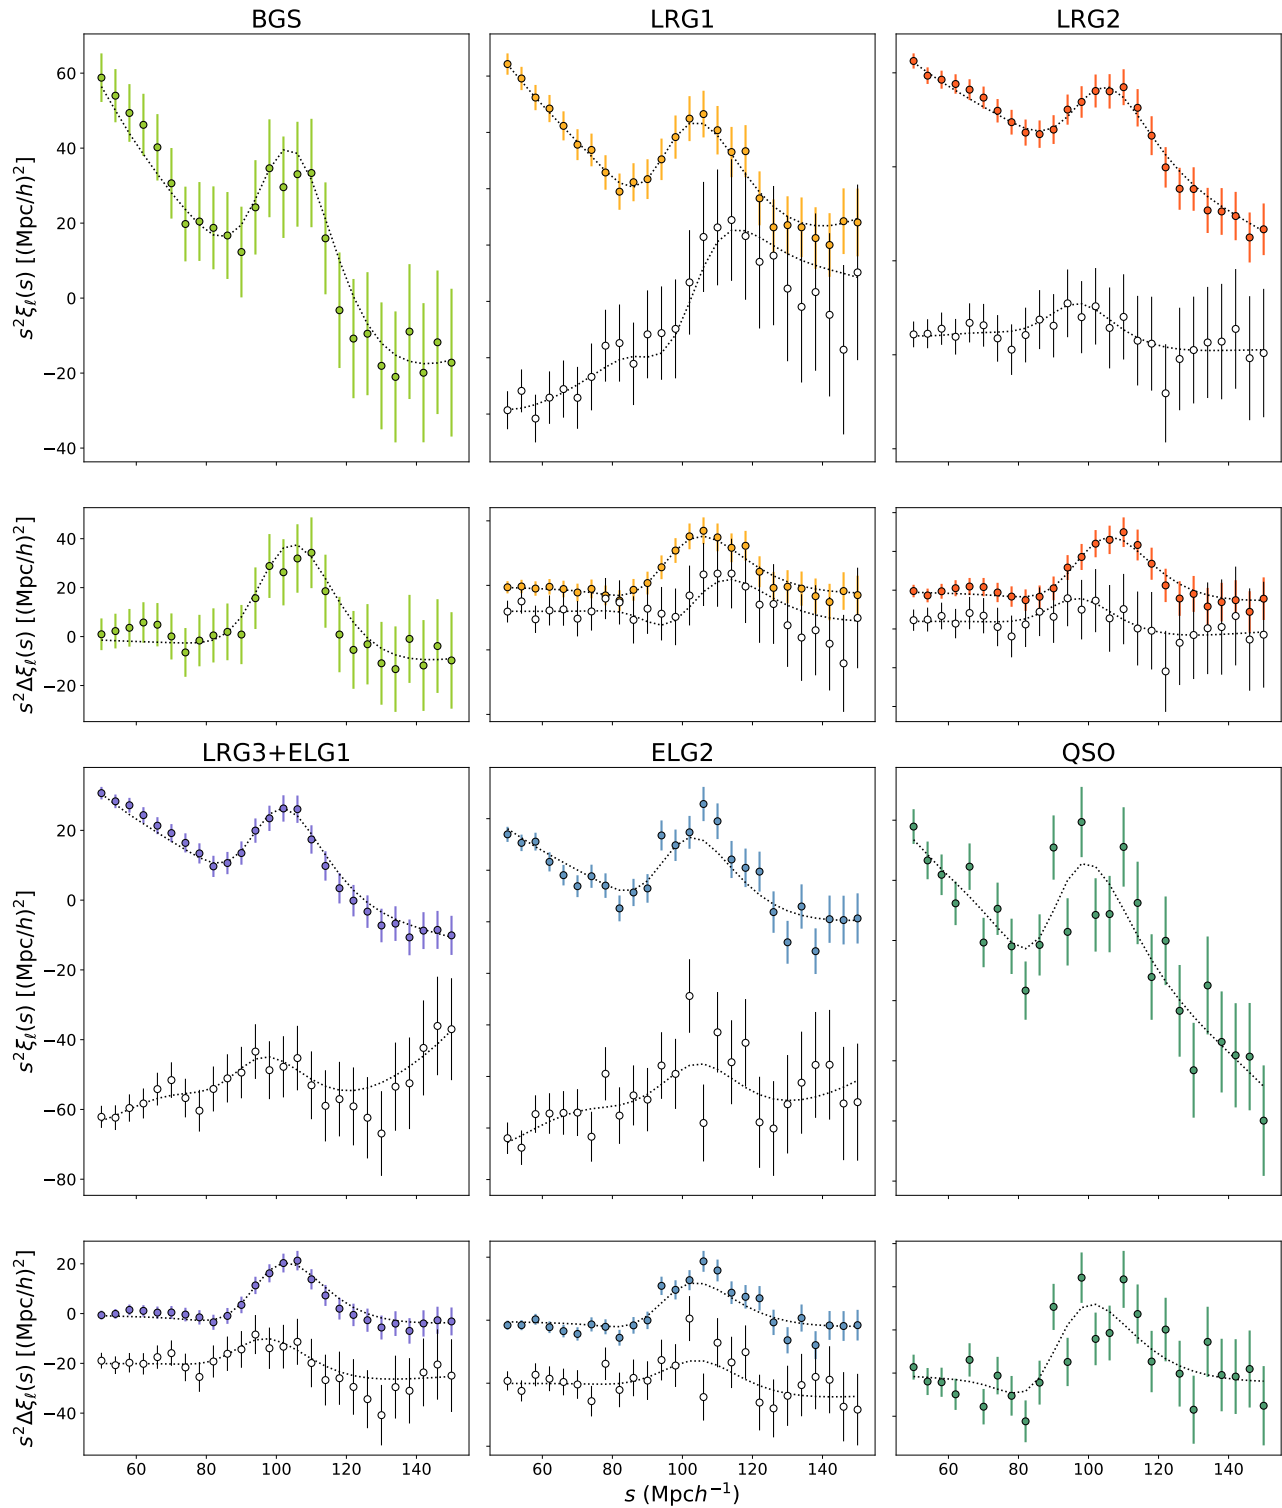


Figure 5.9: The best fit models for the DESI tracers in the DR1 data with β_ϕ varying in the analysis, for comparison to the data. The coloured data shows the best fit models and data for the monopole, and the grey shows the quadrupole. The smaller panels show the isolated BAO peak for each tracer by subtracting the smooth model from the data. For visual clarity a small vertical offset has also been applied to the quadrupole for these panels.

and $r_s h$ when a flat Λ CDM model has been assumed.

Additionally, I show various results when a prior on α and α_{AP} from *Planck* is included (following the approach used by Baumann et al., 2019) by importance sampling the fits to each tracer with the

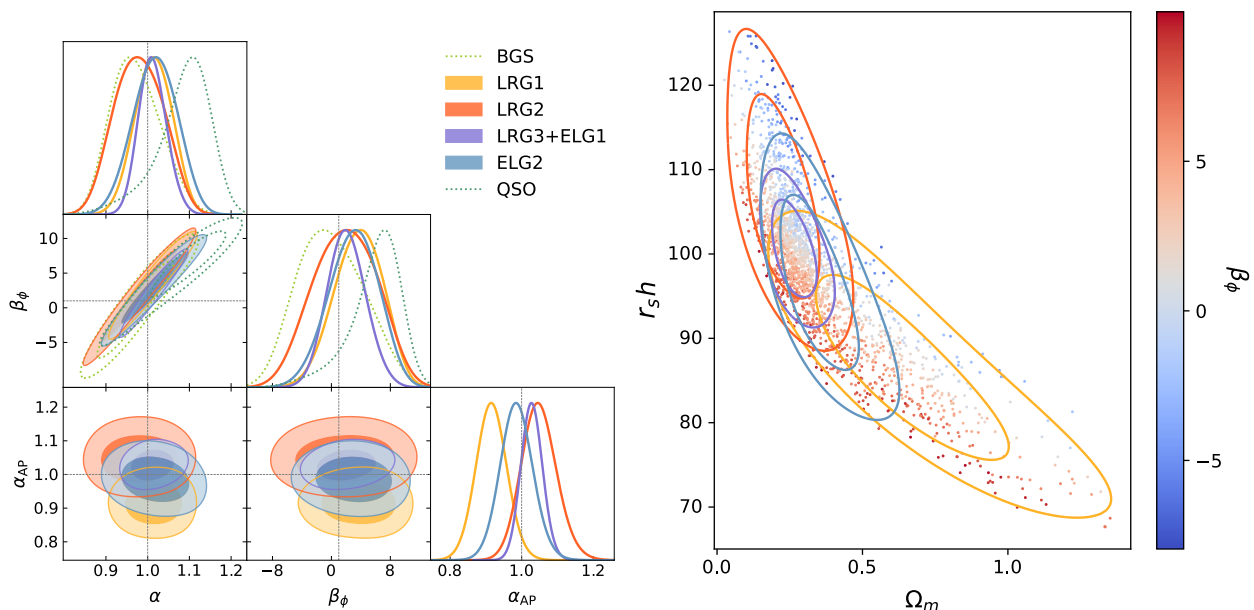


Figure 5.10: Left panel: contours for α , α_{AP} and β_ϕ for the baseline fits to the DESI DR1 data. The fits by QSOs and BGS which are not able to constrain β_ϕ well in the DR1 data are shown as dotted unfilled contours. Right panel: contours for Ω_m , $r_s h$ (Mpc) for the baseline fits to the DESI DR1 data, for individual tracers. The scatter points and colorbar show the value of β_ϕ which has a degeneracy with $r_s h$.

information from the *Planck* 2018 MCMC chains for TT+TE+EE+ low ℓ + lowE.¹⁷ In these cases, the combined fit for β_ϕ is obtained using the same method as previously (equation 5.26), but only after the individual posterior distribution function for each tracer bin has been importance sampled with the *Planck* data. In other words, I have computed a new combined posterior as

$$P_{\text{combined}}(\beta_\phi) \propto \prod_i \int d\alpha_i d\alpha_{AP,i} P_i(\alpha_i, \alpha_{AP,i}, \beta_\phi | \xi_i) \times P_{\text{Planck}}(\alpha_i, \alpha_{AP,i} | \Omega_m, r_s h, w, N_{\text{eff}}), \quad (5.27)$$

in which Ω_m , $r_s h$, N_{eff} and w (the matter density, sound horizon, effective number of neutrino species and in some cases additional parameters such as the dark energy equation of state, EOS) are constrained by *Planck* and i is the set of DESI tracers. Generally for the Λ CDM case in which dark energy is a cosmological constant I assume the curvature density $\Omega_k = 0$ and $\Omega_\Lambda = 1 - \Omega_m$ where Ω_Λ drives the accelerated expansion of the Universe. Each density component has an EOS and this has w fixed to $w = -1$ for a cosmological constant. When adding *Planck*-based priors I perform the same calculation as in equation 5.8 and equation 5.9, but not necessarily limited to the Λ CDM case (in which case there is no longer a unique mapping from cosmological to BAO parameters). We can allow for a model in which the EOS of dark energy varies and thus w is an additional free parameter. I also consider the case in which the dark energy EOS is parameterized by $w(a) = w_0 + w_a(1 - a)$ or when A_{lens} is a free parameter. This is discussed later in this section.

¹⁷This prior is *not* the same as the prior included earlier from *Planck* on N_{eff} . This prior is only for the case we want to include information that impacts the constraints on α and α_{AP} , and allow the BAO data alone to constrain β_ϕ . However I do only include information from *Planck* on α and α_{AP} from chains in which N_{eff} varies freely.

By including the additional information from *Planck* about Ω_m and $r_s h$ for the Λ CDM + N_{eff} case, the uncertainty on α and α_{AP} is reduced for each tracer. Consequently, this allows for better constraints not only on α but on β_ϕ due to their strong degeneracy. This can be seen from the purple and pink dashed lines in the left-hand panel of Figure 5.11, where, interestingly, the best fit value for β_ϕ does not shift noticeably when the information from the *Planck* Λ CDM + N_{eff} chains as prior information, with or without lensing information included. In including information from the *Planck* chains, we do not fold in any explicit CMB information on N_{eff} in order to report only the constraining power available on β_ϕ from the physical shifting of the BAO wiggles; the constraints on N_{eff} from *Planck* include additional information on N_{eff} from the CMB that is degenerate with other cosmological parameters, in the same way that there are additional signals in the galaxy clustering beyond the BAO phase shift that can constrain N_{eff} (see Figures 5.1 and 5.2).

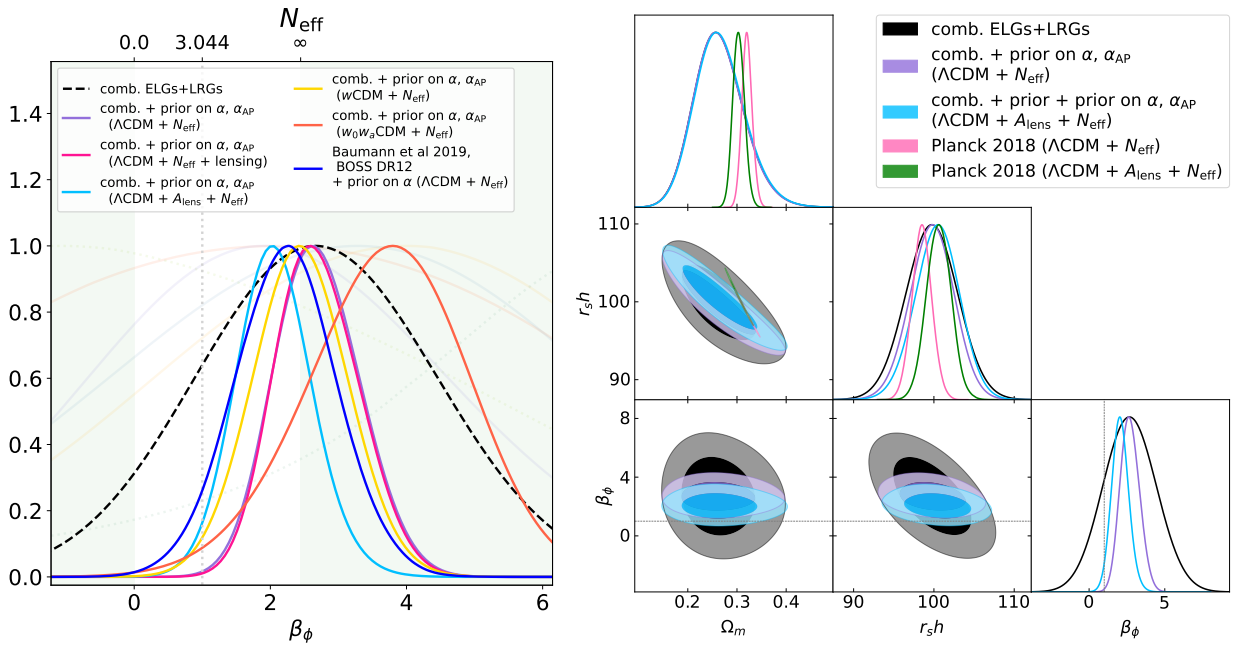


Figure 5.11: Left panel: the 1D posterior fits for β_ϕ , including a combined fit to the tracers (excluding BGS and QSOs, black dashed line), and various combined fits when a CMB prior for α and α_{AP} has been included by importance sampling (corresponding to the fits shown in Table 5.8). The priors all use fits to the *Planck* chains for TT+TE+EE + low ℓ + lowE. I include a prior from fits to Λ CDM + N_{eff} (purple), the same but with lensing information included (pink), the case for a Λ CDM prior (blue), Λ CDM + A_{lens} (red) and finally w CDM (orange). The faint lines show the 1D posteriors of β_ϕ for each tracer that are shown in Figure 5.10. The dark blue line shows the fit to BOSS DR12 by Baumann et al. (2019). Right panel: a combined fit for Ω_m , $r_s h$ and β_ϕ to three LRG bins and the ELG2 bin is shown as a black contour. The fits by QSOs and BGS are not included here. I also show the fits when each tracer has a CMB prior for α and α_{AP} by importance sampling using the *Planck* chains for fits to TT+TE+EE + low ℓ + lowE in the context of Λ CDM + N_{eff} . For comparison the *Planck* chains are also shown for just Ω_m and $r_s h$.

Table 5.8 lists the combined constraints on β_ϕ from the DESI DR1 BAO data, and when the chains have been importance sampled with *Planck* data. For the DESI data alone, the measurement of β_ϕ is consistent with unity within 1σ , and suggests $\beta_\phi > 0$ at $> 68\%$ confidence. Including the *Planck* prior on the BAO parameters (either with or without lensing) sharpens this constraint by a factor of ~ 3

Table 5.8: Constraints on β_ϕ from the combined DESI DR1 BAO tracers, alongside cases when a prior on the BAO α parameters has been included from the *Planck* data by weighting the chains from the BAO fits. I also show the number of (Gaussian) σ away each best fit to β_ϕ is from zero ($N_{\text{eff}} = 0$) and unity ($N_{\text{eff}} = 3.044$).

Data	Ω_m	$r_s h$ (Mpc)	β_ϕ	$\Delta\sigma(\beta_\phi > 0)$	$\Delta\sigma(\beta_\phi > 1)$
DESI DR1 BAO	$0.263^{+0.046}_{-0.054}$	99.7 ± 3.2	2.70 ± 1.70	1.6	1.0
DESI DR1 BAO + <i>Planck</i> prior $\alpha, \alpha_{\text{AP}}$	$0.263^{+0.045}_{-0.054}$	99.7 ± 2.7	$2.70^{+0.60}_{-0.67}$	4.3	2.6
DESI DR1 BAO + <i>Planck</i> prior $\alpha, \alpha_{\text{AP}}$ w/ lensing	$0.263^{+0.045}_{-0.054}$	99.8 ± 2.7	$2.69^{+0.59}_{-0.66}$	4.3	2.7
DESI DR1 BAO + <i>Planck</i> prior $\alpha, \alpha_{\text{AP}}$ w/ varying A_{lens}	$0.265^{+0.046}_{-0.054}$	100.4 ± 2.7	2.05 ± 0.55	3.7	1.9
DESI DR1 BAO + <i>Planck</i> prior $\alpha, \alpha_{\text{AP}}$ w/ varying w	—	—	2.44 ± 0.70	3.4	2.0
DESI DR1 BAO + <i>Planck</i> prior $\alpha, \alpha_{\text{AP}}$ w/ varying w_0, w_a	—	—	$3.7^{+1.2}_{-1.1}$	3.2	2.3

while keeping the maximum posterior value fixed. This leads to a significant detection of a non-zero phase shift ($\beta_\phi > 0$ at $> 4\sigma$), but also leads to a constraint that prefers a value of $N_{\text{eff}} > 3.044$ at greater than 95% confidence. The maximum *a posteriori* value is also large enough that it would correspond to an unphysical value of N_{eff} , although I note that a wide range of N_{eff} is allowed given the non-linear mapping from this to β_ϕ .

At first glance this then seems to suggest a phase shift in the BAO that, when combined with the *Planck* prior, prefers either a larger number of effective relativistic species than predicted by the Standard Model, or the presence of other non-standard physics which introduces a phase shift in the BAO (e.g., non-standard neutrino physics, Kreisch et al. 2020; or non-adiabatic primordial density fluctuations, Baumann et al. 2016). A similar finding was reported by Baumann et al. (2019) for the BOSS DR12 BAO data, who found $\beta_\phi = 1.2 \pm 1.8$ from combining two redshift bins, and $\beta_\phi = 2.22 \pm 0.75$ after including a prior from *Planck* by importance sampling. These two sets of results are remarkably consistent in central value and precision.

Compared to the fits from the BAO data alone or when priors from *Planck* are included in the fits, *Planck* alone prefers a larger Ω_m and smaller $r_s h$. The discrepancy in the best fit for Ω_m is slightly greater than the case for the fits to the BAO DR1 data fits when β_ϕ is fixed, in which case Adame et al. (2025b) reports $\Omega_m = 0.295 \pm 0.015$. However, when β_ϕ is free the data prefers a value for $r_s h$ which is lower but more consistent with the best fit from *Planck*. This may be related to the degeneracy that can be seen between $r_s h$ and β_ϕ in Figure 5.11; this parameter degeneracy arises since both $r_s h$ and β_ϕ shift the BAO wiggles along the k -axis.

To investigate the impact of the *Planck* priors on the fits to β_ϕ further, I also followed the tests done in Baumann et al. (2019) and included a *Planck* prior for the Λ CDM model with A_{lens} (the lensing amplitude) allowed to be free in the analysis. I find a result $\beta_\phi = 2.05 \pm 0.55$, which is both a tighter constraint and closer to unity. The result is that our detection of a non-zero phase shift becomes now

$> 3\sigma$, but the tension with the Standard Model prediction is reduced somewhat from 2.6 to 1.9σ . This fit is still consistent with the result from the BOSS DR12 data at the 1σ level.¹⁸ During this analysis we did not have access to the *Planck* 2021 (PR4) analysis (Tristram et al., 2024), however we note in Montefalcone et al. (2025c) who measure the phase shift in the CMB using both *Planck* 2018 and 2021 analyses find a shift that is more consistent with the standard model expectation.

Lastly, while the *Planck* data does tend to increase the tension between our β_ϕ and the Standard Model prediction, the prior from *Planck* is based on chains from sampling in the Λ CDM model, while DESI data (in combination with CMB and Type Ia Supernovae) moderately prefer a w_0w_a CDM (time-varying Dark Energy) model (Adame et al., 2025b). Given that we are including a prior from a model that is not preferred by DESI constraints on the BAO scaling parameters, this might encourage the combined fit to β_ϕ from DESI with a Λ CDM *Planck* prior to take a higher than expected value of β_ϕ . As such we include fits from a full MCMC fit to a time-varying Dark Energy model with varying N_{eff} . We present results for the w CDM model (in which case we have the normal flat Λ CDM model + N_{eff} but w is allowed to vary as a single free parameter). I find this tends to allow the best fit value of β_ϕ to shift to a value that also reduces the tension with the standard model with $\beta_\phi = 2.44 \pm 0.70$. This may be due to the additional flexibility in the model. In the case for the time-varying EOS model w_0w_a CDM + N_{eff} , where the additional parameters w_0 and w_a vary in the EOS $w(a) = w_0 + (1 - a)w_a$, I find that the additional free parameters lead to weaker constraints and has slightly greater tension with the theoretical expectation for N_{eff} , with $\beta_\phi = 3.7^{+1.2}_{-1.1}$.

5.5 Conclusions

In this work, we have applied the scheme derived in Baumann et al. (2019) to measure the phase shift in the BAOs. This extends upon the standard BAO analysis in DESI that does not include the phase shift amplitude β_ϕ . This phase shift is induced by free-streaming particles that propagate faster than the sound speed of the primordial plasma and thus allows one to measure the impact of free-streaming neutrinos that contribute to N_{eff} in the early Universe.

Compared to the analysis on BAOs by Baumann et al. (2019), we have extended the phase shift scheme for an anisotropic BAO fitting analysis, which we have applied to the DESI DR1 BAO data for galaxies and quasars. In the process, we have also extended the publicly available codes `desilike` to create the fits and likewise `Barry` to validate the fitting methodologies used between two separate codes. We tested our pipeline with `desilike` and `Barry` on highly precise cubic box simulations to ensure the robustness of the fitting schemes when applied to data with low statistical noise. Further to this we tested the pipeline and codes on realistic mocks for the various tracers in the DESI DR1 data release before applying the code to the real DR1 data.

Our analysis on the highly precise cubic box simulations has highlighted the need to address systematics in the BAO fitting methodology for future constraints with the 5 year DESI data. Further investigation will be required in order to determine a more robust way to measure the phase shift for

¹⁸Baumann et al. (2019) find $\beta_\phi = 1.53 \pm 0.83$ for Λ CDM + N_{eff} + A_{lens} .

different broadband fitting methodologies for the power spectrum. We also expect there may be a need to address other differences seen in the fits when varying the analysis pipeline if the measurements of the BAO distortion parameters and β_ϕ are larger relative to a reduced statistical uncertainty in future measurements. Potentially, it will also be necessary to consider in greater detail the impact of the choice of fiducial cosmology or template cosmology. Further to this, while allowing β_ϕ to vary significantly weakens constraints on α due to the strong degeneracy between these parameters, it may be interesting to consider where fixing β_ϕ could lead to small biases in α , particularly as we continue to report more precise measurements with upcoming data from DESI. This may also be especially important to check given the present hints for time-varying dark energy seen in DESI BAO data.

We have measured the phase shift in the DR1 DESI data from the combination of LRGs and ELGs across various redshift bins to be $\beta_\phi = 2.7 \pm 1.7$. This is larger than the measurement of $\beta_\phi = 1.2 \pm 1.8$ in Baumann et al. (2019), however both measurements agree within the 1σ uncertainties. When we add a prior on α and α_{AP} from *Planck* 2018 measurements of the CMB, we obtain $\beta_\phi = 2.7^{+0.60}_{-0.67}$, which as before is larger than the value of 2.22 ± 0.75 given in Baumann et al. (2019) with the same prior included. In cosmological models with only Standard Model particles the phase shift can be interpreted as due to N_{eff} . However here the phase shift measurements of $\beta_\phi \gtrsim 2.44$ do not correspond to a physical value for N_{eff} . The amplitude of the phase shift may hint at physics beyond the standard model since it is not entirely consistent with an interpretation via N_{eff} . However, we emphasize here that the measurement of β_ϕ is for a physical phase shift that is present in the BAOs, and in a non-standard model could be impacted by non-standard neutrino physics (Kreisch et al., 2020) or by non-adiabatic fluctuations (Baumann et al., 2016). Furthermore, our measurements of β_ϕ from DESI BAOs alone are still consistent with physical values of N_{eff} at the 1σ level.

In the future, the analysis can potentially be applied to the DESI Y3 or DESI Y5 data, allowing for these future analyses to extend beyond the standard approach that aims to only measure the dilation parameters. Baumann et al. (2019) forecasts that with the DESI Y5 data it may be potential to obtain $\sigma(\beta_\phi) \sim 0.3$, which is a factor of ~ 6 better than the constraints reported here for the BAO only data, although these forecasts may not include the impact of allowing α_{AP} to vary freely as is done in this work. In order to provide competitive constraints on future data with DESI Y5, it will be necessary to begin to improve the systematics on our measurements that were analysed in detail in section 5.3.

In future work, it will also be interesting to consider the impact of a phase shift due to particles that may have different properties to Standard Model neutrinos, such as interacting neutrinos. It may be expected that such neutrinos could change the functional form of the phase shift that is given by equation 5.15 (which has been studied in Choi et al., 2018) due to the fact that neutrino interactions can change the epoch at which they decouple and free-stream (Kreisch et al., 2020; Camarena et al., 2023a). This may be particularly interesting in the future when stronger constraints on β_ϕ are possible, particularly in light of results that find a bimodal distribution for a parameter that characterizes the interaction strength of strongly interacting neutrinos (Camarena et al., 2023a; Kreisch et al., 2024). It may also be interesting for future work to consider the ability of DESI+*Planck* to constrain the phase shift in the context of more models beyond Λ CDM, or consider adding information from CMB

probes for the phase shift induced by free-streaming relics, such as the constraints found by Follin et al. (2015).

5.6 Extended results with DR2 DESI data

In this section, I include an additional set of results (these results were the product of work by myself and Hugo Rivera-Morales) that extend upon those included in our publication Whitford et al. (2025) with just the DR1 data, using DESI BAO data from the second data release (DR2). I also note here that our detection level of $\beta_\phi > 0$ in DESI DR1 and DR2 has been included in the results of (Elbers et al., 2025). The DR2 data has all the same groups of tracers as shown in the DR1 data, however with a greater number of galaxies we can expect stronger constraints on each parameter; the number of objects per tracer is shown in Table 5.9. Here I present the constraints when we use all the same choices of priors on every free parameter in the analysis as for the DR1 data presented previously. We also use the ‘baseline’ methodology here. However, since the QSOs should allow for a more constraining fit, we allow α_{AP} to be free instead of fixed for this tracer. Furthermore, in the DR1 data we fit the correlation on scales in the range $r = [50, 150] \text{ Mpc } h^{-1}$, but for the DR2 data this is altered to scales of $r = [60, 150] \text{ Mpc } h^{-1}$.

Table 5.9: Summary of DESI DR2 data tracers properties. The LRGs are split across three redshift bins (LRG1, LRG2, LRG3) and ELGs across two (ELG1, ELG2).

Tracer	z	N objects
BGS	0.1 - 0.4	1,188,526
LRG1	0.4 - 0.6	
LRG2	0.6 - 0.8	4,468,483
LRG3	0.8 - 1.1	
ELG1	0.8 - 1.1	6,534,844
ELG2	1.1 - 1.6	
QSO	0.8 - 2.1	1,461,588

5.6.1 Constraints from individual tracers

The fits to the BAO parameters and phase shift are shown in Table 5.10, posteriors for these parameters are shown in Figure 5.12.

I also produced fits using BARRY to validate the fits were consistent between different codes. For the most part, the fits were very consistent between the two codes, although there was a difference in the best fits to α and β_ϕ for the ELG2 tracer bin, that most likely can be attributed to shifts along the degeneracy direction of these parameters. To attempt to improve the consistency in these constraints, I tested altering various different modelling choices in BARRY, as there are some known differences in the modelling of the power spectrum in BARRY and DESILIKE. However, this did not make the fits




Figure 5.12: Posteriors for fits to the BAO and phase shift parameters in the DESI DR2 data for individual tracers.

tracer	β_ϕ	α	α_{AP}
BGS	2.9 ± 2.8	1.007 ± 0.032	NA
LRG1	3.6 ± 2.2	1.024 ± 0.028	0.951 ± 0.027
LRG2	0.0 ± 1.8	$0.963^{+0.021}_{-0.024}$	0.984 ± 0.023
LRG3 + ELG1	2.0 ± 1.3	1.000 ± 0.016	1.026 ± 0.016
ELG2	1.7 ± 1.6	1.000 ± 0.023	1.024 ± 0.024
QSO	$5.3^{+3.5}_{-2.5}$	$1.070^{+0.055}_{-0.043}$	0.993 ± 0.054

Table 5.10: Fits to the BAO and phase shift parameters in the DESI DR2 data for individual tracers, using DESILIKE. The left plot shows the posteriors for each parameter, while the best fit (maximum of the posterior) and error bars are shown in the table on the right.

more consistent. Potentially, since the fits with BARRY have been done using nested sampling with Lange (2023) and fits with DESILIKE have used a variant of MCMC with the Foreman-Mackey et al. (2013) package, some small mathematical differences in the convergence criteria of these codes may allow them to obtain slightly different constraints. I verified that both codes give a reasonable model fit to the DR2 data in either case. We also verified that the results here are consistent with those obtained for the DR2 data in a standard BAO analysis, in which β_ϕ is fixed.

5.6.2 Combined constraints

In Figure 5.13 we show the combined fits to the DESI DR2 tracers. We can see here that the DR2 data provides a constraint, independent of data from *Planck*, which is both stronger than the DR1 constraint

and additionally closer to the expectation $\beta_\phi = 1$. The combined constraint from all tracers (including here BGS and QSOs) gives $\beta_\phi = 2.00 \pm 0.78$, which gives a constraints for $\beta_\phi > 0$ with 2.6σ . This compares to $\beta_\phi = 2.7 \pm 1.7$ for DR1 which gives $\beta_\phi > 0$ at 1.6σ . These results also give $\beta_\phi > 1$ at the 1σ for DR1 and $\beta_\phi > 1$ at 1.3σ for the DR2 data.




Figure 5.13: Combined constraints to β_ϕ for the DESI DR2 data, in comparison to the DR1 data and the BOSS DR12 fits (Baumann et al., 2019).

Finally, when a prior is included from *Planck* on α and α_{AP} , as shown earlier for the DR1 data, we obtain $\beta_\phi = 1.94 \pm 0.28$ for the DR2 data. This compares to $\beta_\phi = 2.70^{+0.60}_{-0.67}$ for the DR1 data. In both cases, adding the information from *Planck* hardly impacts the best-fit value of β_ϕ although the error bars become smaller. The DESI DR2 BAO data + *Planck* thus prefers a $\beta_\phi > 1$ at $\sim 3\sigma$ and $\beta_\phi > 0$ at $\sim 6.9\sigma$. With future DESI data releases, it will be interesting to see if $\beta_\phi > 1$ is still preferred by more constraining data, however more work may need to be done to address how the broadband modelling of the power spectrum impacts the best fit to β_ϕ .

I have explored the constraints on β_ϕ in the context of the combined DESI DR2 BAO + *Planck* datasets; given the increased level of tension between these datasets in the context of Λ CDM shown in (Karim et al., 2025) at 2.3σ , one might suggest that this level of tension indicates the datasets should not be combined in an analysis. However, given the inclusion of β_ϕ generally weakens the cosmology constraints from the DESI BAOs, it is unlikely there is a significant enough level of tension between the datasets that is of concern for this analysis.

The following Chapter 6 is a draft for a manuscript, to eventually be submitted for publication. I have changed the text to use ‘I’ rather than ‘we’ in various places in the text.

Contributor	Statement of contribution	%
Abbé M. Whitford	writing of text	70
	proof-reading	30
	numerical calculations	100
	preparation of figures	100
	initial concept	35
Cullan Howlett	writing of text	5
	proof-reading	20
	supervision, guidance	45
Tamara M. Davis	writing of text	5
	proof-reading	15
	supervision, guidance	25
David Camarena	writing of text	15
	proof-reading	20
	supervision, guidance	20
	initial concept	65
Francis-Yan Cyr-Racine	writing of text	5
	proof-reading	15
	supervision, guidance	10

This project involved a theoretical exploration of the ability to use the change in the phase of BAO and CMB data, that is induced by neutrinos with self-interactions, to constrain the strength of neutrino self-interactions. This idea was suggested first to me by David Camarena while I attended a workshop at the Galileo Galilei Institute of Theoretical Physics in Florence, Italy. I later discussed this project with my supervisors, David and Francis after completing my work on measuring the phase shift that is induced by Standard Model neutrinos in the BAO signal of DESI data. Throughout this project I have received mentorship from David and Francis but the calculations, production of figures and writing of this work has been completed myself.

This research has made use of NASA’s Astrophysics Data System bibliographic services, the astro-ph pre-print archive at <https://arxiv.org/> and the python libraries MATPLOTLIB, NUMPY, SCIPY and PANDAS (Harris et al., 2020; Wes McKinney, 2010; Virtanen et al., 2020b,a; McKinney et al., 2010; Virtanen et al., 2020a; Hunter, 2007). AW thanks the Galileo Galilei Institute of Theoretical Physics in Florence, Italy, for facilitating discussion that led to the research conducted in this paper at the *Neutrino Frontiers* workshop in 2024. David Camarena and Francis-Yan Cyr-Racine were supported by the Robert E. Young Origins of the Universe Chair fund. Francis-Yan Cyr-Racine was also supported by the US National Science Foundation CAREER grant AST-2440096. I also thank Gabriele Montefalcone and the anonymous reviewer of this paper for their useful comments which helped to improve this manuscript.

Chapter 6

Limits on self-interacting neutrinos from the BAO and CMB phase shift

In this chapter, I extend upon the work on Chapter 5, in the context of non-standard neutrino interactions. As discussed earlier, neutrinos with Standard Model interactions free-stream in the early Universe, leaving a distinct phase shift in the pattern of baryon acoustic oscillations (BAO). When isolated, this phase shift allows one to robustly infer the presence of the cosmic neutrino background in BAO and cosmic microwave background (CMB) data independently of other cosmological parameters. While in the context of the standard model, this phase shift follows a known scale-dependent relation, new physics in the cosmic neutrino background could alter the overall shape of this feature. In this chapter, I discuss how changes in the neutrino phase shift could be used to constrain self-interactions among neutrinos. I produce simple models for this phase-shift assuming universal self-interactions, and use these in order to understand what constraining power is available for the strength of such interactions in BAO and CMB data. This work finds that, although challenging, it may be possible to use a detection of the phase to put a more robust limit on the strength of the self-interaction, G_{eff} , which at present suffers from bimodality in cosmological constraints. My forecast analysis reveals that BAO data alone will not provide the precision needed to tightly constrain self-interactions; however, the combined analysis of the phase shift signature in both CMB and BAO can potentially provide a way to detect the impact of new neutrino interactions. These results could be extended upon for models with non-universal interactions.

6.1 Overview

Here I will recap some of the concepts from Chapter 5. As mentioned, in the early Universe, neutrinos quickly decouple from the hot plasma of baryons and photons that allowed sound waves to propagate prior to recombination. These free-streaming neutrinos are referred to as the cosmic neutrino background (CvB). The presence of the CvB impacts the cosmic microwave background (CMB) photons and also the growth of structure we see in the Universe today (see detailed reviews in

Lesgourgues and Pastor, 2006, 2012, 2014). While the CvB itself has not been directly detected due to the nature of neutrinos as very weakly-interacting particles, the impacts of the CvB on cosmological observables have been observed in the CMB (Follin et al., 2015; Aghanim et al., 2020; Montefalcone et al., 2025c), the matter power spectrum (Elgarøy et al., 2002), the Lyman- α forest (Yèche et al., 2017), measurements of baryon acoustic oscillations (BAOs, Baumann et al., 2018, 2019; Adame et al., 2025b; Whitford et al., 2025; Elbers et al., 2025; Karim et al., 2025) and combinations of cosmological data (Tegmark et al., 2004), including terrestrial experiments (Stöcker et al., 2021). Big Bang Nucleosynthesis (BBN) is also able to help constrain neutrino properties (Mangano and Serpico, 2011). These observations constrain the sum of the masses of the different neutrino species $\sum m_\nu$ and the effective number of neutrino species, N_{eff} . The latter parameterises the number of neutrino species and their contribution to the radiation density Ω_r via

$$\begin{aligned}\Omega_r &= \Omega_\nu + \Omega_\gamma \\ &= \Omega_\gamma \left(1 + N_{\text{eff}} \frac{7}{8} \left(\frac{4}{11} \right)^{4/3} \right).\end{aligned}\tag{6.1}$$

N_{eff} is expected to be 3.044 in the Standard Model.¹ As such, altering N_{eff} alters the radiation density, altering the expansion rate in the early Universe. Larger N_{eff} decreases the distance to the sound horizon at recombination, captured by the parameter θ_s in CMB data, and decreases the sound horizon distance to the baryon drag epoch r_s , measured in BAO data. These changes to the expansion rate also impact Silk damping and the growth of structure on scales that enter the horizon prior to matter-radiation equality, as the growth of structure is suppressed in the radiation-domination era. This is reflected in the matter power spectrum and the CMB power spectra, allowing for N_{eff} constraints. Free-streaming neutrinos in the early Universe also induce a distinct imprint on the evolution of the photons and matter perturbations. This imprint, notable as a phase shift in both the CMB and matter power spectrum, is a unique signature of the CvB and can be used to robustly constrain the presence of free-streaming relics (Follin et al., 2015; Baumann et al., 2019; Wallisch, 2019; Montefalcone et al., 2025c) or the lack of those, if there are instead fluid-like relics that do not free-stream (Montefalcone et al., 2025c).

6.1.1 Searching for neutrino physics beyond the Standard Model

Despite these results, neutrinos themselves are still one of the most poorly understood particles, and, for instance, it is not entirely understood how neutrinos gain mass within the context of the Standard Model. While terrestrial experiments are able to constrain the differences in mass between different neutrino species from oscillation experiments (Fukuda et al., 1998; Hirata et al., 1989),² and weak constraints on neutrino masses can be obtained via beta decay or neutrinoless double beta decay (Tanabashi et al., 2018; Zyla et al., 2020; Vitagliano et al., 2020), the tightest upper bounds on neutrino

¹While there are expected to be just three neutrino species or three mass eigenstates, N_{eff} increases to 3.044 to absorb some of the additional energy that is injected into the neutrino distribution during electron-positron annihilation in the early Universe (de Salas and Pastor, 2016; Bennett et al., 2021).

²These experiments measure mass splittings Δm_{ij}^2 between different neutrino eigenstates labelled by i, j , which puts a lower bound on $\sum m_\nu$.

mass to date come from cosmological data that constrain $\sum m_\nu$ (Elbers et al., 2025; Adame et al., 2025b). However, some of these constraints have given extremely puzzling results, with a preference for a zero or negative $\sum m_\nu$ found in the data (Craig et al., 2024; Green and Meyers, 2025; Adame et al., 2025b), although this result is somewhat relaxed now given the results of Elbers et al. (2025).

Beyond neutrino masses, cosmological data can also provide some insight into new physics coupling to neutrinos. Indeed, the early Universe, with its high density and pressure, may allow for the production of yet-unknown weakly coupled light particles Baumann et al. (2018), allowing cosmology to probe interactions that cannot be accessed by terrestrial experiments. The presence of such additional light species (including sterile neutrinos) would change N_{eff} in a way that would be potentially detectable in cosmological data. It may also be possible for cosmological data to search for hints of non-standard interactions in neutrinos.

The parameter N_{eff} captures the effects of light relics which contribute to the radiation density Ω_r , including those that free-stream, but also the effects that may be caused by species that are fluid-like (but do not free-stream). As such, in Montefalcone et al. (2025c), an approach is taken to isolate and separately constrain impacts of free-streaming and fluid-like relics. Similarly, the impacts of dark radiation (which may involved fluid-like components, free-streaming components or radiation with varying decoupling or recoupling mechanisms) are studied in Brinckmann et al. (2023); Saravanan et al. (2025). In other works, there have been searches for signatures of non-standard neutrino interactions in the matter power spectrum and CMB power spectra. In particular, Kreisch et al. (2020); Camarena et al. (2023a); Kreisch et al. (2024); Poudou et al. (2025); He et al. (2024); Park et al. (2019); Oldengott et al. (2017); Venzor et al. (2022) have studied the impacts of neutrinos with self-interactions on the CMB power spectra and matter power spectrum, and there have been some features in the data to hint at strong neutrino self-interactions. For example, Poudou et al. (2025) find that DESI BAO data combined with CMB data from the new *Planck* analysis (Tristram et al., 2024; Rosenberg et al., 2022) prefers neutrinos with strong self-interactions to Λ CDM, with the difference between the minimum χ^2 in each case as -4.3 (this comparison involves 8 model parameters + G_{eff} in the strong self-interaction model). Other analyses have found that standard model neutrinos are still preferred by the data despite such hints (Lancaster et al., 2017; Brinckmann et al., 2021; Choudhury et al., 2021; Camarena and Cyr-Racine, 2025; He et al., 2025).

6.1.2 Studying the phase shift of self-interacting neutrinos: motivation

Searching for these signatures in cosmological data is partially motivated by searching for solutions to the Hubble tension and the σ_8 tension (Kreisch et al., 2020; Poudou et al., 2025; Berryman et al., 2023; Berbig et al., 2020), since allowing for neutrinos that self-interact (or otherwise have delayed free-streaming, a consequence of neutrino self-interactions) alters the sound horizon scale measured from the CMB or BAO data. However, the search for new neutrino interactions are also motivated by the need to explain anomalous results in terrestrial neutrino experiments (Crivellin and Mellado, 2024) and even perhaps to explain the apparent preference for negative neutrino mass (Craig et al., 2024;

Das et al., 2025) or understand constraints in the studies in which additional freedom is allowed in neutrino properties (Pal et al., 2025). Self-interactions can also allow for greater flexibility in fits to inflationary models (Choudhury et al., 2022; Barenboim et al., 2019). The self-interactions in these works generally assume a model in which all neutrinos can share a common parameter to describe the strength of self-interactions, $G_{\text{eff}} = |g_{\nu}^2|/m_{\phi}^2$, for a four-fermion contact interaction. The universal coupling is given by g_{ν} and the mediator mass is given by m_{ϕ} . While several terrestrial experiments (Lyu et al., 2021; Blinov et al., 2019) put strong constraints on this kind of interaction, modelling this interaction in cosmological data can capture the general impacts that neutrinos with more complicated interactions may have, thus it is useful to study. For example, Das et al. (2025) study a scenario in which neutrinos convert into dark radiation after BBN but before recombination. This leaves a signal similar to strongly self-interacting neutrinos in cosmological data, without violating constraints from terrestrial experiments.

However, in many cases, the study of neutrino self-interactions through the G_{eff} parameterization leads to bimodal constraints. In various works (see for example Camarena et al., 2023a; Camarena and Cyr-Racine, 2025; Kreisch et al., 2020, 2024; Berryman et al., 2023; Lancaster et al., 2017; Oldengott et al., 2017; Park et al., 2019) one sees two peaks in the posterior for $\log_{10}(G_{\text{eff}})$; one peak shows a preference for very strong self-interactions with $\log_{10}(G_{\text{eff}}) \sim 1.5$. The second peak prefers sufficiently weaker interactions that are not as strongly constrained and the simultaneous fits to other cosmological parameters are somewhat consistent with Λ CDM. The Λ CDM model and the weaker neutrino self-interaction model are not so clearly distinguished. Due to the nature of the impact that neutrino self-interactions have on the broadband-shape of the matter power spectrum, it is difficult to distinguish whether the peak associated with stronger interactions is simply a result of degeneracies with the parameters n_s and A_s (Berryman et al., 2023) or truly a hint for strong-self interactions. However, the impact that neutrinos have on the phase of the BAO and CMB peaks is a more distinct feature that could be studied that is induced by free-streaming relics.

In this chapter, I continue to explore how cosmological data can be used to search for signatures of neutrinos with non-standard interactions, going beyond a simple measurement of N_{eff} . In particular, I explore the phase shift that neutrinos with self-interactions leave in the CMB and BAO oscillations of the matter power spectrum. I study this in an effort to understand what information is available to constrain non-standard interactions in neutrino species, which is expected to alter the phase shift signal compared to that seen for Standard Model neutrinos. In particular, this may be useful given the bimodal constraints on strongly self-interacting neutrinos seen in the literature, which are driven by degeneracies with A_s and n_s (Berryman et al., 2023). As mentioned, the phase shift in the CMB and BAO is uniquely created by free-streaming relics, and thus is not degenerate with A_s or n_s alone. In section 6.2, I further discuss the neutrino self-interaction model, the cosmological impacts of these interactions and show the phase shift signal in the CMB and BAO, motivating the study of this signal. In section 6.3, I revisit the existing templates for the phase shift signal of Standard Model neutrinos that have been constructed in CMB and BAO data. I build upon these templates in section 6.4 to construct a template for the phase shift that captures the impact of neutrinos with self-interactions.

In sections 6.5, 6.6, I explore the ability to constrain G_{eff} in the BAO and CMB, and conclude in section 6.7. Throughout this work I generally keep $\Omega_b = 0.048$, $h = 0.6763$, $\tau_{\text{re}} = 0.066$, $n_s = 0.9667$, $\ln A_s 10^{10} = 3.064$ unless stated otherwise. To extract the phase shift signal, it is necessary to vary N_{eff} from a fiducial value of 3.044, and Ω_m is varied about the fiducial value of 0.31 to keep the matter-radiation equality scale fixed. For the CMB, it is necessary to alter Y_p about a fiducial value of 0.245 to keep the Silk-damping scale fixed.

6.2 Cosmological impacts of neutrino self-interactions

The impact of Universal self-interactions—parameterized by $\log_{10}(G_{\text{eff}})$ —on the matter power spectrum arises from their impact on neutrino free-streaming. In the Standard Model, neutrinos free-stream from very early times, approximately one second after the Big Bang.³ The phase shift that is caused by free-streaming for Standard Model neutrinos is discussed in more detail in the following section, section 6.3.

Strong self-interactions delay this free-streaming, altering the gravitational potential on different length scales. Consequently, strong self-interactions also alter the phase shift signal of free-streaming neutrinos. I describe how this is so more in the next paragraphs and in section 6.4 where I show the template for the phase shift. To first give context to this, in these self-interaction scenarios, the matter power spectrum shows enhanced power on scales that enter the horizon when self-interactions freeze out and the onset of free streaming begins. Conversely, a suppression on smaller scales appears on scales that enter the horizon while neutrinos are still tightly coupled. In this way, a scale-dependent feature is imprinted, producing effects degenerate with A_s (amplitude) and n_s (tilt).

The additional consequence of modifying the epoch of neutrino free-streaming is that this also modifies the characteristic phase shift observed in power spectrum oscillations. For Standard Model neutrinos, the phase shift approaches zero on the largest scales but has an almost constant amplitude on smaller scales; the amplitude is approximately independent of cosmology except for N_{eff} ; for larger N_{eff} there is a larger net amplitude for the phase shift. When free streaming is delayed due to self-interactions, neutrinos affect the photon–baryon plasma only for modes that cross the horizon after these interactions cease, leading to a scale-dependent change to the oscillation phase. The detailed form of this effect has been less studied in the context of self-interactions and is the focus of this chapter. Figure 6.1 and Figure 6.2 illustrates how the power spectrum and cosmic microwave background temperature-temperature power spectrum and its corresponding phase shift effect change when self-interactions among neutrinos are considered.

While the impact of delayed neutrino free-streaming in the context of universal neutrino self-interactions on the phase shift has not been studied extensively, some similar scenarios have been

³For *massive* neutrinos, they transition from relativistic particles to non-relativistic particles as their thermal velocities decay. As such, the free-streaming has the impact of suppressing small-scale power by carrying away energy, while clustering on scales larger than their free-streaming length. However, this is an impact that occurs much later and is essentially unaffected by their decoupling from other particles (or even delayed free-streaming if they remain coupled for longer) unless the epoch of free-streaming is significantly delayed. This affect is captured by changes to $\sum m_\nu$.




Figure 6.1: The impact of G_{eff} (the interaction strength of neutrinos) on the matter power spectrum. This is directly shown in the first panel, with the relative change to a power spectra with $\log_{10}(G_{\text{eff}}) = -6$ (effectively indistinguishable to ΛCDM) shown in the next panel. In the third panel, we show the BAO oscillations computed by dividing out the smoothed broadband power spectrum. We also alter A_s and n_s before implementing the smoothing procedure, which is explained in further detail in the main text. The shift in the phase $\Delta\phi = \Delta k r_s$ as a function of k is computed as the shifts in the peaks and troughs of the power spectrum oscillations relative to a power spectrum with $\log_{10}(G_{\text{eff}}) = -6$; this is shown in the fourth panel. It is possible to also see this signal as a function of ℓ in the CMB power spectra, by plotting $\Delta\ell$ (or $\theta_s\Delta\ell$).




Figure 6.2: The effect of neutrino self-interactions and thus delayed free-streaming on the CMB temperature-temperature (TT) power spectrum; this is directly shown by plotting power spectra with various G_{eff} in the top panel, with the relative difference shown in the second panel. The final panel shows the shift in the peaks and troughs of the CMB oscillations, $\Delta\ell$, which is the phase shift caused by neutrino self-interactions.

explored in some detail. Choi et al. (2018) shows the results of analytical and numerical calculations of a phase shift in the CMB power spectra (TT, EE) when an additional species contributes ΔN_{eff} . In that work, the additional species are given delayed free-streaming compared to Standard Model neutrinos. They show that the phase and amplitude of CMB power spectra (TT, EE) are affected, with a dependence on ΔN_{eff} and the epoch at which free-streaming occurs. The decoupling species induces a scale-dependent change in the phase shift and leads to a drop in the phase signal amplitude (relative to the Standard Model neutrino phase shift) on small scales. These scales correspond to those smaller than the horizon when the species decouple. More recently, Montefalcone et al. (2025b) has studied the phase shift for neutrinos that have self-interactions or dark-matter couplings, causing neutrinos to behave as a fluid prior to decoupling and propagating faster than the sound speed. They find in the case of gradual decoupling of neutrinos from self-interactions or dark-matter couplings over a finite length of time, the phase shift has a roughly constant amplitude dependence, with a small scale-dependence that they choose to not include in their modelling or analysis since the effect is

small. In the analysis presented I include a subtle scale dependence in the phase shift measurements, which may be detectable with highly sensitive future experiments. Additionally, Saravanan et al. (2025) studies in some detail the impact of altering the radiation density and the fraction of it which free-streams or is instead fluid-like; we expect that the impact on the phase shift of increasing the fluid-like radiation density should be similar to increasing the strength of neutrino self-interactions for CMB and BAO data, although their discussion gives a more general description rather than for the case we study here.

I numerically study changes to the phase shift in both the BAO oscillation signal and the CMB power spectra (TT, EE, TE) in the simplistic model for a four-fermion contact interaction. I explore this when N_{eff} or G_{eff} is varied. To reiterate, the goal is to understand whether there is some ability to use a measurement of the phase shift in these spectra to constrain G_{eff} (or any non-standard model particles expected to induce similar imprints on the phase-shift). To do this, I use the code CLASS-PT of Camarena et al. (2023a)⁴ which has implemented a four-fermion contact interaction between neutrinos in the CLASS (Blas et al., 2011) Boltzmann solver, and allows one to compute the matter power spectrum and CMB power spectra with the neutrino self-interaction strength varied via G_{eff} . This is of interest because a detection of the phase shift allows for a detection of the CvB that is more robust to changes in other cosmological parameters; in refs. Baumann et al. (2019); Montefalcone et al. (2025c), a template has been constructed for the phase shift scale dependence due to Standard Model neutrinos that is independent of cosmological parameters; the amplitude of the phase shift (and thus the template) is simply scaled with changes to N_{eff} , allowing for one to infer the presence of the CvB, and constrain N_{eff} .

While we can use the broadband shape of the matter power spectrum or CMB power spectra to study neutrino self-interactions, the impacts of changing the neutrino free-streaming properties are very degenerate with A_s and n_s given the accessible scales in the data. This has led to various works attempting to measure G_{eff} from cosmological data suffering from bimodality in the constraints. However, it may be possible for the phase shift to assist in constraining G_{eff} in a manner that does not suffer from this degeneracy. Unsurprisingly, I show the effect of altering G_{eff} for fixed N_{eff} on the phase $\Delta\ell$ (or Δk) is similar to the effect seen in Choi et al. (2018) on the phase shift when free-streaming is delayed. For neutrinos with more greatly delayed free-streaming (for larger G_{eff}), there is a decreasing amplitude for the phase shift and a turnover scale emerges which may be related to the scale identified in Choi et al. (2018) as the horizon corresponding to the delayed neutrino free-streaming. In section 6.4, I fit a template to capture the phase shift specifically in the case we are varying G_{eff} , building on the template developed for Standard Model neutrinos in Baumann et al. (2016, 2019) in the BAO signal, and by Montefalcone et al. (2025c) in the CMB power spectra.

⁴<https://github.com/davidcato/class-interacting-neutrinos-PT>

6.3 The phase shift of Standard Model neutrinos

Relativistic neutrinos move close to the speed of light c , while the BAO sound speed is $\sim c/\sqrt{3}$. The parameter N_{eff} quantifies the impact of any radiation species with this free-streaming behavior. Particles that propagate faster than the BAO sound speed during BAO formation induce a characteristic phase shift in the power spectra ($\propto \Delta k$ in matter and $\propto \Delta \ell$ in the CMB), producing a robust signature of the CνB (Bashinsky and Seljak, 2004; Baumann et al., 2016; Green and Ridgway, 2020). This shift arises because free-streaming particles alter the gravitational potential, introducing a change on the phase of BAO oscillations. Bashinsky and Seljak (2004) showed that modifying the driving of a harmonic oscillator – an analogy for BAO perturbations – changes both the amplitude and phase of the oscillations. Since neutrinos move faster than the sound speed, they generate metric perturbations beyond the sound horizon, shifting the physical scale of sound waves that freeze at recombination toward larger scales (smaller k and ℓ).

On small scales, the amplitude of the phase shift $\Delta\phi$ is proportional to $\varepsilon_{\nu} = \frac{N_{\text{eff}}}{N_{\text{eff}} + \alpha_{\nu}}$, where $\alpha_{\nu} = \frac{8}{7} \left(\frac{11}{4}\right)^{4/3}$. However, on large scales the amplitude of the phase shift approaches zero. The phase shift was first parameterized and measured in CMB data by Follin et al. (2015), then subsequently in BAO data by Baumann et al. (2019); Wallisch (2019) and more recently in the DESI DR1 BAO data by Whitford et al. (2025) and CMB data by Montefalcone et al. (2025c). Updated constraints from DESI BAOs using the second DESI data release (DR2) can be found in Elbers et al. (2025). The CMB constraints by Montefalcone et al. (2025c) show the presence of a non-zero phase shift at a significance of 14σ , while the updated BAO constraints in Elbers et al. (2025), with the addition of priors from Planck data, find a non-zero detection at a significance of 6.9σ .

6.3.1 Baryon acoustic oscillations

For Standard Model neutrinos, the impact of neutrino free-streaming is to shift the Fourier-space BAO signal toward smaller k . Bashinsky and Seljak (2004) showed that the shift in the phase, $\Delta\phi = \Delta k r_s$, is $\Delta\phi \approx 0.1912\pi\varepsilon_{\nu}$ for large k . An improved calculation is given by Green and Ridgway (2020), however a template was numerically fit to the phase shift by Baumann et al. (2018, 2019), giving

$$\begin{aligned} \Delta\phi &= \beta_{\phi}(N_{\text{eff}})f(k) \\ &= \frac{\varepsilon_{\nu}}{\varepsilon_{\nu,\text{fid}}} \frac{\phi_{\infty}}{1 + (k_*/k)^{\xi}}. \end{aligned} \quad (6.2)$$

where $\beta_{\phi} = \varepsilon_{\nu}/\varepsilon_{\nu,\text{fid}}$, and $\varepsilon_{\nu,\text{fid}}$ is ε_{ν} for some fiducial cosmology. This template provides an effective representation of the phase shift that captures the most important features and is independent of the cosmological model. The amplitude is captured by β_{ϕ} , which is relative to the choice of N_{eff} in a fiducial model (labelled by ‘fid’). The function $f(k)$ captures the scale dependence of the phase shift, with $\phi_{\infty} = 0.227$, $k_* = 0.0324h \text{ Mpc}^{-1}$, and $\xi = 0.872$.

To reveal the phase-shift signature of free-streaming radiation in the BAO signal independently of other cosmological effects, the oscillatory component must be extracted. This is done by smoothing




Figure 6.3: The phase shift function $f(k)$ extracted from the BAO signal. $f(k)$ has been extracted for 120 spectra by computing $\Delta k/r_s$ relative to a spectrum with $N_{\text{eff}} = 3.044$, for a range of values of N_{eff} indicated by the colour bar, with each normalized by $\beta_\phi(N_{\text{eff}})$. The dotted line shows the fitting function given in Baumann et al. (2018, 2019), and the black points show the average value of $f(k)$ from the 120 spectra. The power spectra were computed using CLASS (Blas et al., 2011).⁶

the matter power spectrum $P(k)$ to obtain $P_{\text{sm}}(k)$, and then defining $\mathcal{O}(k) = \frac{P(k)}{P_{\text{sm}}(k)} - 1$. Then, one can repeat this process for a range of power spectra with varying N_{eff} . To ensure weak dependence to other effects, in each case the cold dark matter density needs to be chosen to fix the scale of matter-radiation equality without changing Ω_b , using $a_{\text{eq}} = \frac{\Omega_r}{\Omega_m}$. We can set Ω_{cdm} for a cosmology with $N_{\text{eff, fid}}$, $\Omega_{b, \text{fid}}$, $\Omega_{\text{cdm, fid}}$ as

$$\Omega_{\text{cdm}} = \left(\frac{\alpha_v + N_{\text{eff, fid}}}{\alpha_v + N_{\text{eff}}} \right) \Omega_m - \Omega_b. \quad (6.3)$$

Furthermore, changes to the sound horizon scale need to be accounted for by rescaling k for each of the power spectra by $\frac{r_s}{r_{s, \text{fid}}}$, where r_s can be computed by CLASS for each cosmology. Then, the shift in the peaks and troughs of the BAO oscillations Δk can be compared to a fiducial model, and the average of the shifts in the peaks and troughs as a function of k can be fit to give the function $f(k)$.⁵ The extracted phase shift signal from theoretical power spectra can be seen in Figure 6.3.

6.3.2 Cosmic microwave background

The phase shift seen in the CMB power spectra has a functional form similar to that seen in BAO data. The functional form seen in the temperature (TT), polarization (EE), and cross temperature-polarization (TE) power spectra is well captured by the template fit in Montefalcone et al. (2025c),

$$\begin{aligned} \Delta\ell &= A_v f_v(\ell) \\ &= \frac{(\epsilon_v - \epsilon_{v, N_{\text{eff}}=3.044})}{(\epsilon_{v, N_{\text{eff}}=1.0} - \epsilon_{v, N_{\text{eff}}=3.044})} \frac{\ell_\infty}{1 + (\ell/\ell_*)^\xi}. \end{aligned} \quad (6.4)$$

The phase shift is given by $\Delta\phi = \Delta\ell\theta_s$, where θ_s is the sound horizon at recombination. The amplitude of the phase shift $A_v = (\epsilon_v - \epsilon_{v, N_{\text{eff}}=3.044})/(\epsilon_{v, N_{\text{eff}}=1.0} - \epsilon_{v, N_{\text{eff}}=3.044})$ is normalized with respect to

⁵More detail on this algorithm can be found in Baumann et al. (2018).




Figure 6.4: The scale dependent function $f_v(\ell)$ that captures the phase shift in the CMB power spectra, calculated by averaging the phase from 100 power spectra with varying N_{eff} after dividing by the amplitude A_v , defined in Equation (6.4). The red line shows the template of Montefalcone et al. (2025c), and the data points show the phase shift extracted from each of the spectra, with the shaded regions showing the spread in the averaged data points.

the difference in the amplitude of the shift for $N_{\text{eff}} = 3.044$ relative to $N_{\text{eff}} = 1.0$. The function $f_v(\ell)$ is shown in Figure 6.4, along with the phase shift extracted from the CMB power spectra using an approach similar to the procedure used for the BAO signal; this is also described in detail in Follin et al. (2015); Hou et al. (2013); Montefalcone et al. (2025c). While the choice of normalization is arbitrary, we would like to express the normalization in a manner similar to that for the BAO; we capture the change in the phase shift amplitude with the parameter $\beta_\phi = \frac{\mathcal{E}_v}{\mathcal{E}_{v,\text{fid}}}$, which is unity when the fiducial cosmological model and the true cosmology match. Therefore we can equivalently write Equation (6.4) as

$$\begin{aligned} \Delta\ell &= \frac{\mathcal{E}_{v,N_{\text{eff}}=3.044}}{(\mathcal{E}_{v,N_{\text{eff}}=1.0} - \mathcal{E}_{v,N_{\text{eff}}=3.044})} \frac{(\mathcal{E}_v - \mathcal{E}_{v,N_{\text{eff}}=3.044})}{\mathcal{E}_{v,N_{\text{eff}}=3.044}} f_v(\ell) \\ &= \left(\frac{\mathcal{E}_{v,N_{\text{eff}}=3.044}}{\mathcal{E}_{v,N_{\text{eff}}=1.0} - \mathcal{E}_{v,N_{\text{eff}}=3.044}} \right) [\beta_\phi - 1] f_v(\ell). \end{aligned} \quad (6.5)$$

This parametrizes the phase shift amplitude in terms of β_ϕ , analogous to the BAO phase shift description, and the factor of $\mathcal{E}_{v,N_{\text{eff}}=3.044}/(\mathcal{E}_{v,N_{\text{eff}}=1.0} - \mathcal{E}_{v,N_{\text{eff}}=3.044})$ is simply a constant factor. This choice specifically forces the fiducial cosmology to correspond to $N_{\text{eff}} = 3.044$ in the denominator of β_ϕ .⁷

As for the BAO, in order to isolate the phase shift signature independently of other cosmological parameters, one needs to alter Ω_{cdm} to fix the epoch of matter-radiation equality as N_{eff} is varied, and

⁷One also needs to account for the fact that in this writing of the phase shift $\Delta\ell$ is actually negative if $N_{\text{eff}} > 1.0$ (for the true cosmology), because the constant factor $\mathcal{E}_{v,N_{\text{eff}}=3.044}/(\mathcal{E}_{v,N_{\text{eff}}=1.0} - \mathcal{E}_{v,N_{\text{eff}}=3.044})$ in front of the equation is negative. However, the phase shift $\Delta\phi$ in a cosmology with a positive number of neutrino species, that shifts $\mathcal{C}(\ell) \rightarrow \mathcal{C}(\ell + \Delta\phi)$, relative to a cosmology with $N_{\text{eff}} = 0$, should shift the peaks to smaller values of ℓ , which is only true if $\Delta\phi$ is positive. Therefore, the above equation acquires a negative sign above, unless one directly maps $\mathcal{C}(\ell)$ evaluated at ℓ , to $\ell + \Delta\phi$, rather than mapping ℓ to $\mathcal{C}(\ell + \Delta\phi)$.

the change in scale θ_s (computed by CLASS) is used to rescale the ℓ 's to fix the sound horizon at the epoch of recombination. In the CMB, one must also fix the diffusion damping scale r_D , which can be controlled by Y_p , the helium abundance (mass fraction). I fix the ratio r_D/r_s using the approximate relation (Hou et al., 2013)

$$\frac{r_D}{r_s} = \frac{\theta_D}{\theta_s} \approx \frac{(1 + \rho_\nu/\rho_\gamma)^{0.28}}{\sqrt{1 - Y_p}}. \quad (6.6)$$

The ratio ρ_ν/ρ_γ can be computed as $N_{\text{eff}}a_\nu^{-1}$. The power of 0.28, seen in the numerator, may not be optimal for all cosmological models, but I found this choice worked sufficiently well for the choice of fiducial cosmology. I note the oscillations seen in both Figure 6.3 and Figure 6.4 contain both physical and noise contributions, with the physical components arising since the phase shift template is only a leading order approximation; nonetheless, the templates capture the important features we can reasonably hope to constrain.

6.4 The phase shift with non-standard model neutrino physics

In this section I explore how the phase shift changes in a power spectrum generated for a cosmology with non-standard model neutrinos. As mentioned before, I consider the case of self-interacting neutrinos following a four-fermion interaction with universal coupling, whose effective strength is parameterized by G_{eff} . I apply the same algorithms as discussed above for both the BAO and the CMB to see the phase shift, but now for the case in which G_{eff} is varied between $[-6, 0.5]$. For very small G_{eff} , I expect the power spectrum to match the case of Λ CDM with Standard Model neutrinos. This should recover the phase shift signal seen in the previous sections, with the amplitude modulated by the value of N_{eff} , as demonstrated in Appendix A.7. Like the phase shift introduced by Standard Model neutrinos, I also expect the impact of G_{eff} on the phase shift to be roughly independent of z as the physics that causes the phase shift in the early Universe remains unaffected by the physics of the late Universe.

In practice, I fit multiple templates to the phase shift signal by varying the power spectra with N_{eff} to obtain a function analogous to $f(k)$, for various fixed choices of G_{eff} .⁸ This approach allows one to easily obtain the averaged non-linear dependence of the phase shift for different fixed values of G_{eff} . G_{eff} changes the functional form of the phase shift due to delayed free-streaming. Like the algorithm described previously, I compute the shift Δk for many values of N_{eff} while dividing out the change in amplitude β_ϕ . The main impact of varying N_{eff} is simply to linearly rescale the amplitude of the phase shift. This is expected to be the case regardless of the value of G_{eff} since the neutrino interactions are universal in the model I use; for alternative cases where interactions are not universal, it would be necessary to study the phase shift that arises in these circumstances separately. While the model studied here and variations of self-interaction models are already very strongly constrained by terrestrial experiments (Lyu et al., 2021; Blinov et al., 2019; Belotsky et al., 2001; Berkov et al., 1987,

⁸It may also be possible to alternatively fix N_{eff} , and instead fit a function analogous to $f(k)$ due to changing G_{eff} , and repeat this for different values of N_{eff} . This is illustrated in Figure 6.1. However, I found fitting a template to this to be much more difficult due to the non-linear impact of G_{eff} .

1988), it allows us to test and understand whether there is information available regarding non-standard interactions in the phase shift of the power spectra that could be accessible with future experiments.

6.4.1 Baryon Acoustic Oscillations




Figure 6.5: The phase shift signal for different fixed values of G_{eff} . The red line shows a fit through the extracted phase shift and the blue line shows the template for Standard Model neutrinos (Baumann et al., 2018; Wallisch, 2019). The dashed black lines correspond to the standard model theoretical prediction derived by Bashinsky and Seljak (2004). The numerical calculations are noisy for $-2 < \log_{10}(G_{\text{eff}}) < -1$, likely because there is more power on scales close to the BAO scale for these interaction strengths which introduces noise into the dewiggling process. Additionally, for all plots on scales $> 0.4 h \text{Mpc}^{-1}$ the BAO oscillations are very small, leading to a noisier signal. As such we fit only for $k \leq 0.4 h \text{Mpc}^{-1}$. The power spectra were generated using CLASS-PT.

I apply the same algorithm as discussed in section 6.3.1 for the BAO signal, but now multiple times for G_{eff} varied between $[-6, 0.5]$. As we expect, for very small G_{eff} the power spectrum and phase

shift is effectively the same as for the case of Λ CDM with Standard Model neutrinos. In Figure 6.5 I show the numerically extracted phase shift with varying G_{eff} , and a fit through it. In practice, we fit the phase shift considering the functional form of $f(k)$ for the Standard Model case (Baumann et al., 2019; Wallisch, 2019) modulated by extra variations on amplitude and scale-dependence due to the delayed free streaming. This contrasts with the fits shown in Montefalcone et al. (2025b) for the CMB, in which they fit only the amplitude change and neglect the scale dependent change in the phase shift signal. Carefully studying Figures 2, 6 and 7 in their work shows a similar scale-dependent change in the signal. In line with the scale-dependent damping and shifting ‘peak’ seen in the results of Choi et al. (2018), we find that such changes can be well-represented by a changes in amplitude A and additional exponential damping parameterized by B , both functions of G_{eff} , such we have

$$\Delta\phi = A(G_{\text{eff}})f(k)e^{kr_s B(G_{\text{eff}})}. \quad (6.7)$$

In each case, the parameters A and B have been fit using the package EMCEE (Foreman-Mackey et al., 2013),⁹ and this equation and the figure shows the phase shift amplitude for $N_{\text{eff}} = 3.044$ relative to a spectrum with $N_{\text{eff}} = 0$. I found that using EMCEE gave results fairly consistent with using the SCIPY.OPTIMIZE.MINIMIZE algorithm (Virtanen et al., 2020a,b; Wes McKinney, 2010) but also allowed the ability to estimate uncertainties on A and B . More generally, we could write the shift in the ks in the matter power spectrum oscillations relative to Λ CDM with $N_{\text{eff}} = 3.044$ as

$$\mathcal{O}(k) \rightarrow \mathcal{O}(k + [A(G_{\text{eff}})\beta_\phi(N_{\text{eff}})e^{kB(G_{\text{eff}})r_s} - 1]f(k)/r_s). \quad (6.8)$$

In the limit that $G_{\text{eff}} \rightarrow 0$, we have the standard model phase shift as $A \rightarrow 1$ and $B \rightarrow 0$, giving

$$\mathcal{O}(k) \rightarrow \mathcal{O}(k + [\beta_\phi(N_{\text{eff}}) - 1]f(k)/r_s), \quad (6.9)$$

which in turn reduces to the standard BAO wiggles in the case where N_{eff} is equal to the fiducial value. In Figure 6.6 I show the fit to the parameters A and B as a function of $\log_{10}(G_{\text{eff}})$.¹⁰ The change to the phase shift signal with varying $\log_{10}(G_{\text{eff}})$ is well captured by the parameter A that modulates the amplitude of the signal, and B that modulates an exponential damping. As shown by the uppermost panel of Figure 6.5, the fit appears less accurate for very strong self-interactions. Although improved modeling could be explored in future work, the discrepancy is likely due to numerical noise, since the phase-shift signal is suppressed at strong interaction strengths.

The variation of the amplitude A with G_{eff} is expected to be at least partially degenerate with the effect of N_{eff} , which also modulates the phase-shift amplitude through β_ϕ . This degeneracy will weaken the constraints on G_{eff} . Given that the BAO signal only puts weak constraints on N_{eff} alone from the phase shift at present, this suggests that at present it is unlikely that the BAO will be able to strongly constrain G_{eff} from just the phase shift signal alone (although we quantify this more clearly in Section 6.5). However, it may be possible to combine BAO and CMB information to obtain better

⁹<https://emcee.readthedocs.io/en/stable/#>

¹⁰I employ 6th-order polynomial fits to represent A and B as functions of $\log_{10}(G_{\text{eff}})$. I have found this particular choice to be the most suitable for our Forecast analysis. Other approaches, as a rational construction of these functions, might outperform our choice when applied to real data.




Figure 6.6: The fit to the phase shift amplitude and exponential damping parameters A and B in Equation 6.7 as a function of $\log_{10}(G_{\text{eff}})$. I use a 6th order polynomial to fit a smooth function through the points extracted from the phase shift calculations. The points in the fit for the template as a function of k has been weighted by the uncertainty on the individual fits for each parameter A and B . The error bars have been scaled by the χ^2 goodness-of-fit to each point.

constraints. Additionally, the signal is quite subtle: it does not change quickly for $\log_{10}(G_{\text{eff}}) \geq -0.5$; while for $\log_{10}(G_{\text{eff}}) \leq -3.5$ there is no discernible impact on the phase shift at all. This is also inline with the results presented by Montefalcone et al. (2025b) for the CMB, who find the phase shift approaches a constant in the limit of fluid-like neutrinos (corresponding to stronger interactions). For weaker interaction strengths, neutrino free-streaming is not significantly delayed.

I note that there was difficulty in dewiggling the power spectrum and thus extracting the phase shift signal for $-2 \leq \log_{10}(G_{\text{eff}}) \leq -1$. The power spectrum wiggles and smooth broadband signal could not be robustly separated by simply applying the dewiggling algorithm of Hinton et al. (2016) or Wallisch (2019). I identified this as being caused by the additional power that the neutrino self-interactions leave in the full shape matter power spectrum due to the delayed free-streaming; see the second panel in Figure 6.1. Relative to the Λ CDM power spectrum, there is a ‘bump’ on scales close to the BAO scale for $-2 \leq \log_{10}(G_{\text{eff}}) \leq -1$, which possibly makes it more difficult to fit a smooth polynomial through the spectrum (whether in real space or Fourier space). However, I found the best approach to improve the robustness of these methodologies was to simply vary A_s and n_s in a manner that partially cancels the additional power that enters the spectrum on those scales. This allowed the dewiggling algorithm of Wallisch (2019) to work more robustly. This solution in altering A_s and n_s does not impact the phase shift of the BAO oscillations since A_s and n_s have no impact on the sound horizon scale or positions of the peaks. I used the following equations to vary A_s and n_s ,¹¹

¹¹It is apparent in equation 6.10 that there is a degeneracy between A_s and n_s with G_{eff} ; however, to reiterate this degeneracy only applies to the impact that G_{eff} has on the broadband signal and not the oscillatory signal. This allows us to isolate the phase shift in the oscillatory signal and study the impact that is solely due to G_{eff} , allowing us to study the impact of G_{eff} independent of A_s and n_s .

$$\begin{aligned}\ln(10^{10}A_s) &= -0.28 \log_{10} G_{\text{eff}} + 2.62, \\ n_s &= -0.076 \log_{10} G_{\text{eff}} + 0.835.\end{aligned}\tag{6.10}$$

This relation is derived using the correlation between A_s and n_s with $\log_{10} G_{\text{eff}}$ from a set of MCMC chains from real data in (Camarena et al., 2023a). While A_s and n_s are able to generally capture the impact that G_{eff} has on the broadband shape, in future work it would be worth exploring an approach that instead minimizes the difference between a power spectrum with varying A_s and n_s in a limited range of k with respect to a power spectrum with changes to G_{eff} . This would allow one to explore if a more precise model that goes beyond a simple linear relation. Additionally, one could perform tests to ensure that there is no ‘leakage’ of the oscillatory signal when applying such a relation; however, given the agreement seen between the amplitude and scale dependence of the phase shift signal in the BAO with that in the CMB in the next sections with respect to changes in G_{eff} (which also matches the expectation for such a signal seen in Choi et al., 2018; Montefalcone et al., 2025b), this suggests our algorithm has not introduced a spurious signal into our phase shift template.

6.4.2 Cosmic microwave background

As before for the BAO signal, I explore how the phase shift signal changes in the CMB power spectra for neutrinos with non-standard interactions by varying G_{eff} . Figure 6.7 shows how the phase shift signal is impacted by self-interactions, and a model fit through the modified phase shift signal. The phase shift can be seen very clearly in the CMB data; I refit the parameters A and B as before, keeping the template $f_\nu(\ell)$ from Equation 6.4 and modulating the amplitude via A and an exponential damping through B for each value of $\log_{10}(G_{\text{eff}})$:

$$\Delta\ell = A(G_{\text{eff}})f_\nu(\ell)e^{(B(G_{\text{eff}})/50\theta_s\ell)}\tag{6.11}$$

Similar to our fits for the BAO, we find the parameter B captures a small but significant scale-dependent change to the phase shift amplitude, which is neglected in Montefalcone et al. (2025b) despite being visible in some of their figures. In Figure 6.8 I show the fit to these parameters as a function of G_{eff} . As before, I use EMCEE to fit the parameters. The function $f_\nu(\ell)$ has been defined in Montefalcone et al. (2025c) to be the phase shift for $N_{\text{eff}} = 3.044$ relative to $N_{\text{eff}} = 1.0$. I re-express the shift in the ℓ s more generally in order to be able to show the shift relative to an arbitrary number of neutrinos (N_{eff}) if desired,

$$C_\ell(\ell) \rightarrow C_\ell(\ell + [\beta(N_{\text{eff}})A(G_{\text{eff}})e^{(B(G_{\text{eff}})/50\theta_s\ell)} - 1]C_\nu f_\nu(\ell)).\tag{6.12}$$

Written like this, one can always choose β_ϕ such that $\beta_\phi = 1$ at some fiducial N_{eff} , $\beta_\phi = \frac{\mathcal{E}_\nu}{\mathcal{E}_{\nu,\text{fiducial}}}$. Here C_ν is simply a constant equal to $\frac{-\mathcal{E}_{\nu,N_{\text{eff}}=3.044}}{(\mathcal{E}_{\nu,N_{\text{eff}}=1.0} - \mathcal{E}_{\nu,N_{\text{eff}}=3.044})}$.

Unsurprisingly, the effect of altering G_{eff} for fixed N_{eff} on the phase $\Delta\ell$ (or Δk) is similar to the effect seen in Choi et al. (2018) on the phase shift when free-streaming is delayed. For neutrinos with more greatly delayed free-streaming (for larger G_{eff}), there is a decreasing amplitude for the phase

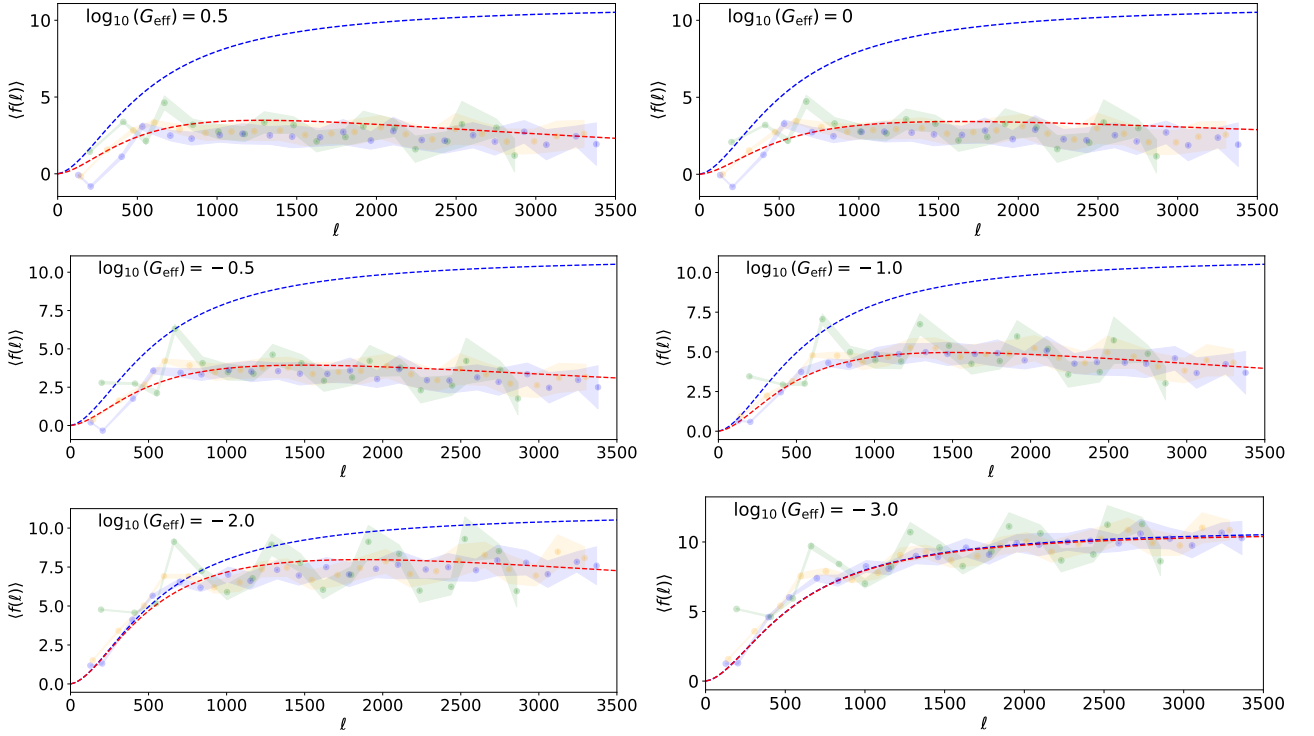


Figure 6.7: The phase shift signal for different fixed values of G_{eff} in the CMB power spectra. The red line shows a fit through the extracted phase shift and the blue line shows the template for Standard Model neutrinos (Montefalcone et al., 2025c). The green, blue and yellow regions show the shift in ℓ for the TT, TE and EE power spectra respectively. The power spectra were generated using CLASS-PT. Correlations between each of the power spectra have been neglected in these fits, but I include cross-correlations later for estimating the constraining power of the spectra for $\log_{10}(G_{\text{eff}})$.




Figure 6.8: The fit to A and B as a function of G_{eff} . We recover $A = 1$ and $B = 0$ at small values of G_{eff} , where we expect the power spectrum and phase shift from this model to agree with the expectation of Λ CDM. In each panel, the blue dashed line shows a polynomial fit through the parameters; for smaller G_{eff} I fit a simple weighted polynomial, as for the BAO, but I only fit a subset of the points (shown by the blue dashed line) and model $A = 1$ and $B = 0$ as constants for larger values of G_{eff} . Some of the noise in the fit to A and B is likely due to the degeneracy between these parameters. For comparison the BAO polynomial fits are also shown by the red dotted lines.

shift and a turnover scale emerges which may be related to the scale identified in Choi et al. (2018) as the horizon corresponding to the delayed neutrino free-streaming.

6.5 Can we constrain G_{eff} from the phase of the BAO?

Having studied the impact of neutrino self-interactions on the phase shift, I now examine whether its detection can be used to constrain non-standard neutrino interactions. Unlike constraints on G_{eff} that rely on the direct study of cosmological observables, such as the CMB and matter power spectra, analyzing the distinct imprint of phase shift has the advantage of being less degenerate with other cosmological parameters, e.g., A_s and n_s . However, we *can* expect it to be highly degenerate with the impact of N_{eff} on the amplitude of the phase, or any other parameter that alters r_s . Despite this limitation, the non-linear dependence of G_{eff} on k may still provide additional constraining power.

A typical BAO analysis involves comparing the BAO oscillations $\mathcal{O}(k)$ to a template for $\mathcal{O}(k)$ that is computed at some fixed ‘template’ cosmology, and then finding the best fit to the ‘distortion’ parameters α_{\parallel} and α_{\perp} , which shift the positions of the BAO oscillations in the template to match them to the data. These parameters ultimately capture two impacts; the first is the Alcock-Paczynski (AP) effect (Alcock and Paczyński, 1979), which leads to the observed distances to galaxies (and thus the BAO feature) to differ from the true distances when one assumes a fiducial cosmology to transform galaxy redshifts into physical distances. The second is the impact of the choice of cosmology in the BAO template, which leads the physical scale r_s to differ from that of the true cosmology. As such, α_{\parallel} and α_{\perp} allow one to determine how much the fiducial cosmology and template cosmology differ from the true cosmology, and constrain parameters that alter the expansion rate.¹² The growth rate of structure multiplied by σ_8 , $f\sigma_8$ and the galaxy bias b also enter as free parameters in the analysis; these parameters change the impact of redshift space distortions through the Kaiser effect (Kaiser, 1987). Including the parameter β_{ϕ} extends upon a typical BAO analysis, but allows one to capture the impact of free-streaming neutrinos on the BAO phase (see more details in Baumann et al. (2019); Whitford et al. (2025)). Here, we go a step further by including the impact of G_{eff} on the BAO phase via the parameters $A(G_{\text{eff}})$ and $B(G_{\text{eff}})$.

¹²These parameters enter by shifting $\mathcal{O}(k) \rightarrow \mathcal{O}(k')$; see more details in Whitford et al. (2025) for the equations relating the BAO distortions parameters to k, k' , and for the equations that define $\alpha_{\parallel}, \alpha_{\perp}$.

6.5.1 Profile of χ^2

I start with considering the BAO oscillations measured from a cosmic-variance (or volume) limited survey, in order to determine if any information is available in the best-case scenario. I first consider what ability there is to constrain G_{eff} alone from the BAO phase, which is illustrated by Figure 6.9a. This illustrates the constraining power for a volume-limited survey that is able to measure the BAO feature in redshift bins of width $\Delta z = 0.1$ spanning $z = 0 - 2$, with a sky area of $\sim 30,000$ square degrees. I consider the constraining power in a survey with $k_{\text{min}}, k_{\text{max}} = 0.001h, 0.5h \text{ Mpc}^{-1}$ with bins of width $\Delta k = 0.0025h \text{ Mpc}^{-1}$. In this first simple analysis, I fix G_{eff} and allow the usual BAO parameters in each redshift bin to vary ($f\sigma_8, b\sigma_8, \alpha_{\parallel}, \alpha_{\perp}$). In section 6.5.2, I obtain an alternative calculation of the significance of a measurement compared to ΛCDM . For BAO data with a true cosmology corresponding to some choice of G_{eff} , I effectively plot the ‘profile likelihood’ — the difference between the χ^2 goodness-of-fit of the data assuming a ΛCDM model compared to models with different G_{eff} . The $\Delta\chi^2$ is essentially the difference between a null hypothesis as ΛCDM and a model with varying G_{eff} , $\Delta\chi^2 = \chi^2_{\Lambda\text{CDM}} - \chi^2$, allowing us to determine a confidence interval for a measurement. In order to compute each $\Delta\chi^2$, I find the optimal $\{f\sigma_8, b\sigma_8, \alpha_{\parallel}, \alpha_{\perp}\}$, in each redshift bin for our ‘data’ when one assumes ΛCDM or various models with varying G_{eff} .




Figure 6.9: a) The difference between the χ^2 goodness-of-fit with ΛCDM and the χ^2 goodness-of-fit for a model with neutrino self-interactions parameterized by G_{eff} , as a function of G_{eff} . This shows the profile for BAO data with a fiducial model that is indicated by the legend. It is the best case result, for a volume-limited survey. b) A plot of the BAO signal for a cosmology in which $\log_{10}(G_{\text{eff}}) = -3$, after subtracting the BAO signal from cosmologies with varying phase shifts due to variations in G_{eff} or β_ϕ . The grey shaded region shows the theoretical error bars for a volume-limited survey, in which the volume covers the range $z = 0 - 2$, and the bins are of width $\Delta k = 0.0025h \text{ Mpc}^{-1}$.

The minimum χ^2 is obtained with a simple optimization with respect to the ‘data’, using the scipy algorithm `SCIPY.OPTIMIZE.MINIMIZE`. I model the variance of the power spectrum at each z using the result of Tegmark (1997) as $\frac{2}{N_{\text{modes}}}(P(k, \mu, z))^2$, where N_{modes} is the number of modes that can be observed in a particular volume, which depends on \bar{n} , the number density of galaxies. For a volume-limited survey, the terms involving the number densities vanish. In this analysis, we expect to

see a peak at the model G_{eff} that matches the truth most closely. A sharper peak indicates better ability to constrain G_{eff} . For our calculations I use the template Equation 6.7 to model the impact of altering G_{eff} on the BAO signal, explicitly ensuring that only information available in the phase shift about G_{eff} is included in the calculation. This is also demonstrated by Figure 6.9b which shows the difference between the BAO spectrum when the phase is altered due to variations in G_{eff} (and for comparison β_ϕ), with the shaded region showing error bars on the BAO power spectrum.

As shown in Figure 6.9a, the results suggest that strong interactions (SIv, $\log_{10}(G_{\text{eff}}) > -2$) are able to be constrained by the cosmic-variance limited BAO data—we will explore what significance these results correspond here and also in the next section. In such cases, the $\Delta\chi^2$ peaks about the value of G_{eff} that matches the fiducial model employed to generate ‘data’. Moderate interactions (MIv, $\log_{10}(G_{\text{eff}}) \leq -2$) are difficult to distinguish from the case of no neutrino self-interactions, given the nearly flat slope of $\Delta\chi^2$. Nevertheless, a preference against strong interactions remains possible, due to the sharp drop in $\Delta\chi^2$ at larger G_{eff} . These results align with expectations: the phase-shift signal is more pronounced for strong interactions, while very weak interactions have little impact.

To assess the statistical significance of our results, I compute the probability associated with $\Delta\chi^2$ under the null hypothesis. Assuming that $\Delta\chi^2$ follows a χ^2 distribution with 80 degrees of freedom (20 redshift bins \times 4 parameters per bin), the probability p of obtaining a larger $\Delta\chi^2$ value under the null hypothesis (Λ CDM) is $p = 1 - \text{CDF}_{k=1}(\Delta\chi^2)$, where CDF denotes the cumulative distribution function of the χ^2 distribution. The corresponding significance in units of standard deviations, σ , can then be obtained by applying the inverse normal CDF at $1 - p$ (for a one-tailed test).

I find that a cosmological scenario with strongly interacting neutrinos (dark blue line in Figure 6.9a) following $\log_{10}(G_{\text{eff}}) = 0.0$ can be detected via phase shift analysis with a significance of $\sim 4\sigma$. However, if the self-interaction strength of the data is slightly weaker with $\log_{10}(G_{\text{eff}}) = -1.5$ (green line), the Λ CDM and self-interacting neutrino models become essentially indistinguishable. The same is true for moderate interacting scenarios with ($\log_{10}(G_{\text{eff}}) = -3$, cyan line). Nonetheless, in the case of a Universe with $\log_{10}(G_{\text{eff}}) = -3$ or standard model neutrinos, an analysis of the phase shift allows to rule out much stronger interactions, $\log_{10}(G_{\text{eff}}) = 0.0$, at a significance also of $\sim 3.8\sigma$. Overall, this suggests that we cannot strongly constrain any interactions with the BAO phase shift alone unless they are very strong interactions, or we are trying to show that strong interactions are not preferred by the data. In this case, the BAO signal can be useful. However, we may at least be able to obtain information from the BAO that could complement CMB constraints, which we expect to have much better constraining power; this will be explored in section 6.6. I will also explore how our ignorance about β_ϕ will impact these constraints in the next section.

6.5.2 Fisher matrix forecasts

To understand the constraining power of the BAO phase shift when G_{eff} (and β_ϕ) are free parameters, I study fisher matrix forecasts (see some more information on the Fisher information in Appendix A.6) for a volume-limited BAO survey. For a galaxy survey our observable is the Fourier-space galaxy

density field $\delta(k)$ which has a mean $\mu = \bar{\delta}(k) = 0$. We can model the covariance Σ for this as $\Sigma = P(k, z, \mu)$, where the shot noise is ~ 0 for a volume-limited survey.¹³ $P(k)$ is the matter power spectrum in redshift space. The fisher information for our observable, which is a multivariate Gaussian distribution of our parameter vector \mathbf{x}_i (containing our cosmological parameters of interest) can be written as

$$\begin{aligned} F_{ij} &= \frac{d\mu}{d\theta_i} \Sigma^{-1} \frac{d\mu}{d\theta_j} + \frac{1}{2} \text{tr} \left(\Sigma^{-1} \frac{d\Sigma}{d\theta_i} \Sigma^{-1} \frac{d\Sigma}{d\theta_j} \right) \\ &= \frac{1}{2} \text{tr} \left(\Sigma^{-1} \frac{d\Sigma}{d\theta_i} \Sigma^{-1} \frac{d\Sigma}{d\theta_j} \right). \end{aligned} \quad (6.13)$$

The first term is zero for $\mu = \bar{\delta}$. Fisher information is additive, and thus one can integrate the information over all modes k , redshifts z and angles μ observed in the survey to obtain the information over a 3D volume. Thus this equation becomes a triple integral to capture the information over all angles and accessible scales,

$$F_{ij} = \int_0^{r_{\text{max}}(z)} d^3r \int_{k_{\text{min}}}^{k_{\text{max}}} d^3k \int_0^1 d\mu \frac{1}{2} \text{tr} \left(\Sigma^{-1} \frac{d\Sigma}{d\theta_i} \Sigma^{-1} \frac{d\Sigma}{d\theta_j} \right). \quad (6.14)$$

To compute the Fisher information, I modified the forecasting code GOFISH (which produces the power spectrum with CAMB, Lewis and Challinor, 2011) to include the phase shift amplitude for standard model neutrinos, β_ϕ , and $\log_{10}(G_{\text{eff}})$.¹⁴ The original code was developed to produce forecasts for the DESI survey, so more details about the modelling in redshift-space can be found in Adame et al. (2024). The Cramér-Rao bound states that the minimal (statistical) variance of a parameter given an observable is bound by the inverse of the fisher information (more details can be seen in Appendix A.6). This allows one to determine a forecast for a parameter of interest from a galaxy survey with a given signal-to-noise ratio.

I assume a volume-limited survey as previously described in section 6.5.1, with a sky area of 30,000 square degrees, with 20 redshift bins from $z = 0$ to $z = 2$. I allow $k_{\text{min}}, k_{\text{max}} = 0.001h, 0.5h \text{ Mpc}^{-1}$. For our Fisher forecasts I set the galaxy bias $b_g = 1$ at $z = 0$ and allow it to vary proportionally to $\frac{\sigma_8(0)}{\sigma_8(z)}$ with z . For all forecasts, I allow the following to be free parameters: $b_g(z)_i, \sigma_8(z)_i$ for each redshift bin i , $f\sigma_8(z)_i$ where f is the growth rate of structure, and α_{\parallel} and α_{\perp} (the anisotropic BAO distortion parameters). I compute the constraining power on $\log_{10}(G_{\text{eff}})$, for different choices of $\log_{10}(G_{\text{eff}})$. I also consider the case that β_ϕ (which changes due to N_{eff} and increases the amplitude of the phase shift signal) is allowed to vary freely. Results are shown in Table 6.1. For a more realistic scenario, I also consider the ability of DESI to constrain neutrino self-interactions from the BAO signal alone in the full five-year data, which is shown in Table 6.2. For very large error bars (≥ 15) we simply state that G_{eff} is not constrained in these tables. For our DESI forecasts, I use a sky area of 14,000 square degrees and set $k_{\text{min}}, k_{\text{max}} = 0.001h, 0.3h \text{ Mpc}^{-1}$. For the number densities of galaxies in BGS, ELGs, LRGs and QSO with z and their bias, I follow the values provided in Adame et al. (2024). For multitracer galaxy surveys, the Fisher Information can be computed by simply extending

¹³ Σ and μ take a matrix form and vector form when computing the Fisher information for multiple tracers.

¹⁴The modified code can be found at <https://github.com/abbew25/GoFish>.

μ to a vector $\boldsymbol{\mu} = [\delta_0, \delta_1, \dots, \delta_N]$ for N tracers. Σ is extended to a $N \times N$ covariance matrix $\boldsymbol{\Sigma}_{ij}$ in which the diagonal terms give the variance on each tracer's density field and the covariance between different tracers is given by the off-diagonals. I treat the constraints on $b_g(z)_i \sigma_8(z)_i$, $f \sigma_8(z)_i$, $\alpha_{\parallel}(z)$ and $\alpha_{\perp}(z)$ independently for each tracer in each redshift bin, which allows for stronger constraints by cross-correlating information between different tracers.

Table 6.1: Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a volume-limited survey. Constraints are shown when one fixes $\beta_{\phi} = 1$ (for $N_{\text{eff}} = 3.044$) or when it is allowed to vary (effectively allowing N_{eff} to change the BAO phase amplitude freely). In brackets, I show the number of σ each constraint gives for $\log_{10}(G_{\text{eff}})$ relative to $\log_{10}(G_{\text{eff}}) = -3$ (I choose this value since it cannot realistically be distinguished from smaller choices of G_{eff}).

Model	$\sigma(\log_{10}(G_{\text{eff}})) _{\beta_{\phi}=1}$	$\sigma(\log_{10}(G_{\text{eff}}))_{\beta_{\phi}\text{ free}}$
SIV, $\log_{10}(G_{\text{eff}}) = 0.25$	± 0.86 (3.96)	± 6.59 (0.0)
SIV, $\log_{10}(G_{\text{eff}}) = 0$	± 0.86 (3.49)	± 13.47 (0.0)
SIV, $\log_{10}(G_{\text{eff}}) = -1.5$	± 1.62 (0.93)	± 8.82 (0.0)
MIv, $\log_{10}(G_{\text{eff}}) = -3$	unconstrained	unconstrained

Table 6.2: Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for the DESI 5-year dataset. The columns are the same as those in Table 6.1.

Model	$\sigma(\log_{10}(G_{\text{eff}})) _{\beta_{\phi}=1}$	$\sigma(\log_{10}(G_{\text{eff}}))_{\beta_{\phi}\text{ free}}$
SIV, $\log_{10}(G_{\text{eff}}) = 0.25$	± 2.07 (1.57)	± 13.76 (0.0)
SIV, $\log_{10}(G_{\text{eff}}) = 0$	± 2.19 (1.37)	unconstrained
SIV, $\log_{10}(G_{\text{eff}}) = -1.5$	± 4.24 (0.0)	12.3 (0.0)

The results in Table 6.1 are consistent with our results from the χ^2 profile analysis; it is possible to obtain constraints on G_{eff} , but only for very strong interactions, and in the case one can assume that N_{eff} is known well enough to fix $\beta_{\phi} = 1$. Interestingly, the detection significance of $\sim 3.5\sigma$ for the case of a model with strong interactions ($\log_{10}(G_{\text{eff}}) = 0$) is similar to what was obtained from the χ^2 profile analysis, which found a detection significance of $\sim 4\sigma$. For weaker interactions, or when we allow $\beta_{\phi}(N_{\text{eff}})$ to vary, the constraining power decreases. This is largely because the derivative $\frac{d\mathcal{O}(k)}{dG_{\text{eff}}}$ is extremely small since the phase shift does not change much for weak self-interactions (see Figure 6.6; the Fisher information directly depends on this), and due to the fact that the amplitude modulations by $A(G_{\text{eff}})$ is degenerate with β_{ϕ} . Moreover, the modulation of the phase-shift amplitude by G_{eff} is less significant than the effect of varying N_{eff} and thus β_{ϕ} . However, larger values of N_{eff} imply larger β_{ϕ} and thus a greater amplitude of the phase shift signal, in which case the amplitude modulation by G_{eff} may become more important, restoring some constraining power on G_{eff} .

While these results suggest that we are able to gain little information about G_{eff} at all if only weak interactions occur, the results from Figure 6.9a suggest there is sufficient information to conduct a model comparison with real data, providing a potential avenue to rule out strong interactions (or more extreme scenarios). It is possible to estimate the significance level at which stronger interactions can be disfavoured using the Fisher information \mathbf{F} ; one can approximately compute a $\Delta\chi^2$ between two models as $\Delta\chi^2 \approx \Delta\boldsymbol{\theta}^T \mathbf{F} \Delta\boldsymbol{\theta}$. If one takes the Fisher information from our constraints for a model

with $\log_{10}(G_{\text{eff}}) = -3$, and compares this to a model with $\log_{10}(G_{\text{eff}}) = 0$, this gives $\Delta\chi^2 \approx 6.75$, corresponding to a preference of the weaker interaction model over the other at the level of $\sim 0.8\sigma$ (here the number of degrees of freedom is just one for our single parameter that changes between the two models, G_{eff} , leaving all other parameters fixed). Compared to our profile analysis, this is smaller (I found earlier a significance of $\sim 3.8\sigma$ compared to ΛCDM), although this calculation does not take into account variation of other parameters.

The results of table 6.2, that forecast for upcoming data from DESI, find the constraints are (as one can expect) weaker than those shown for a volume-limited survey, but there is still some ability to constrain G_{eff} from the phase shift for very strong interactions if we assume a fixed value of N_{eff} ; otherwise, as for the volume-limited case, the constraints are not as useful from BAO data alone. In Section 6.5, I will explore whether there are improved constraints from upcoming BAO data from a survey such as DESI, if it is combined with the constraining power of a CMB experiment.

6.6 Can we constrain G_{eff} from the phase of the CMB power spectra?

6.6.1 Profile of χ^2

As for the BAO, I conduct a calculation of an effective ‘profile likelihood’ of the χ^2 goodness-of-fit to the data, by computing $\Delta\chi^2$ between the model of the null hypothesis (ΛCDM) and models with varying G_{eff} . As before for the BAO, the $\Delta\chi^2$ involves varying a number of parameters that are standard for a CMB analysis while leaving G_{eff} fixed. I vary G_{eff} across simulated datasets generated with different true values of G_{eff} . In each case, we expect the likelihood to peak near the input value. A sharper peak indicates a stronger detection. For simplicity, in this section I just consider the information in the CMB temperature-temperature (TT) power spectrum, and consider a cosmic-variance limited survey (instrument noise is neglected in order to understand the constraining power in the best case scenario) with the same sky area of ~ 30000 square degrees. I consider ℓ in the range $\ell = [2, 2500]$, and use the raw CMB power spectra. As before, a Fisher matrix analysis of the ability to constrain will also be given in the next section, which will include information from the polarization and temperature-polarization cross power spectra.

To calculate the χ^2 I compare the ‘observed’ C_ℓ s to each model; I assume the variance on C_ℓ as $\frac{2}{(2\ell+1)f}C_\ell^2$, where f is the fraction of sky coverage. The minimum χ^2 is found with respect to the ‘data’ also using `SCIPY.OPTIMIZE.MINIMIZE`. For our analysis, I allow the following cosmological parameters to vary: $\{h, \omega_b = \Omega_b h^2, \omega_c = \Omega_c h^2, \ln(10^{10} A_s), n_s, \tau_{\text{re}}\}$. Since generating the power spectrum is slow in our minimization algorithm, I use the `COSMOPOWER` emulator for the CMB TT power spectrum (Spurio Mancini et al., 2021) (and as for the BAO, I use our template to include the impact of G_{eff} on the phase). Figure 6.10a shows the profile likelihood, and Figure 6.10b shows a comparison of how the phase shift changes between different models with varying G_{eff} and β_ϕ . The grey shaded region shows the uncertainty on a CV limited CMB experiment TT power spectrum.




Figure 6.10: a) The difference between the χ^2 goodness-of-fit with ΛCDM and the χ^2 goodness-of-fit for a model with neutrino self-interactions parameterized by G_{eff} . As in Figure 6.9a, this shows the profile (but with CMB data) for a fiducial model that is indicated by the legend. This is the best case result for a cosmic-variance limited survey. b) A plot of the CMB TT spectrum in a cosmology for which $\log_{10}(G_{\text{eff}}) = -3$, after subtracting the CMB signal from cosmologies with varying phase shifts due to variation in G_{eff} or β_{ϕ} . The grey shaded region shows the theoretical error bars for a cosmic-variance limited CMB experiment, in which the sky area covers $\sim 30,000$ square degrees.

Studying Figure 6.10a, we can see that the $\Delta\chi^2$ visually indicates an ability to distinguish ΛCDM and models with neutrino self-interactions in the CMB, with similar trends to that seen for the BAO. The $\Delta\chi^2$ reaches a maximum for the most strongly interacting model (shown as the dark blue line in the figure for the case the ‘data’ corresponds to a model with $\log_{10}(G_{\text{eff}}) = 0.0$). The green line that corresponds to a model with $\log_{10}(G_{\text{eff}}) = -1.5$ shows the expected peak (albeit a small peak) about -1.5. For models with weaker interactions, there is poorer ability to differentiate from ΛCDM , but there is some ability to distinguish these models from stronger interactions. Following the same procedure as before to interpret the $\Delta\chi^2$ as a detection significance (now from a χ^2 distribution with $k = 6$ degrees of freedom), I find for the case the data has a cosmology corresponding to $\log_{10}(G_{\text{eff}}) = 0.0$, the detection significance is $\sim 6.2\sigma$ with respect to ΛCDM . Similar to the case shown for the BAO, the model with $\log_{10}(G_{\text{eff}}) = -1.5$ or smaller values cannot be distinguished from ΛCDM based on this analysis. However, taking the model in which $\log_{10}(G_{\text{eff}}) = -3$, we can show that $\log_{10}(G_{\text{eff}}) = -3$ is preferred over stronger interactions, corresponding to $\log_{10}(G_{\text{eff}}) = 0.0$ at a significance of $\sim 6.6\sigma$. Overall, this result is similar to what we find in the BAO profile analysis, except that the detection significance is greater for very strong interactions.

6.6.2 Fisher matrix forecasts

In the Fisher matrix analysis, I include not only the TT spectrum, but also the EE and TE spectra to determine the full constraining power of the CMB in different cases. I also allow G_{eff} to be a free parameter (and in separate cases β_{ϕ}). For a CMB experiment, our observable is the Fourier-space

temperature fluctuation field across different angles, $\frac{\Delta T(\theta)}{\bar{T}}$; the mean vanishes, like in the case for the matter field δ . However, here we use the covariance matrix of the CMB power spectrum rather than the map itself. This depends on the CMB auto-power spectrum of temperature fluctuations, C_ℓ^{TT} , giving

$$\Sigma_{TT} = \frac{2}{(2\ell+1)f} (C_\ell^{TT} + N_\ell^{TT})^2. \quad (6.15)$$

N_ℓ^{TT} is the noise on the CMB power spectrum that comes from the instruments. Obtaining the full fisher information involves summing over all the accessible angular modes ℓ and power spectra $\{a, b\}$ as

$$F_{ij} = \sum_{a,b} \sum_{\ell_{\min}}^{\ell_{\max}} \left(\frac{dC_\ell^a}{d\theta_i} \Sigma_{a,b}^{-1} \frac{dC_\ell^b}{d\theta_j} \right). \quad (6.16)$$

Similar to the case of computing the Fisher information for multiple galaxy survey tracers, the analysis can be extended to include information from the polarization power spectrum and temperature-polarization cross-power spectrum, thus we have to consider the power spectra for $\{TT, TE, EE\}$. In this case the covariance matrix is given by

$$\Sigma = \frac{2}{(2\ell+1)f_{\text{sky}}} \times \begin{bmatrix} (C_\ell^{TT} + N_\ell^{TT})^2 & C_\ell^{TE} C_\ell^{TE} & (C_\ell^{TT} + N_\ell^{TT}) C_\ell^{TE} \\ C_\ell^{TE} C_\ell^{TE} & (C_\ell^{EE} + N_\ell^{EE})^2 & (C_\ell^{EE} + N_\ell^{EE}) C_\ell^{TE} \\ (C_\ell^{TT} + N_\ell^{TT}) C_\ell^{TE} & (C_\ell^{EE} + N_\ell^{EE}) C_\ell^{TE} & \frac{1}{2} (C_\ell^{TE} C_\ell^{TE} + (C_\ell^{TT} + N_\ell^{TT})(C_\ell^{EE} + N_\ell^{EE})) \end{bmatrix}. \quad (6.17)$$

We assume instrument noise isn't correlated between the spectra, so $N_\ell^{TE} = 0$. As in the previous section, I produce forecasts for a survey with a sky area of $\sim 30,000$ square degrees, which is cosmic-variance limited, giving $N_\ell^{TT} = N_\ell^{EE} = 0$. I use the range $\ell = [2, 2500]$. Aside from G_{eff} I include the following free-parameters in the forecast; $100\theta_s$, $100\omega_b = 100\Omega_b h^2$, $\omega_c = \Omega_c h^2$, $\ln(10^{10} A_s)$, n_s , τ_{re} . I also include $\beta_\phi(N_{\text{eff}})$ in separate forecasts. As for the BAO, I explicitly include information about G_{eff} from the phase shift only by using the template to change the CMB phase shift due to variations in G_{eff} and β_ϕ when these parameters are varied. In order to produce the forecasts, I modified the GOFISH code to produce forecasts for CMB data, which is able to use CAMB or CLASS.¹⁵ For our forecasts, I neglect CMB lensing and use the raw CMB power spectra. In Table 6.3, I show forecasts for a cosmic-variance limited experiment. We are able to obtain forecasts for more weakly interacting models compared to the case for the BAO; the CMB is much more constraining. As is the case for the BAO, it is easier to constrain the interactions from the phase shift when they are stronger, as $\frac{dC(\ell)}{dG_{\text{eff}}}$ is larger.

These results suggest that upcoming CMB experiments will be able to use the phase shift to constrain G_{eff} , even if the constraints will be weaker given realistic noise properties for upcoming surveys; for a cosmic-variance limited survey, the CMB will be able to achieve extremely strong constraints ($> 5\sigma$) on neutrino self-interactions if they exist. For very weak interactions or a fit to cosmology that is consistent with Λ CDM, it should be possible to disfavour stronger interactions using

¹⁵<https://github.com/abbew25/CMB-fish>

Table 6.3: Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a cosmic-variance limited CMB measurement of the TT, EE and TE power spectra. The columns are the same as those in Table 6.1. The detection significance is given in σ relative to a model with $\log_{10}(G_{\text{eff}}) = -3.5$, in which the phase shift is essentially indistinguishable from ΛCDM .

Model	$\sigma(\log_{10}(G_{\text{eff}})) _{\beta_\phi=1}$	$\sigma(\log_{10}(G_{\text{eff}}))_{\beta_\phi\text{ free}}$
SIv, $\log_{10}(G_{\text{eff}}) = 0$	$\pm 0.02(> 5)$	$\pm 0.06(> 5)$
SIv, $\log_{10}(G_{\text{eff}}) = -1.5$	$\pm 0.04(> 5)$	$\pm 0.11(> 5)$
MIv, $\log_{10}(G_{\text{eff}}) = -3$	$\pm 0.27(1.85)$	$\pm 0.49(1.02)$

just the CMB. In Table 6.4 I show forecasts for a *Planck*-like experiment, using the noise properties assumed in Font-Ribera et al. (2014). I assume $f = 0.7$, $\ell = [2, 2500]$ for the TT spectrum and $[2, 2000]$ for the EE and TE spectrum, and again I use the unlensed spectra.

Table 6.4: Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a *Planck*-like CMB measurement of the TT, EE and TE power spectra. The columns are the same as those in Table 6.1. The detection significance is given in σ relative to a model with $\log_{10}(G_{\text{eff}}) = -3.5$, in which the phase shift is essentially indistinguishable from ΛCDM .

Model	$\sigma(\log_{10}(G_{\text{eff}})) _{\beta_\phi=1}$	$\sigma(\log_{10}(G_{\text{eff}}))_{\beta_\phi\text{ free}}$
SIv, $\log_{10}(G_{\text{eff}}) = 0$	$\pm 0.15(> 5)$	$\pm 0.95(3.68)$
SIv, $\log_{10}(G_{\text{eff}}) = -1.5$	$\pm 0.37(> 5)$	$\pm 0.77(2.59)$
MIv, $\log_{10}(G_{\text{eff}}) = -3.0$	$\pm 1.07(0.46)$	$\pm 3.36(0.0)$

The constraints here for G_{eff} in the CMB are comparable to those for neutrino self-interactions in other works (Kreisch et al., 2020; Poudou et al., 2025; Camarena and Cyr-Racine, 2025) without suffering from the degeneracy with A_s and n_s that requires the analysis to be split into two regions for SI and MI modes. Our analysis shows that with CMB data, the phase shift information may be a promising approach to measure or disfavour self-interactions that delay neutrino free-streaming. Similar to the case for the CMB, self-interactions will be poorly constraining in the case that $\log_{10}(G_{\text{eff}}) \leq -3$ due to the fact the derivative $\frac{dC(\ell)}{dG_{\text{eff}}}$ approaches zero very quickly. However, it will likely be possible that model comparisons can help to disfavour strong interactions models with real data.

6.6.3 Combined BAO and CMB forecasts

Finally, I consider the combined constraining power of BAO data from a galaxy survey with CMB data. This is mainly to test whether the phase shift information in the BAO is able to provide any useful additional constraining power for G_{eff} from the phase shift information available in a CMB experiment alone. In Table 6.5, I show the forecasts for constraints on G_{eff} from the combined information in DESI YR5 BAO data with data from a *Planck*-like CMB experiment. To compute the combined constraints, I take the Fisher information matrices from the forecasts for the BAO experiment from DESI, and use a Jacobian transformation to convert the Fisher matrix for the parameters $\{\alpha_{i,\parallel}(z), \alpha_{i,\perp}(z), f\sigma_{i,8}(z), \beta_\phi, \log_{10}(G_{\text{eff}})\}$ (where i is an index running over different redshift bins and tracers) to a new Fisher matrix with $\{100\theta_s, \Omega_b h^2, \Omega_c h^2, A_s, n_s, \tau_{\text{re}}, \beta_\phi, \log_{10}(G_{\text{eff}})\}$. Fisher

information is additive, so in this form the Fisher information from the DESI data forecasts and the *Planck*-like experiment can simply be added together.

Table 6.5: Fisher forecasts for the ability to constrain G_{eff} , for a few different models, for a *Planck*-like CMB measurement of the TT, EE and TE power spectra, combined with a measurement of data from DESI YR5. The columns are the same as those in Table 6.1; the brackets show the relative improvement on the constraints from a CMB experiment alone.

Model	$\sigma(\log_{10}(G_{\text{eff}})) _{\beta_\phi=1}$	$\sigma(\log_{10}(G_{\text{eff}}))_{\beta_\phi \text{ free}}$
SIv, $\log_{10}(G_{\text{eff}}) = 0$	$\pm 0.15(1.0)$	$\pm 0.93(1.02)$
SIv, $\log_{10}(G_{\text{eff}}) = -1.5$	$\pm 0.35(1.06)$	$\pm 0.69(1.11)$
MIv, $\log_{10}(G_{\text{eff}}) = -3.0$	$\pm 1.04(1.03)$	$\pm 2.29(1.47)$

The results in Table 6.5 show that the BAO information is indeed able to improve the constraints by a factor of $\sim 1 - 1.47$ (up to a 47% improvement) depending on the strength of the self-interactions. Since the CMB data is very constraining on its own for stronger self-interactions, the combined Fisher information from both probes together has the least improvement in this case. For weaker constraints from the CMB, the BAO information is more useful to help improve the constraints. The results here suggest that even though the BAO information alone may be less useful than CMB data for constraining self-interactions, BAO data can add useful information to the constraints from the phase shift that are worthwhile, especially for distinguishing weaker interactions.

6.7 Conclusion

In summary, in this work I have explored the possibility of using the phase shift information from the BAO signal and the CMB power spectra to constrain the strength of neutrino self-interactions. Due to neutrino self-interactions in the early Universe, delayed neutrino free-streaming alters the typical phase shift signal expected from Standard Model neutrinos.

I explore the phase shift, both in the BAO signal and the CMB power spectra, in the context of neutrinos with universal self-interactions parameterized by G_{eff} . It is interesting that I show that the impact on the phase shift with increasing G_{eff} has a similar impact on the phase shift signal, both in the CMB and BAO, that is shown in the previous work of Choi et al. (2018). Choi et al. (2018) studies the CMB phase shift in the presence of free-streaming light relics with a delayed free-streaming epoch relative to Standard Model neutrinos, and shows that the CMB power spectra can constrain a delay in free-streaming light relics and consequently a change to N_{eff} in the context of a model that involves copies of the Standard Model (and thus additional neutrinos with delayed free-streaming). These results help to validate those I have shown.

I construct a template for the phase shift that captures the impact of altering G_{eff} , both for the BAO oscillatory spectrum and the CMB power spectra, by building on the existing templates by Baumann et al. (2018, 2019) and Montefalcone et al. (2025c). These new templates allow one to model the phase shift for both alterations to N_{eff} through β_ϕ and $\log_{10}(G_{\text{eff}})$ through two parameters, $A(G_{\text{eff}})$ and $B(G_{\text{eff}})$. $A(G_{\text{eff}})$ modulates the amplitude of the phase shift, while $B(G_{\text{eff}})$ introduces a scale dependent

exponential damping into the phase shift signal. The templates can be used to simultaneously or separately fit β_ϕ and $\log_{10}(G_{\text{eff}})$ in real data or produce forecasts for the information available in just the phase shift as I have done in this chapter. The results show that both the BAO and the CMB are able to constrain G_{eff} from the phase shift information, but the constraining power is strongly dependent on G_{eff} itself. While the BAO constraints alone are quite weak in most cases and likely will not be useful for any realistic case in the future, the BAO does provide complementary information to CMB constraints, improving the constraining power by up to 47% with the larger improvements corresponding to smaller G_{eff} . On its own, a *Planck*-like measurement of the phase shift in the CMB should be able to constrain G_{eff} with an uncertainty between $\pm 0.15 - 1.07$, for $\log_{10}(G_{\text{eff}})$ between $[-3, 0.0]$. However, for weaker interactions, it is very difficult for BAO or CMB data to differentiate between weaker interaction strengths, since it is not possible to distinguish changes to the phase in such cases. As such, it will be likely that in this case, stronger interaction models can simply be disfavoured by the phase shift information in any data; only upper bounds on $\log_{10}(G_{\text{eff}})$ may be possible. In future work, it will be possible to fit data from CMB experiments to constrain self-interactions from the phase shift as was done in Montefalcone et al. (2025b), although their constraints are phrased in terms of a decoupling redshift rather than G_{eff} . Another difference is that the scale dependence of the phase shift we include above which is neglected in Montefalcone et al. (2025b) given they find it to be a negligible effect in an analysis on current CMB data. They also considers a broader range of temperature dependence for the neutrino interactions than we do, as implemented in their nuCLASS¹⁶ Boltzmann solver. In their Appendix B, one can see how changing the temperature dependence of the interactions impacts the phase shift signature, with the scale dependence being negligible for weaker temperature scaling but becoming more significant for interactions with a steeper temperature dependence. The temperature dependence for the interaction rate is also studied in Pal et al. (2025). It would be interesting to determine how the subtle scale dependence we have included in our phase-shift template impacts the constraints on the neutrino decoupling redshift and G_{eff} .

Overall, it is interesting to study how different neutrino interactions may change the phase in the CMB or BAO, which is a robust signature of the CvB, and constraints from cosmological data can be complementary to particle physics experiments which may not be able to access interactions between particles in the conditions that occur in the early Universe. As opposed to studying the general effects of neutrino self-interactions via the G_{eff} parameterization which changes the broadband shape of the matter power spectrum in a manner that may lead to strong degeneracies with A_s and n_s , studying the phase shift alone may help improve constraints by isolating a robust impact of neutrino species that free-stream at early times or instead decouple at a delayed time in the case of self-interactions. While I found obtaining strong constraints from the BAO alone would be extremely challenging, there is potential for interesting constraints from CMB data (and combinations of this data with BAO information).

There may also be additional information from the phase shift that could be explored in other probes. In future work, it may be worth exploring forecasts in which the helium abundance Y_p , which

¹⁶<https://github.com/subhajitghosh-phy/nuCLASS.git>

describes the mass abundance of helium in the total helium and hydrogen, could be included in the CMB forecasts. This may weaken constraints because Y_p changes the silk damping scale and thus is expected to have some degeneracy with $\beta_\phi(N_{\text{eff}})$. In Libanore et al. (2025), it is shown that the 21 cm power spectrum is able to constrain G_{eff} ; these constraints could be complementary to those obtained from the phase shift in CMB and BAO data. In Montefalcone et al. (2025a), the phase shift is studied in the 21 cm power spectrum.

Various interesting extensions could be considered for this work. Firstly, this work only explores the phase shift as it arises for neutrinos with universal self-interactions; a model that is largely disfavored by experimental data (Lyu et al., 2021; Blinov et al., 2019). However, it is possible to have non-universal self-interactions, and it might be expected this would change the significance of the effect self-interactions has on the phase shift if only some neutrinos have delayed free-streaming. In future work, it may be possible to develop a template for such models. Interestingly, Das and Ghosh (2021) shows fits to *Planck* data for G_{eff} in which only one or two neutrino species self-interact, and this provides a better fit to the data than the case of three neutrinos self-interacting; in contrast, Brinckmann et al. (2021) does not find any preference for neutrino self-interactions when studying varying numbers of self-interacting or free-streaming species. Another possible extension could be to consider models of resonant neutrino self-interactions (Noriega et al., 2025) or models in which neutrinos ‘self-recouple’ at a later time instead of having delayed free-streaming (Berryman et al., 2023). In the case of resonant self-interactions explored in Noriega et al. (2025), for specific mediator masses m_ϕ , the neutrinos self-interact preferably at $z \sim 4 \times 10^4 \frac{m_\phi}{\text{eV}}$, which corresponds to the epoch of BAO propagation. Consequently, the signal of self-interactions can be boosted, and we might expect this to also have a different impact on the phase shift that could be explored. Additionally, for some mediator masses it may be possible for neutrinos to ‘self-recouple’; in this case, modes that enter the causal horizon before the neutrinos self-recouple will experience the same phase shift that we expect to see for Standard Model neutrinos, while those that enter afterwards may have a similar impact on the phase shift that we see in the case for self-interactions explored in this work.

Chapter 7

Conclusion

7.1 Conclusion

Over the years, a wealth of data has been collected for the study of cosmology. This has allowed us to make fundamental insights that have developed and improved the concordance model, Λ CDM. Measurements of galaxy rotation curves and clusters (Zwicky, 1933, 1937; Rubin and Ford Jr, 1970), and the CMB (Penzias and Wilson, 1965; Mather et al., 1994; Fixsen et al., 1996; Aghanim et al., 2020) have allowed us to infer the presence of dark matter. Galaxy cluster measurements (Efstathiou et al., 1990), SNe Type Ia (Phillips, 1993; Riess et al., 1998; Perlmutter et al., 1999; Abbott et al., 2024) and BAOs (Cole et al., 2005; Eisenstein et al., 2005; Dawson et al., 2012; Adame et al., 2025b,a; Abbott et al., 2019) have allowed us to infer the presence of dark energy.

Further to this, we have probed the large-scale structure to understand the matter distribution, the growth of structure, and extract more information about cosmological parameters across different scales. This has added to the picture of how the cosmic web has come to be at the present day. Galaxy surveys like 2dFGS (Cole et al., 2005), 6dFGS (Jones et al., 2009), SDSS (York et al., 2000; Dawson et al., 2012; Beutler et al., 2017), and now DESI (Aghamousa et al., 2016a,b; Adame et al., 2025b) have revolutionized cosmology with an unprecedented amount of data for probing the galaxy correlation function, the power spectrum, and making measurements of parameters such as the matter density Ω_m , the expansion rate H_0 and the sum of neutrino masses $\sum m_\nu$. Peculiar velocity data such as CosmicFlows-4 (Tully et al., 2023) has allowed us to perform additional tests of Λ CDM and measure the growth of structure in the local Universe. We will be able to construct more detailed maps and peculiar velocity surveys with upcoming data from LSST,¹ Euclid,² SKAO³, WALLABY⁴ and 4MOST.⁵

At present, the Λ CDM model is still treated as the concordance model. However, in Chapter 2

¹<https://rubinobservatory.org/>

²https://www.esa.int/Science_Exploration/Space_Science/Euclid

³<https://www.skao.int/en>

⁴<https://wallaby-survey.org/>

⁵<https://www.eso.org/public/teles-instr/paranal-observatory/surveytelescopes/vista/>

⁴most/

we discussed how recent analyses have suggested that Λ CDM may no longer remain the model that best describes the data. Recent measurements of SNe Ia (Rubin et al., 2025; Abbott et al., 2024) and combined measurements of BAO (Adame et al., 2025a,b) and CMB data (Aghanim et al., 2020) find that a w_0w_a CDM model, in which dark energy varies in its strength over time, may give a better fit. At this point though, the measurements have not yet reached a significance to definitively confirm this, and there is a chance future analyses with additional data will confirm or deny this conclusion.

This result may not necessarily be surprising. The Λ CDM model can be considered an ‘effective model’; it includes the impact of dark matter and dark energy, but it does not necessarily include a framework to predict or understand these constituents or how they should behave in relation to the expansion and growth of structure. It has – until now – best described our Universe based on our empirical observations, and is motivated by predictions of a constant ‘vacuum energy’ by quantum field theory (QFT) and its simplicity as a model. However, given the massive discrepancy between the QFT predictions and the observed dark energy (Weinberg, 1989), scientists should remain open to the possibility of alternative models or simply exploring the extent of our understanding of additional *known* parameters such as neutrino properties that can be included in our analyses and impact constraints (Adame et al., 2025b; Elbers et al., 2025).

Other smaller problems have also arisen in our understanding of cosmology. In particular, anomalies in the local Universe may help point us towards better models. As of now, there are measurements of large bulk flows (Kashlinsky et al., 2008; Watkins et al., 2009; Feldman et al., 2010; Watkins and Feldman, 2015b; Peery et al., 2018; Howlett et al., 2022; Watkins et al., 2023; Whitford et al., 2023; Duangchan et al., 2025) that are not able to be predicted by Λ CDM. However, there are some contradictory measurements that find results consistent with Λ CDM (Nusser and Davis, 2011; Hong et al., 2014; Ma and Pan, 2014; Hoffman et al., 2015; Qin et al., 2018, 2021) or not as strongly in tension (Courtois et al., 2025). These contradictions remain unresolved; attempting to understand these measurements may help us connect measurements of structure growth in the local Universe to the overall picture of the universal expansion history. At this point in time, it is not entirely possible to know if anomalous bulk flow measurements suggest an issue with the Λ CDM model and the assumptions in the Cosmological Principle, or if instead, noisy peculiar velocity measurements and statistical fluctuations are behind a large-amplitude bulk flow measurement.

In Chapter 4 of this thesis, I studied in great detail the influence of survey geometry in the context of bulk flow measurements made with two well-known bulk flow estimators. My results found that in general, the uncertainties derived on bulk flow amplitudes from these estimators may underestimate the real variance of the measurement. This conclusion is based on tests in mock datasets. Correcting the size of the uncertainties is likely the most important result of this work, given the uncertainty is crucial in determining the level of tension a measurement has with any given model, and thus the significance of the ‘bulk flow tension’. Consequently, bulk flow measurements made with real data should always be compared with a measurement applied to mock data through the same pipeline. Anomalous bulk flow measurements are unlikely to be direct consequence of survey geometry, but non-linearities in the galaxy peculiar velocity field may lead to larger uncertainty in the measurements.

An additional area of uncertainty relates to the properties of neutrinos, one of the most poorly understood particles in the Standard Model. Their influence on the cosmological expansion rate and growth of structure is an ongoing area of research. While the effects of neutrinos in cosmology is subtle, their impact is degenerate with parameters that affect the expansion rate, such as H_0 . Therefore, quantifying this impact is important in the era of precision cosmology to reduce cosmological tensions, and to further explore new cosmological models. The motivation for studying these impacts is also driven by uncertainties about how neutrinos derive mass in the context of the Standard Model, motivating the possibility of models in which neutrinos have new interactions with Standard Model or ‘dark sector’ particles (thus potentially exploring models of dark matter). On top of this, I mentioned in Chapter 2 an interesting result has been found relating to neutrino properties in the DESI data; a combined analysis of CMB data (Aghanim et al., 2020) with DESI BAO data (Adame et al., 2025b) found a preference for the sum of neutrino masses $\sum m_\nu$ to be below the minimum possible value of ~ 0.058 eV determined by neutrino oscillation experiments, although the constraints are more flexible in the context of time-varying dark energy (Elbers et al., 2025). Regardless, the results give room for cosmologists to suggest mechanisms that would allow for such measurements to be derived from cosmological data (Craig et al., 2024; Green and Meyers, 2025). Constraints in the literature also suggest the possibility of neutrinos with non-standard self-interactions (Kreisch et al., 2020, 2024; Camarena et al., 2023a; Camarena and Cyr-Racine, 2025; Poudou et al., 2025; He et al., 2025). In the models often explored in the literature, the neutrinos universally (all species) self-interact; this is used as an ‘effective model’ to explore possible non-standard interactions. However, the constraints usually have a bimodal distribution, likely driven by parameter degeneracies with A_s (the power spectrum amplitude) and n_s (the spectral index). If neutrinos really do have non-standard properties, it is then beneficial to study this both for improving constraints on other cosmological parameters of interest, but also for improving our understanding of neutrinos in the Standard Model, or even dark matter candidates.

One manner in which we can explore non-standard neutrino physics in cosmology is via the parameter N_{eff} , which parameterises the effective number of neutrino species. Neutrino species with non-standard properties, or the presence of additional species beyond those known in the Standard Model, are captured by this parameter in cosmological data. The parameter N_{eff} is degenerate with H_0 , $\sum m_\nu$ and other parameters that vary the expansion rate, but free-streaming radiation induces a unique phase shift in both the CMB power spectra and the BAO wiggles. In Chapter 5, I explore the ability of BAO data from DESI to constrain the phase shift amplitude of Standard Model neutrinos β_ϕ , which can be related to N_{eff} (Whitford et al., 2025; Elbers et al., 2025). My work contributed to the development of software to make more precise measurements of the phase shift with future data, to infer the presence of the CvB in the BAOs, and additionally to help constrain neutrino properties. In this work, it was found that more study will need to be done in the future to improve BAO fitting pipelines, given the phase shift measurement is somewhat sensitive to the choice of the modelling of the broadband power spectrum; this is more important for future data that will have greater constraining power. The results with the DESI data found that the phase shift amplitude is larger than would be

expected in the context of Λ CDM (corresponding to a larger N_{eff} , which is in contradiction with CMB measurements), but this could plausibly be due to an upwards fluctuation in the data. When allowing the lensing amplitude in the Planck data (used as a prior in the constraints, Aghanim et al., 2020) to vary freely, or allowing for time-varying dark energy with the w CDM model, the constraints shift to an amplitude that is closer to the Standard Model expectation. It may also be important to include the impact of β_ϕ simply for future constraints on the standard BAO parameters for measuring the expansion history, because the degeneracy of β_ϕ with α_{iso} may also alter constraints on parameters such as w_0 and w_a .

Finally, in Chapter 6 I studied the impacts of neutrinos in cosmology further, but specifically in the context of a model in which neutrinos have non-standard self-interactions. This is an extension to Chapter 5, as I construct templates (building on existing ones) for the phase shift to capture the impact of not just N_{eff} but also the self-interaction strength G_{eff} . If neutrinos have self-interactions, this alters the epoch at which they begin to free-stream; changes to G_{eff} delay free-streaming and thus alter the amplitude of the phase shift relative to neutrinos with Standard Model behaviour. I show how this effect alters the phase shift in the BAO and CMB power spectra, building on work by Choi et al. (2018). I also show that it may be possible to use this phase shift effect to show or disfavour a preference for neutrinos with strong self-interactions in BAO or the CMB data. This may help to constrain G_{eff} in a manner that does not have degeneracy with A_s or n_s . This will help us move forward to test more promising models with future cosmological data.

7.2 Future work

7.2.1 Bulk flows

There are a number of extensions regarding bulk flows for future work, some of which are in progress. Primarily, we can expect to see new peculiar velocity (PV) data arising from surveys including SKA, 4MOST, WALLABY, DESI and LSST, increasing the depth of PV measurements. In particular, the 4MOST and WALLABY surveys will map out the sky in the southern hemisphere, helping to even out the depths of PV data available across different regions of the sky, providing a better ‘3D’ picture of the local peculiar velocity field. Currently, the CosmicFlows-4 dataset is ‘top heavy’ with much greater depth of data in the northern sky. The new data may also help us to look beyond large structures such as Shapley (Proust et al., 2006) and ‘the Great Attractor’ (Radburn-Smith et al., 2006; Kraan-Korteweg et al., 2017), which currently appear to source the large bulk flows seen in the literature; if the bulk flow amplitude is reduced on scales beyond these structures, this will indicate how much of a role they play.

If large bulk flow measurements persist with the addition of new data, it may indicate cosmologists should revisit assumptions surrounding the validity of the Cosmological Principle, relativistic effects, or the role that back-reaction plays in local measurements of the large-scale structure (this has been studied recently in Giani et al., 2025). It may be that the assumption of homogeneity and isotropy

breaks down on small scales, such that our measurements in the local Universe cannot correctly be interpreted with this assumption. There has also been recent work attempting to study more suitable metrics for such cases, for example Camarena et al. (2023b) studies the inhomogeneous ‘LTB’ model. Additionally, there have been studies to show that the dipole seen in radio galaxy catalogues is larger than that seen in the CMB dipole (Singal, 2023, 2025; Dam et al., 2023); the CMB dipole is generally interpreted to be due to the motion of our solar system with respect to the CMB (which includes the Sun’s motion through the Milky Way and the Milky Way’s motion through the Local Group), but we expect this to match the dipole seen in number counts of radio objects. The larger-than-expected dipole may be due to systematics, but could be interpreted as a breakdown of the Cosmological Principle. The FLRW metric also predicts the expansion of the Universe, and thus the distances to objects as a function of z , in the scenario that structure growth itself does not significantly impact this expansion (back-reaction). To some extent, this is the case because back-reaction is a difficult impact to model. However, if back-reaction is indeed significant on local scales, it may lead to apparently large bulk flows, or anisotropic measurements of H_0 across the sky, which has been measured in X-ray clusters (Kashlinsky et al., 2008). Regarding relativistic effects, Heinesen (2023) shows that neglecting gravitational redshift or time-evolution of the velocity field may lead to the bulk flow to be overestimated by as much as 15% at $z \sim 0.1$; it is thus practical to consider whether other relativistic effects are at play in our calculations.

Further to this, while my research in Chapter 4 suggests the bulk flow estimators themselves are unlikely to lead to *inaccurate* bulk flows, the estimators themselves are not overly *precise* when applied to realistic data that is noisy and contains non-linearities in the measurements. With additional data, even though the number of objects increasing should reduce Poisson noise, the noise due to the intrinsic scatter in scaling relations such as the Fundamental Plane or Tully-Fisher relation will not improve; the scatter in the individual PV measurements increases with increasing depth using these distance estimators. As such, it will be more important and challenging than ever to make improvements to bulk flow estimators to try increase our precision in bulk flow estimates.

Concluding the discussion of future work regarding bulk flows, it should be noted that it is unlikely that a preference for a time-varying dark energy models will alter the bulk flow amplitude. For these models, the matter power spectrum $P(k)$ does not vary significantly and even does not provide strong constraining power of w_0, w_a in the DESI data compared to the BAO data; from the relation between the power spectrum and bulk flow amplitude seen in equation 4.15, it seems that a very large change is required in the matter power spectrum to lead to a large change in the bulk flow amplitude.

7.2.2 The phase shift induced by free-streaming relics

As discussed in Chapter 5 and 6, the CMB is able to put very strong constraints on N_{eff} via a constraint on the amplitude of the phase shift for Standard Model free streaming relics (Follin et al., 2015; Montefalcone et al., 2025c). However, it will nonetheless be interesting to see future constraints on the amplitude β_ϕ in future BAO data, such as in the final DESI data release (the Year 5 data), as it

shows a detection of the CvB in LSS data. It will also be important to see whether the amplitude of β_ϕ that is detected is more consistent with the expectation of $N_{\text{eff}} = 3.044$. As mentioned earlier, this measurement may require further work on the modelling of the broadband power spectrum to ensure it does not bias the constraints, and it will be important to include β_ϕ as a free parameter in future work given the degeneracy between β_ϕ and α_{iso} as constraints become more precise. Forecasts show (Baumann et al., 2018) that a cosmic-variance limited BAO survey should be able to constrain $\sigma(\beta_\phi) \sim 0.1$ corresponding to $\sigma(N_{\text{eff}}) \sim 0.5$.

I additionally mentioned above that future work can also involve a measurement of a phase shift that results in the case of light relics that do not immediately free-stream, following the work of Chapter 6. The application of this template to real data will provide a test of models that involve non-standard light relics in the early Universe, in a manner that removes degeneracies with parameters such as A_s and n_s . This will help us to determine whether it is worth moving forward with further study of models in which there are non-standard neutrino interactions or other more complex models that mimic this signal. With upcoming data from DESI combined with CMB data, one could potentially make an interesting constraint to support or disfavour a phase shift caused by strong neutrino self-interactions, or any model that mimics this phase shift signal. If it is indeed found that there is any preference at all for strong self-interactions, it will be worth exploring more realistic models of self-interactions, such as in the case they are non-universal or they have resonances at particular redshifts. These results will be complementary to particle physics experiments.

Bibliography

- Abareshi, B., Aguilar, J., Ahlen, S., Alam, S., Alexander, D. M., Alfarsy, R., Allen, L., Prieto, C. A., Alves, O., Ameel, J., et al. (2022). Overview of the instrumentation for the Dark Energy Spectroscopic Instrument. *The Astronomical Journal*, 164(5):207.
- Abbott, T., Abdalla, F., Alarcon, A., Allam, S., Andrade-Oliveira, F., Annis, J., Avila, S., Banerji, M., Banik, N., Bechtol, K., et al. (2019). Dark Energy Survey Year 1 Results: Measurement of the Baryon Acoustic Oscillation scale in the distribution of galaxies to redshift 1. *Monthly Notices of the Royal Astronomical Society*, 483(4):4866–4883.
- Abbott, T., Acevedo, M., Aguena, M., Alarcon, A., Allam, S., Alves, O., Amon, A., Andrade-Oliveira, F., Annis, J., Armstrong, P., et al. (2024). The Dark Energy Survey: Cosmology Results With ~ 1500 New High-redshift Type Ia Supernovae Using The Full 5-year Dataset. *The Astrophysical Journal Letters*, 973(1):L14.
- Abdalla, E., Abellán, G. F., Aboubrahim, A., Agnello, A., Akarsu, Ö., Akrami, Y., Alestas, G., Aloni, D., Amendola, L., Anchordoqui, L. A., et al. (2022). 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.
- Abolfathi, B., Aguado, D., Aguilar, G., Prieto, C. A., Almeida, A., Ananna, T. T., Anders, F., Anderson, S. F., Andrews, B. H., Anguiano, B., et al. (2018). The fourteenth data release of the Sloan Digital Sky Survey: First spectroscopic data from the extended Baryon Oscillation Spectroscopic Survey and from the second phase of the Apache Point Observatory Galactic Evolution Experiment. *The Astrophysical Journal Supplement Series*, 235(2):42.
- Adame, A., Aguilar, J., Ahlen, S., Alam, S., Aldering, G., Alexander, D., Alfarsy, R., Prieto, C. A., Alvarez, M., Alves, O., et al. (2024). Validation of the scientific program for the Dark Energy Spectroscopic Instrument. *The Astronomical Journal*, 167(2):62.
- Adame, A., Aguilar, J., Ahlen, S., Alam, S., Alexander, D., Alvarez, M., Alves, O., Anand, A., Andrade, U., Armengaud, E., et al. (2025a). DESI 2024 III: Baryon Acoustic Oscillations from Galaxies and Quasars. *Journal of Cosmology and Astroparticle Physics*, 2025(04):012.

- Adame, A., Aguilar, J., Ahlen, S., Alam, S., Alexander, D., Alvarez, M., Alves, O., Anand, A., Andrade, U., Armengaud, E., et al. (2025b). DESI 2024 VI: Cosmological constraints from the measurements of baryon acoustic oscillations. *Journal of Cosmology and Astroparticle Physics*, 2025(02):021.
- Agarwal, S., Feldman, H. A., and Watkins, R. (2012). Testing the minimum variance method for estimating large-scale velocity moments. *Monthly Notices of the Royal Astronomical Society*, 424(4):2667–2675.
- Aghamousa, A., Aguilar, J., Ahlen, S., Alam, S., Allen, L. E., Prieto, C. A., Annis, J., Bailey, S., Balland, C., Ballester, O., et al. (2016a). The DESI experiment part I: science, targeting, and survey design. *arXiv preprint arXiv:1611.00036*.
- Aghamousa, A., Aguilar, J., Ahlen, S., Alam, S., Allen, L. E., Prieto, C. A., Annis, J., Bailey, S., Balland, C., Ballester, O., et al. (2016b). The DESI experiment part II: instrument design. *arXiv preprint arXiv:1611.00037*.
- Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A. J., Barreiro, R., Bartolo, N., Basak, S., et al. (2020). Planck 2018 results-VI. Cosmological parameters. *Astronomy & Astrophysics*, 641:A6.
- Agrawal, A., Okumura, T., and Futamase, T. (2019). Constraining neutrino mass and dark energy with peculiar velocities and lensing dispersions of Type Ia supernovae. *Physical Review D*, 100(6):063534.
- Alcock, C. and Paczyński, B. (1979). An evolution free test for non-zero cosmological constant. *Nature*, 281(5730).
- Amendola, L. and Quartin, M. (2021). Measuring the Hubble function with standard candle clustering. *Monthly Notices of the Royal Astronomical Society*, 504(3):3884–3889.
- Andersen, P., Davis, T. M., and Howlett, C. (2016). Cosmology with peculiar velocities: observational effects. *Monthly Notices of the Royal Astronomical Society*, 463(4):4083–4092.
- Ballinger, W., Peacock, J., and Heavens, A. (1996). Measuring the cosmological constant with redshift surveys. *Monthly Notices of the Royal Astronomical Society*, 282(3):877–888.
- Barenboim, G., Denton, P. B., and Oldengott, I. M. (2019). Constraints on inflation with an extended neutrino sector. *Physical Review D*, 99(8):083515.
- Bashinsky, S. and Seljak, U. (2004). Neutrino perturbations in CMB anisotropy and matter clustering. *Physical Review D*, 69(8).
- Baumann, D., Beutler, F., Flauger, R., Green, D., Slosar, A., Vargas-Magaña, M., Wallisch, B., and Yeche, C. (2019). First constraint on the neutrino-induced phase shift in the spectrum of baryon acoustic oscillations. *Nature Physics*, 15(5):465–469.

- Baumann, D., Green, D., Meyers, J., and Wallisch, B. (2016). Phases of New Physics in the CMB. *Journal of Cosmology and Astroparticle Physics*, 2016(01):007.
- Baumann, D., Green, D., and Wallisch, B. (2018). Searching for light relics with large-scale structure. *Journal of Cosmology and Astroparticle Physics*, 2018(08):029.
- Belotsky, K., Sudarikov, A., and Khlopov, M. Y. (2001). Constraint on anomalous 4ν interaction. *Physics of Atomic Nuclei*, 64(9):1637–1642.
- Bennett, J. J., Buldgen, G., De Salas, P. F., Drewes, M., Gariazzo, S., Pastor, S., and Wong, Y. Y. Y. (2021). Towards a precision calculation of N_{eff} in the Standard Model II: Neutrino decoupling in the presence of flavour oscillations and finite-temperature QED. *JCAP*, 04:073.
- Berbig, M., Jana, S., and Trautner, A. (2020). The hubble tension and a renormalizable model of gauged neutrino self-interactions. *Physical Review D*, 102(11):115008.
- Berkov, A., Nikitin, Y. P., Sudarikov, A., and Khlopov, M. Y. (1987). Possible experimental search for anomalous 4 neutrino interaction. *Yad. Fiz.*, 46:1729–1737.
- Berkov, A., Nikitin, Y. P., Sudarikov, A., and Khlopov, M. Y. (1988). Possible manifestations of anomalous 4 neutrino interaction. *Sov. J. Nucl. Phys.*, 48:497–501.
- Bernal, J. L., Smith, T. L., Boddy, K. K., and Kamionkowski, M. (2020). Robustness of baryon acoustic oscillation constraints for early-Universe modifications of Λ CDM cosmology. *Physical Review D*, 102(12):123515.
- Berryman, J. M., Blinov, N., Brdar, V., Brinckmann, T., Bustamante, M., Cyr-Racine, F.-Y., Das, A., de Gouvêa, A., Denton, P. B., Dev, P. B., et al. (2023). Neutrino self-interactions: A white paper. *Physics of the Dark Universe*, 42:101267.
- Beutler, F., Seo, H.-J., Saito, S., Chuang, C.-H., Cuesta, A. J., Eisenstein, D. J., Gil-Marín, H., Grieb, J. N., Hand, N., Kitaura, F.-S., et al. (2017). The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: anisotropic galaxy clustering in Fourier space. *Monthly Notices of the Royal Astronomical Society*, 466(2):2242–2260.
- Blas, D., Lesgourgues, J., and Tram, T. (2011). The cosmic linear anisotropy solving system (CLASS). Part II: Approximation schemes. *Journal of Cosmology and Astroparticle Physics*, 2011(07):034.
- Blinov, N., Kelly, K. J., Krnjaic, G., and McDermott, S. D. (2019). Constraining the self-interacting neutrino interpretation of the Hubble tension. *Physical Review Letters*, 123(19):191102.
- Brinckmann, T., Chang, J. H., Du, P., and LoVerde, M. (2023). Confronting interacting dark radiation scenarios with cosmological data. *Physical Review D*, 107(12):123517.
- Brinckmann, T., Chang, J. H., and LoVerde, M. (2021). Self-interacting neutrinos, the Hubble parameter tension, and the cosmic microwave background. *Physical Review D*, 104(6):063523.

- Bunn, E. F. and Hogg, D. W. (2009). The kinematic origin of the cosmological redshift. *American Journal of Physics*, 77(8):688–694.
- Burkey, D. and Taylor, A. N. (2004). Prospects for galaxy–mass relations from the 6dF Galaxy Survey. *Monthly Notices of the Royal Astronomical Society*, 347(1):255–268.
- Camarena, D. and Cyr-Racine, F.-Y. (2025). Strong constraints on a simple self-interacting neutrino cosmology. *Physical Review D*, 111(2):023504.
- Camarena, D., Cyr-Racine, F.-Y., and Houghteling, J. (2023a). Confronting self-interacting neutrinos with the full shape of the galaxy power spectrum. *Physical Review D*, 108(10):103535.
- Camarena, D., Marra, V., Sakr, Z., Nesseris, S., Da Silva, A., García-Bellido, J., Fleury, P., Lombriser, L., Martinelli, M., Martins, C. J. A. P., et al. (2023b). Euclid: Testing the Copernican principle with next-generation surveys. *Astronomy & Astrophysics*, 671:A68.
- Carroll, S. M. (2019). *Spacetime and geometry*. Cambridge University Press.
- Castro, T., Quartin, M., and Benitez-Herrera, S. (2016). Turning noise into signal: learning from the scatter in the Hubble diagram. *Physics of the dark universe*, 13:66–76.
- Chaussidon, E., Yèche, C., Palanque-Delabrouille, N., Alexander, D. M., Yang, J., Ahlen, S., Bailey, S., Brooks, D., Cai, Z., Chabanier, S., et al. (2023). Target Selection and Validation of DESI Quasars. *The Astrophysical Journal*, 944(1):107.
- Chen, S.-F., Howlett, C., White, M., McDonald, P., Ross, A. J., Seo, H.-J., Padmanabhan, N., Aguilar, J., Ahlen, S., Alam, S., et al. (2024a). Baryon Acoustic Oscillation Theory and Modelling Systematics for the DESI 2024 results. *Monthly Notices of the Royal Astronomical Society*, 534:544–574.
- Chen, X., Ding, Z., Paillas, E., Nadathur, S., Seo, H., Chen, S., Padmanabhan, N., White, M., de Mattia, A., McDonald, P., et al. (2024b). Extensive analysis of reconstruction algorithms for DESI 2024 baryon acoustic oscillations. *arXiv preprint arXiv:2411.19738*.
- Choi, G., Chiang, C.-T., and LoVerde, M. (2018). Probing decoupling in dark sectors with the cosmic microwave background. *Journal of Cosmology and Astroparticle Physics*, 2018(06):044.
- Choudhury, S. R., Hannestad, S., and Tram, T. (2021). Updated constraints on massive neutrino self-interactions from cosmology in light of the h_0 tension. *Journal of Cosmology and Astroparticle Physics*, 2021(03):084.
- Choudhury, S. R., Hannestad, S., and Tram, T. (2022). Massive neutrino self-interactions and inflation. *Journal of Cosmology and Astroparticle Physics*, 2022(10):018.
- Chuang, C.-H., Kitaura, F.-S., Prada, F., Zhao, C., and Yepes, G. (2015). EZmocks: extending the Zel’dovich approximation to generate mock galaxy catalogues with accurate clustering statistics. *Monthly Notices of the Royal Astronomical Society*, 446(3):2621–2628.

- Cole, S., Percival, W. J., Peacock, J. A., Norberg, P., Baugh, C. M., Frenk, C. S., Baldry, I., Bland-Hawthorn, J., Bridges, T., Cannon, R., et al. (2005). The 2dF Galaxy Redshift Survey: power-spectrum analysis of the final data set and cosmological implications. *Monthly Notices of the Royal Astronomical Society*, 362(2):505–534.
- Collaboration, A., Robitaille, T. P., Tollerud, E. J., Greenfield, P., Droettboom, M., Bray, E., Aldcroft, T., Davis, M., Ginsburg, A., Price-Whelan, A. M., Kerzendorf, W. E., et al. (2013). Astropy: A community Python package for astronomy. *Astronomy & Astrophysics*, 558:A33.
- Courtois, H. M., Mould, J., Hollinger, A. M., Dupuy, A., and Zhang, C.-P. (2025). In search for the Local Universe dynamical homogeneity scale with CF4++ peculiar velocities. *arXiv preprint arXiv:2502.01308*.
- Courtois, H. M., Tully, R. B., Fisher, J. R., Bonhomme, N., Zavodny, M., and Barnes, A. (2009). The Extragalactic Distance Database: All Digital HI Profile Catalog. *The Astronomical Journal*, 138(6):1938–1956.
- Craig, N., Green, D., Meyers, J., and Rajendran, S. (2024). No ν_s is good news. *Journal of High Energy Physics*, 2024(9):1–33.
- Crivellin, A. and Mellado, B. (2024). Anomalies in particle physics and their implications for physics beyond the standard model. *Nature Reviews Physics*, 6(5):294–309.
- Dam, L., Lewis, G. F., and Brewer, B. J. (2023). Testing the cosmological principle with CatWISE quasars: a bayesian analysis of the number-count dipole. *Monthly Notices of the Royal Astronomical Society*, 525(1):231–245.
- Das, A., Dev, P., Gao, C., Ghosh, S., and Kim, T. (2025). Impostor Among ν_s : Dark Radiation Masquerading as Self-Interacting Neutrinos. *arXiv preprint arXiv:2506.08085*.
- Das, A. and Ghosh, S. (2021). Flavor-specific interaction favors strong neutrino self-coupling in the early universe. *Journal of Cosmology and Astroparticle Physics*, 2021(07):038.
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., Anderson, S. F., Aubourg, É., Bailey, S., Barkhouser, R. H., Bautista, J. E., Beifiori, A., Berlind, A. A., et al. (2012). The Baryon oscillation spectroscopic survey of SDSS-III. *The Astronomical Journal*, 145(1):10.
- de Salas, P. F. and Pastor, S. (2016). Relic neutrino decoupling with flavour oscillations revisited. *JCAP*, 07:051.
- Deppisch, F. F., Graf, L., Rodejohann, W., and Xu, X.-J. (2020). Neutrino self-interactions and double beta decay. *Physical Review D*, 102(5):051701.
- Di Valentino, E., Mena, O., Pan, S., Visinelli, L., Yang, W., Melchiorri, A., Mota, D. F., Riess, A. G., and Silk, J. (2021). In the realm of the Hubble tension—a review of solutions. *Classical and Quantum Gravity*, 38(15):153001.

- Djorgovski, S. and Davis, M. (1987). Fundamental properties of elliptical galaxies. *The Astrophysical Journal*, 313:59–68.
- Dressler, A., Lynden-Bell, D., Burstein, D., Davies, R. L., Faber, S., Terlevich, R., and Wegner, G. (1987). Spectroscopy and photometry of elliptical galaxies. I-A new distance estimator. *Astrophysical Journal, Part 1 (ISSN 0004-637X)*, vol. 313, Feb. 1, 1987, p. 42-58., 313:42–58.
- Duangchan, C., Valade, A., Libeskind, N., and Steinmetz, M. (2025). Prior-free cosmological parameter estimation of Cosmicflows-4. *arXiv preprint arXiv:2507.22236*.
- Efstathiou, G., Sutherland, W. J., and Maddox, S. (1990). The cosmological constant and cold dark matter. *Nature*, 348(6303):705–707.
- Eisenstein, D. J., Zehavi, I., Hogg, D. W., Scoccamarro, R., Blanton, M. R., Nichol, R. C., Scranton, R., Seo, H.-J., Tegmark, M., Zheng, Z., et al. (2005). Detection of the baryon acoustic peak in the large-scale correlation function of SDSS luminous red galaxies. *The Astrophysical Journal*, 633(2):560.
- Elbers, W., Aviles, A., Noriega, H., Chebat, D., Menegas, A., Frenk, C., Garcia-Quintero, C., Gonzalez, D., Ishak, M., Lahav, O., et al. (2025). Constraints on Neutrino Physics from DESI DR2 BAO and DR1 Full Shape. *arXiv preprint arXiv:2503.14744*.
- Elgarøy, Ø., Lahav, O., Percival, W., Peacock, J., Madgwick, D., Bridle, S., Baugh, C., Baldry, I., Bland-Hawthorn, J., Bridges, T., et al. (2002). New upper limit on the total neutrino mass from the 2 degree field galaxy redshift survey. *Physical Review Letters*, 89(6):061301.
- Esteban, I., Pandey, S., Brdar, V., and Beacom, J. F. (2021). Probing secret interactions of astrophysical neutrinos in the high-statistics era. *Physical Review D*, 104(12):123014.
- Fakhouri, H., Boone, K., Aldering, G., Antilogus, P., Aragon, C., Bailey, S., Baltay, C., Barbary, K., Baugh, D., Bongard, S., et al. (2015). Improving cosmological distance measurements using twin Type Ia supernovae. *The Astrophysical Journal*, 815(1):58.
- Feldman, H. A., Kaiser, N., and Peacock, J. A. (1994). Power spectrum analysis of three-dimensional redshift surveys. *The Astrophysical Journal*, 426:23.
- Feldman, H. A., Watkins, R., and Hudson, M. J. (2010). Cosmic flows on $100 h^{-1}$ Mpc scales: standardized minimum variance bulk flow, shear and octupole moments. *Monthly Notices of the Royal Astronomical Society*, 407(4):2328–2338.
- Fisher, J. and Tully, R. (1981). Upper limits on the space density of intergalactic neutral hydrogen clouds. *The Astrophysical Journal*, 243:L23–L26.
- Fixsen, D. J., Cheng, E. S., Gales, J. M., Mather, J. C., Shafer, R. A., and Wright, E. L. (1996). The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set. *The Astrophysical Journal*, 473(2):576.

- Follin, B., Knox, L., Millea, M., and Pan, Z. (2015). First detection of the acoustic oscillation phase shift expected from the cosmic neutrino background. *Physical Review Letters*, 115(9):091301.
- Font-Ribera, A., McDonald, P., Mostek, N., Reid, B. A., Seo, H.-J., and Slosar, A. (2014). DESI and other Dark Energy experiments in the era of neutrino mass measurements. *Journal of Cosmology and Astroparticle Physics*, 2014(05):023.
- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J. (2013). emcee: the MCMC hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306.
- Forero-Sánchez, D., Rashkovetskyi, M., Alves, O., de Mattia, A., Nadathur, S., Zarrouk, P., Gil-Marín, H., Ding, Z., Yu, J., Andrade, U., et al. (2024). Analytical and EZmock covariance validation for the DESI 2024 results. *arXiv preprint arXiv:2411.12027*.
- Friedman, A. (1922). On the Curvature of space. *Z. Phys.*, 10:377–386.
- Fukuda, Y., Hayakawa, T., Ichihara, E., Inoue, K., Ishihara, K., Ishino, H., Itow, Y., Kajita, T., Kameda, J., Kasuga, S., et al. (1998). Evidence for oscillation of atmospheric neutrinos. *Physical review letters*, 81(8):1562.
- Garcia-Quintero, C., Mena-Fernández, J., Rocher, A., Yuan, S., Hadzhiyska, B., Alves, O., Rashkovetskyi, M., Seo, H., Padmanabhan, N., Nadathur, S., et al. (2025). HOD-dependent systematics in Emission Line Galaxies for the DESI 2024 BAO analysis. *Journal of Cosmology and Astroparticle Physics*, 2025(01):132.
- Giani, L., Howlett, C., Ruggeri, R., Bianchini, F., Said, K., and Davis, T. M. (2023). Cross-correlating radial peculiar velocities and CMB lensing convergence. *Journal of Cosmology and Astroparticle Physics*, 2023(05):002.
- Giani, L., Von Martens, R., and Camilleri, R. (2025). Novel Approach to Cosmological Nonlinearities as an Effective Fluid. *Physical Review Letters*, 135(7):071004.
- Gil-Marín, H., Percival, W. J., Verde, L., Brownstein, J. R., Chuang, C.-H., Kitaura, F.-S., Rodríguez-Torres, S. A., and Olmstead, M. D. (2016). The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: RSD measurement from the power spectrum and bispectrum of the DR12 BOSS galaxies. *Monthly Notices of the Royal Astronomical Society*, 460(4):4188–4209.
- Green, D. and Meyers, J. (2025). Cosmological preference for a negative neutrino mass. *Physical Review D*, 111(8):083507.
- Green, D. and Ridgway, A. K. (2020). The phase of the BAO on observable scales. *Journal of Cosmology and Astroparticle Physics*, 2020(12):050.
- Hadzhiyska, B., White, M. J., Chen, X., Garrison, L. H., DeRose, J., Padmanabhan, N., Garcia-Quintero, C., Mena-Fernández, J., Chen, S.-F., Seo, H.-J., et al. (2023). Mitigating the noise of DESI mocks using analytic control variates. *The Open Journal of Astrophysics*, 6.

- Hahn, C., Wilson, M. J., Ruiz-Macias, O., Cole, S., Weinberg, D. H., Moustakas, J., Kremin, A., Tinker, J. L., Smith, A., Wechsler, R. H., et al. (2023). The DESI bright galaxy survey: final target selection, design, and validation. *The Astronomical Journal*, 165(6):253.
- Hamuy, M. and Pinto, P. A. (2002). Type II supernovae as standardized candles. *The Astrophysical Journal*, 566(2):L63–L65.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825):357–362.
- Hartlap, J., Simon, P., and Schneider, P. (2007). Why your model parameter confidences might be too optimistic. Unbiased estimation of the inverse covariance matrix. *Astronomy & Astrophysics*, 464(1):399–404.
- Haynes, M. P., Giovanelli, R., Kent, B. R., Adams, E. A., Balonek, T. J., Craig, D. W., Fertig, D., Finn, R., Giovanardi, C., Hallenbeck, G., et al. (2018). The Arecibo Legacy Fast ALFA Survey: The ALFALFA Extragalactic HI Source Catalog. *The Astrophysical Journal*, 861(1):49.
- Haynes, M. P., Giovanelli, R., Martin, A. M., Hess, K. M., Saintonge, A., Adams, E. A., Hallenbeck, G., Hoffman, G. L., Huang, S., Kent, B. R., et al. (2011). The Arecibo Legacy Fast ALFA survey: the α .40 HI source catalog, its characteristics and their impact on the derivation of the HI mass function. *The astronomical journal*, 142(5):170.
- He, A., An, R., Ivanov, M. M., and Gluscevic, V. (2024). Self-interacting neutrinos in light of large-scale structure data. *Physical Review D*, 109(10):103527.
- He, A., Ivanov, M. M., Bird, S., An, R., and Gluscevic, V. (2025). A Fresh Look at Neutrino Self-Interactions With the Lyman- α Forest: Constraints from EFT and PRIYA. *arXiv preprint arXiv:2503.15592*.
- Heinesen, A. (2023). Bulk flows in the local Universe and the importance of relativistic effects. *Physical Review D*, 108(10):103530.
- Heymans, C., Tröster, T., Asgari, M., Blake, C., Hildebrandt, H., Joachimi, B., Kuijken, K., Lin, C.-A., Sánchez, A. G., Van Den Busch, J. L., et al. (2021). KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints. *Astronomy & Astrophysics*, 646:A140.
- Hinton, S. R. (2016). ChainConsumer. *The Journal of Open Source Software*, 1:00045.
- Hinton, S. R., Howlett, C., and Davis, T. M. (2020). BARRY and the BAO model comparison. *Monthly Notices of the Royal Astronomical Society*, 493(3):4078–4093.

- Hinton, S. R., Kazin, E., Davis, T. M., Blake, C., Brough, S., Colless, M., Couch, W. J., Drinkwater, M. J., Glazebrook, K., Jurek, R. J., et al. (2016). Measuring the 2D baryon acoustic oscillation signal of galaxies in WiggleZ: cosmological constraints. *Monthly Notices of the Royal Astronomical Society*, 464(4):4807–4822.
- Hirata, K. S., Kajita, T., Kifune, T., Kihara, K., Nakahata, M., Nakamura, K., Ohara, S., Oyama, Y., Sato, N., Takita, M., et al. (1989). Observation of B-8 solar neutrinos in the Kamiokande-II detector. *Physical review letters*, 63(1):16.
- Hoffman, Y., Courtois, H. M., and Tully, R. B. (2015). Cosmic bulk flow and the local motion from Cosmicflows-2. *Monthly Notices of the Royal Astronomical Society*, 449(4):4494–4505.
- Hoffman, Y., Nusser, A., Valade, A., Libeskind, N. I., and Tully, R. B. (2021). From Cosmicflows distance moduli to unbiased distances and peculiar velocities. *Monthly Notices of the Royal Astronomical Society*, 505(3):3380–3392.
- Hoffman, Y., Valade, A., Libeskind, N. I., Sorce, J. G., Tully, R. B., Pfeifer, S., Gottlöber, S., and Pomaredé, D. (2024). The large-scale velocity field from the Cosmicflows-4 data. *Monthly Notices of the Royal Astronomical Society*, 527(2):3788–3805.
- Hong, T., Springob, C. M., Staveley-Smith, L., Scrimgeour, M. I., Masters, K. L., Macri, L. M., Koribalski, B. S., Jones, D. H., and Jarrett, T. H. (2014). 2MTF–IV. A bulk flow measurement of the local Universe. *Monthly Notices of the Royal Astronomical Society*, 445(1):402–413.
- Hou, Z., Keisler, R., Knox, L., Millea, M., and Reichardt, C. (2013). How massless neutrinos affect the cosmic microwave background damping tail. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 87(8):083008.
- Howlett, C., Manera, M., and Percival, W. J. (2015). L-PICOLA: A parallel code for fast dark matter simulation. *Astronomy and Computing*, 12:109–126.
- Howlett, C., Said, K., Lucey, J. R., Colless, M., Qin, F., Lai, Y., Tully, R. B., and Davis, T. M. (2022). The Sloan Digital Sky Survey Peculiar Velocity Catalogue. *Monthly Notices of the Royal Astronomical Society*, 515(1):953–976.
- Howlett, C., Staveley-Smith, L., and Blake, C. (2017a). Cosmological forecasts for combined and next-generation peculiar velocity surveys. *Monthly Notices of the Royal Astronomical Society*, 464(3):2517–2544.
- Howlett, C., Staveley-Smith, L., Elahi, P. J., Hong, T., Jarrett, T. H., Jones, D. H., Koribalski, B. S., Macri, L. M., Masters, K. L., and Springob, C. M. (2017b). 2MTF–VI. Measuring the velocity power spectrum. *Monthly Notices of the Royal Astronomical Society*, 471(3):3135–3151.
- Huang, G.-y., Ohlsson, T., and Zhou, S. (2018). Observational constraints on secret neutrino interactions from big bang nucleosynthesis. *Physical Review D*, 97(7):075009.

- Hubble, E. (1929). A relation between distance and radial velocity among extra-galactic nebulae. *Proceedings of the national academy of sciences*, 15(3):168–173.
- Huchtmeier, W. and Richter, O.-G. (1989). HI observations of galaxies in the Virgo cluster of galaxies. II-Global parameters of the galaxies. *Astronomy and Astrophysics*, 210:1–17.
- Humphreys, E., Reid, M. J., Moran, J. M., Greenhill, L. J., and Argon, A. L. (2013). Toward a new geometric distance to the active galaxy NGC 4258. III. Final results and the Hubble constant. *The Astrophysical Journal*, 775(1):13.
- Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95.
- Huterer, D. (202020). Lecture notes michigan cosmology summer school.
- Jackson, J. (1972). A critique of Rees’s theory of primordial gravitational radiation. *Monthly Notices of the Royal Astronomical Society*, 156(1):1P–5P.
- Jarrett, T., Chester, T., Cutri, R., Schneider, S., Skrutskie, M., and Huchra, J. (2000). 2MASS extended source catalog: overview and algorithms. *The Astronomical Journal*, 119(5):2498–2531.
- Jones, D. H., Read, M. A., Saunders, W., Colless, M., Jarrett, T., Parker, Q. A., Fairall, A. P., Mauch, T., Sadler, E. M., Watson, F. G., et al. (2009). The 6dF Galaxy Survey: final redshift release (DR3) and southern large-scale structures. *Monthly Notices of the Royal Astronomical Society*, 399(2):683–698.
- Kaiser, N. (1987). Clustering in real space and in redshift space. *Monthly Notices of the Royal Astronomical Society*, 227(1):1–21.
- Kaiser, N. (1988). Theoretical implications of deviations from Hubble flow. *Monthly Notices of the Royal Astronomical Society*, 231(1):149–167.
- Karim, M. A., Aguilar, J., Ahlen, S., Alam, S., Allen, L., Prieto, C. A., Alves, O., Anand, A., Andrade, U., Armengaud, E., et al. (2025). DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints. *arXiv preprint arXiv:2503.14738*.
- Kashlinsky, A., Atrio-Barandela, F., Kocevski, D., and Ebeling, H. (2008). A measurement of large-scale peculiar velocities of clusters of galaxies: results and cosmological implications. *The Astrophysical Journal*, 686(2):L49–L52.
- Koda, J., Blake, C., Davis, T., Magoulas, C., Springob, C. M., Scrimgeour, M., Johnson, A., Poole, G. B., and Staveley-Smith, L. (2014). Are peculiar velocity surveys competitive as a cosmological probe? *Monthly Notices of the Royal Astronomical Society*, 445(4):4267–4286.
- Kolb, E. W. and Turner, M. S. (1987). Supernova 1987A and the secret interactions of neutrinos. *Physical Review D*, 36(10):2895.

- Kourkchi, E., Tully, R. B., Eftekharzadeh, S., Llop, J., Courtois, H. M., Guinet, D., Dupuy, A., Neill, J. D., Seibert, M., Andrews, M., et al. (2020). Cosmicflows-4: The Catalog of 10,000 Tully–Fisher Distances. *The Astrophysical Journal*, 902(2):145.
- Kraan-Korteweg, R. C., Cluver, M. E., Bilicki, M., Jarrett, T. H., Colless, M., Elagali, A., Böhringer, H., and Chon, G. (2017). Discovery of a supercluster in the Zone of Avoidance in Vela. *Monthly Notices of the Royal Astronomical Society: Letters*, 466(1):L29–L33.
- Kreisch, C. D., Cyr-Racine, F.-Y., and Doré, O. (2020). Neutrino puzzle: Anomalies, interactions, and cosmological tensions. *Physical Review D*, 101(12):123505.
- Kreisch, C. D., Park, M., Calabrese, E., Cyr-Racine, F.-Y., An, R., Bond, J. R., Doré, O., Dunkley, J., Gallardo, P., Gluscevic, V., et al. (2024). Atacama Cosmology Telescope: The persistence of neutrino self-interaction in cosmological measurements. *Physical Review D*, 109(4):043501.
- Lancaster, L., Cyr-Racine, F.-Y., Knox, L., and Pan, Z. (2017). A tale of two modes: Neutrino free-streaming in the early universe. *Journal of Cosmology and Astroparticle Physics*, 2017(07):033.
- Lange, J. U. (2023). Nautilus: boosting Bayesian importance nested sampling with deep learning. *Monthly Notices of the Royal Astronomical Society*, 525(2):3181–3194.
- Lasker, J., Rosell, A. C., Myers, A., Ross, A., Bianchi, D., Hanif, M., Kehoe, R., de Mattia, A., Napolitano, L., Percival, W., et al. (2024). Production of alternate realizations of desi fiber assignment for unbiased clustering measurement in data and simulations. *arXiv preprint arXiv:2404.03006*.
- Leavitt, H. S. and Pickering, E. C. (1912). Periods of 25 Variable Stars in the Small Magellanic Cloud. *Harvard College Observatory Circular*, vol. 173, pp. 1-3, 173:1–3.
- Lee, M. G., Freedman, W. L., and Madore, B. F. (1993). The tip of the red giant branch as a distance indicator for resolved galaxies. *The Astrophysical Journal*, 417:553.
- Lemaître, G. (1927). Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Société Scientifique de Bruxelles*, 47:49–59.
- Lesgourgues, J. and Pastor, S. (2006). Massive neutrinos and cosmology. *Physics Reports*, 429(6):307–379.
- Lesgourgues, J. and Pastor, S. (2012). Neutrino mass from Cosmology. *Advances in High Energy Physics*, 2012(1):608515.
- Lesgourgues, J. and Pastor, S. (2014). Neutrino cosmology and Planck. *New Journal of Physics*, 16(6):065002.
- Lewis, A. (2019). Getdist: a python package for analysing monte carlo samples. *arXiv preprint arXiv:1910.13970*.

- Lewis, A. and Challinor, A. (2011). CAMB: Code for anisotropies in the microwave background. *Astrophysics source code library*, pages ascl-1102.
- Libanore, S., Ghosh, S., Kovetz, E. D., Boddy, K. K., and Raccanelli, A. (2025). Joint 21-cm and CMB Forecasts for Constraining Self-Interacting Massive Neutrinos. *arXiv preprint arXiv:2504.15348*.
- Lyu, K.-F., Stamou, E., and Wang, L.-T. (2021). Self-interacting neutrinos: Solution to Hubble tension versus experimental constraints. *Physical Review D*, 103(1):015004.
- Ma, Y.-Z., Gordon, C., and Feldman, H. A. (2011). Peculiar velocity field: Constraining the tilt of the Universe. *Phys. Rev. D*, 83:103002.
- Ma, Y.-Z. and Pan, J. (2014). An estimation of local bulk flow with the maximum-likelihood method. *Monthly Notices of the Royal Astronomical Society*, 437(2):1996–2004.
- Maksimova, N. A., Garrison, L. H., Eisenstein, D. J., Hadzhiyska, B., Bose, S., and Satterthwaite, T. P. (2021). AbacusSummit: a massive set of high-accuracy, high-resolution N-body simulations. *Monthly Notices of the Royal Astronomical Society*, 508(3):4017–4037.
- Mangano, G. and Serpico, P. D. (2011). A robust upper limit on N_{eff} from BBN, circa 2011. *Physics Letters B*, 701(3):296–299.
- Mather, J. C., Cheng, E., Cottingham, D. A., Eplee Jr, R., Fixsen, D. J., Hewagama, T., Isaacman, R., Jensen, K., Meyer, S. S., Noerdlinger, P. D., et al. (1994). Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument. *The Astrophysical Journal*, 420(2):439–444.
- McKinney, W., van der Walt, S., and Millman, J. (2010). Proceedings of the 9th Python in Science Conference.
- Mena-Fernandez, J., Garcia-Quintero, C., Yuan, S., Hadzhiyska, B., Alves, O., Rashkovetskyi, M., Seo, H., Padmanabhan, N., Nadathur, S., Howlett, C., et al. (2025). HOD-Dependent Systematics for Luminous Red Galaxies in the DESI 2024 BAO Analysis. *Journal of Cosmology and Astroparticle Physics*, 2025(01):133.
- Migkas, K., Pacaud, F., Schellenberger, G., Erler, J., Nguyen-Dang, N., Reiprich, T., Ramos-Ceja, M., and Lovisari, L. (2021). Cosmological implications of the anisotropy of ten galaxy cluster scaling relations. *Astronomy & Astrophysics*, 649:A151.
- Montefalcone, G., Cruz, H. A. G., Munoz, J. B., Kovetz, E. D., and Kamionkowski, M. (2025a). Tracing the neutrino-induced phase shift in the 21-cm spectrum. *arXiv preprint arXiv:2509.03595*.
- Montefalcone, G., Ghosh, S., Boddy, K. K., Ho, D. W. R., and Tsai, Y. (2025b). Directly probing neutrino interactions through cmb phase shift measurements. *arXiv preprint arXiv:2509.20363*.
- Montefalcone, G., Wallisch, B., and Freese, K. (2025c). Free-Streaming Neutrinos and Their Phase Shift in Current and Future CMB Power Spectra. *arXiv preprint arXiv:2501.13788*.

- Moon, J., Valcin, D., Rashkovetskyi, M., Saulder, C., Aguilar, J. N., Ahlen, S., Alam, S., Bailey, S., Baltay, C., Blum, R., et al. (2023). First detection of the BAO signal from early DESI data. *Monthly Notices of the Royal Astronomical Society*, 525(4):5406–5422.
- Noriega, H. E., De-Santiago, J., Garcia-Arroyo, G., Venzor, J., and Pérez-Lorezana, A. (2025). Resonant neutrino self-interactions: insights from the full shape galaxy power spectrum. *arXiv preprint arXiv:2506.07994*.
- Nusser, A. (2014). An inconsistency in the standard maximum likelihood estimation of bulk flows. *The Astrophysical Journal*, 795(1):3.
- Nusser, A. (2016). On methods of estimating cosmological bulk flows. *Monthly Notices of the Royal Astronomical Society*, 455(1):178–184.
- Nusser, A. and Davis, M. (2011). The cosmological bulk flow: consistency with Λ CDM and $z \approx 0$ constraints on σ_8 and γ . *The Astrophysical Journal*, 736(2):93.
- O. Alves et al. (2024). Analytical covariance matrices of DESI galaxy power spectra. *in preparation*.
- Oldengott, I. M., Tram, T., Rampf, C., and Wong, Y. Y. (2017). Interacting neutrinos in cosmology: exact description and constraints. *Journal of Cosmology and Astroparticle Physics*, 2017(11):027.
- Ostriker, J. P. and Steinhardt, P. J. (1995). The observational case for a low-density universe with a non-zero cosmological constant. *Nature*, 377(6550):600–602.
- Paillas, E., Ding, Z., Chen, X., Seo, H., Padmanabhan, N., de Mattia, A., Ross, A., Nadathur, S., Howlett, C., Aguilar, J., et al. (2025). Optimal Reconstruction of Baryon Acoustic Oscillations for DESI 2024. *Journal of Cosmology and Astroparticle physics*, 2025(01):142.
- Pal, S., Samanta, R., and Pal, S. (2025). Exploring neutrino interactions in light of present and upcoming galaxy surveys. *Journal of Cosmology and Astroparticle Physics*, 2025(03):047.
- Park, M., Kreisch, C. D., Dunkley, J., Hadzhiyska, B., and Cyr-Racine, F.-Y. (2019). λ cdm or self-interacting neutrinos: How cmb data can tell the two models apart. *Physical Review D*, 100(6):063524.
- Peery, S., Watkins, R., and Feldman, H. A. (2018). Easily interpretable bulk flows: continuing tension with the standard cosmological model. *Monthly Notices of the Royal Astronomical Society*, 481(1):1368–1375.
- Penzias, A. A. and Wilson, R. W. (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. *The Astrophysical Journal*, 142:419–421.
- Percival, W. J., Friedrich, O., Sellentin, E., and Heavens, A. (2022). Matching Bayesian and frequentist coverage probabilities when using an approximate data covariance matrix. *Monthly Notices of the Royal Astronomical Society*, 510(3):3207–3221.

- Percival, W. J., Ross, A. J., Sánchez, A. G., Samushia, L., Burden, A., Crittenden, R., Cuesta, A. J., Magana, M. V., Manera, M., Beutler, F., et al. (2014). The clustering of Galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: including covariance matrix errors. *Monthly Notices of the Royal Astronomical Society*, 439(3):2531–2541.
- Pérez-Fernández, A., Medina-Varela, L., Ruggeri, R., Vargas-Magaña, M., Seo, H., Padmanabhan, N., Ishak, M., Aguilar, J., Ahlen, S., Alam, S., et al. (2025). Fiducial-Cosmology-dependent systematics for the DESI 2024 BAO Analysis. *Journal of Cosmology and Astroparticle Physics*, 2025(01):144.
- Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P., Castro, P. G., Deustua, S., Fabbro, S., Goobar, A., Groom, D. E., et al. (1999). Measurements of Ω and Λ from 42 high-redshift supernovae. *The Astrophysical Journal*, 517(2):565.
- Philcox, O. H., Eisenstein, D. J., O’Connell, R., and Wiegand, A. (2020). RASCALC: a jackknife approach to estimating single-and multitracer galaxy covariance matrices. *Monthly Notices of the Royal Astronomical Society*, 491(3):3290–3317.
- Phillips, M. M. (1993). The absolute magnitudes of Type IA supernovae. *The Astrophysical Journal*, 413(2):105–108.
- Poudou, A., Simon, T., Montandon, T., Teixeira, E. M., and Poulin, V. (2025). Self-interacting neutrinos in light of recent CMB and LSS data. *arXiv preprint arXiv:2503.10485*.
- Proust, D., Quintana, H., Carrasco, E. R., Reisenegger, A., Slezak, E., Muriel, H., Dünner, R., Sodré Júnior, L., Drinkwater, M. J., Parker, Q. A., et al. (2006). The Shapley Supercluster: the largest matter concentration in the local Universe. *Messenger*, 124:30–31.
- Qin, F. (2021). Gaussianization of peculiar velocities and bulk flow measurement. *Research in Astronomy and Astrophysics*, 21(10):242.
- Qin, F., Howlett, C., and Staveley-Smith, L. (2019). The redshift-space momentum power spectrum–II. Measuring the growth rate from the combined 2MTF and 6dFGSv surveys. *Monthly Notices of the Royal Astronomical Society*, 487(4):5235–5247.
- Qin, F., Howlett, C., Staveley-Smith, L., and Hong, T. (2018). Bulk flow in the combined 2MTF and 6dFGSv surveys. *Monthly Notices of the Royal Astronomical Society*, 477(4):5150–5166.
- Qin, F., Parkinson, D., Howlett, C., and Said, K. (2021). Cosmic Flow Measurement and Mock Sampling Algorithm of Cosmicflows-4 Tully- Fisher Catalog. *The Astrophysical Journal*, 922(1):59.
- Quartin, M., Amendola, L., and Moraes, B. (2022). The 6×2 pt method: supernova velocities meet multiple tracers. *Monthly Notices of the Royal Astronomical Society*, 512(2):2841–2853.
- Radburn-Smith, D. J., Lucey, J., Woudt, P., Kraan-Korteweg, R., and Watson, F. (2006). Structures in the Great Attractor region. *Monthly Notices of the Royal Astronomical Society*, 369(3):1131–1142.

- Raichoor, A., Moustakas, J., Newman, J. A., Karim, T., Ahlen, S., Alam, S., Bailey, S., Brooks, D., Dawson, K., de la Macorra, A., et al. (2023). Target selection and validation of DESI emission line galaxies. *The Astronomical Journal*, 165(3):126.
- Rashkovetskyi, M., Forero-Sánchez, D., de Mattia, A., Eisenstein, D., Padmanabhan, N., Seo, H., Ross, A., Aguilar, J., Ahlen, S., Alves, O., et al. (2025). Semi-analytical covariance matrices for two-point correlation function for DESI 2024 data. *Journal of Cosmology and Astroparticle Physics*, 2025(01):145.
- Renk, J. J., Stöcker, P., Bloor, S., Hotinli, S., Balázs, C., Bringmann, T., Gonzalo, T. E., Handley, W., Hoof, S., Howlett, C., et al. (2021). CosmoBit: A GAMBIT module for computing cosmological observables and likelihoods. *Journal of Cosmology and Astroparticle Physics*, 2021(02):022.
- Rest, A., Scolnic, D., Foley, R., Huber, M., Chornock, R., Narayan, G., Tonry, J., Berger, E., Soderberg, A. M., Stubbs, C. W., et al. (2014). Cosmological Constraints from Measurements of Type Ia Supernovae discovered during the first 1.5 yr of the Pan-STARRS1 Survey. *The Astrophysical Journal*, 795(1):44.
- Riemer-Sørensen, S., Parkinson, D., and Davis, T. M. (2013). What is half a neutrino? Reviewing cosmological constraints on neutrinos and dark radiation. *Publications of the Astronomical Society of Australia*, 30:e029.
- Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., Gilliland, R. L., Hogan, C. J., Jha, S., Kirshner, R. P., et al. (1998). Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116(3):1009.
- Rocher, A., Ruhlmann-Kleider, V., Burtin, E., Yuan, S., de Mattia, A., Ross, A. J., Aguilar, J., Ahlen, S., Alam, S., Bianchi, D., et al. (2023). The desi one-percent survey: exploring the halo occupation distribution of emission line galaxies with abacussummit simulations. *Journal of Cosmology and Astroparticle Physics*, 2023(10):016.
- Rosenberg, E., Gratton, S., and Efstathiou, G. (2022). CMB power spectra and cosmological parameters from Planck PR4 with CamSpec. *Monthly Notices of the Royal Astronomical Society*, 517(3):4620–4636.
- Rubin, D., Aldering, G., Betoule, M., Fruchter, A., Huang, X., Kim, A. G., Lidman, C., Linder, E., Perlmutter, S., Ruiz-Lapuente, P., et al. (2025). Union through UNITY: Cosmology with 2000 SNe Using a Unified Bayesian Framework. *The Astrophysical Journal*, 986(2):231.
- Rubin, V. C. and Ford Jr, W. K. (1970). Rotation of the Andromeda nebula from a spectroscopic survey of emission regions. *Astrophysical Journal*, 159:379.
- Said, K., Colless, M., Magoulas, C., Lucey, J. R., and Hudson, M. J. (2020). Joint analysis of 6dFGS and SDSS peculiar velocities for the growth rate of cosmic structure and tests of gravity. *Monthly Notices of the Royal Astronomical Society*, 497(1):1275–1293.

- Saravanan, M. M., Brinckmann, T., Loverde, M., and Weiner, Z. J. (2025). Abundance and properties of dark radiation from the cosmic microwave background. *arXiv preprint arXiv:2503.04671*.
- Schneider, P. (2006). *Extragalactic astronomy and cosmology: an introduction*, volume 146. Springer.
- Scrimgeour, M. I., Davis, T., Blake, C., James, J. B., Poole, G. B., Staveley-Smith, L., Brough, S., Colless, M., Contreras, C., Couch, W., et al. (2012). The WiggleZ Dark Energy Survey: the transition to large-scale cosmic homogeneity. *Monthly Notices of the Royal Astronomical Society*, 425(1):116–134.
- Scrimgeour, M. I., Davis, T. M., Blake, C., Staveley-Smith, L., Magoulas, C., Springob, C. M., Beutler, F., Colless, M., Johnson, A., Jones, D. H., et al. (2016). The 6dF Galaxy Survey: bulk flows on 50-70 h^{-1} Mpc scales. *Monthly Notices of the Royal Astronomical Society*, 455(1):386–401.
- Sheth, R. K. and Diaferio, A. (2001). Peculiar velocities of galaxies and clusters. *Monthly Notices of the Royal Astronomical Society*, 322(4):901–917.
- Singal, A. K. (2023). Discordance of dipole asymmetries seen in recent large radio surveys with the cosmological principle. *Monthly Notices of the Royal Astronomical Society*, 524(3):3636–3646.
- Singal, A. K. (2025). Solar peculiar motion inferred from dipole anisotropy in redshift distribution of quasars appears to lie along the Galactic Centre direction. *Scientific Reports*, 15(1):31805.
- Sorce, J. G., Gottlöber, S., Yepes, G., Hoffman, Y., Courtois, H. M., Steinmetz, M., Tully, R. B., Pomaredé, D., and Carlesi, E. (2016). Cosmicflows constrained local universe simulations. *Monthly Notices of the Royal Astronomical Society*, 455(2):2078–2090.
- Springob, C. M., Haynes, M. P., Giovanelli, R., and Kent, B. R. (2005). A digital archive of HI 21 centimeter line spectra of optically targeted galaxies. *The Astrophysical Journal Supplement Series*, 160(1):149–162.
- Springob, C. M., Magoulas, C., Colless, M., Mould, J., Erdoğdu, P., Jones, D. H., Lucey, J. R., Campbell, L., and Fluke, C. J. (2014). The 6dF Galaxy Survey: peculiar velocity field and cosmography. *Monthly Notices of the Royal Astronomical Society*, 445(3):2677–2697.
- Spurio Mancini, A., Piras, D., Alsing, J., Joachimi, B., and Hobson, M. P. (2021). CosmoPower: emulating cosmological power spectra for accelerated Bayesian inference from next-generation surveys.
- Stöcker, P., Balázs, C., Bloor, S., Bringmann, T., Gonzalo, T. E., Handley, W., Hotinli, S., Howlett, C., Kahlhoefer, F., Renk, J. J., et al. (2021). Strengthening the bound on the mass of the lightest neutrino with terrestrial and cosmological experiments. *Physical Review D*, 103(12):123508.
- Strauss, M. A. and Willick, J. A. (1995). The density and peculiar velocity fields of nearby galaxies. *Physics reports*, 261(5-6):271–431.

- Tanabashi, M., Hagiwara, K., Hikasa, K., Nakamura, K., Sumino, Y., Takahashi, F., Tanaka, J., Agashe, K., Aielli, G., AMSLER, C., et al. (2018). Review of Particle Physics: particle data groups. *Physical Review D*, 98(3):1–1898.
- Tegmark, M. (1997). Measuring cosmological parameters with galaxy surveys. *Physical Review Letters*, 79(20):3806.
- Tegmark, M., Strauss, M. A., Blanton, M. R., Abazajian, K., Dodelson, S., Sandvik, H., Wang, X., Weinberg, D. H., Zehavi, I., Bahcall, N. A., et al. (2004). Cosmological parameters from SDSS and WMAP. *Physical review D*, 69(10):103501.
- Tonry, J. and Schneider, D. P. (1988). A new technique for measuring extragalactic distances. *Publications of the Astronomical Society of the Pacific*, 100(632):1227.
- Tristram, M., Banday, A. J., Douspis, M., Garrido, X., Górski, K., Henrot-Versillé, S., Hergt, L., Ilić, S., Keskitalo, R., Lagache, G., et al. (2024). Cosmological parameters derived from the final Planck data release (PR4). *Astronomy & Astrophysics*, 682:A37.
- Tully, R. B., Courtois, H. M., and Sorce, J. G. (2016). COSMICFLOWS-3. *The Astronomical Journal*, 152(2):50.
- Tully, R. B. and Fisher, J. R. (1977). A new method of determining distances to galaxies. *Astronomy and Astrophysics*, 54(3):661–673.
- Tully, R. B., Kourkchi, E., Courtois, H. M., Anand, G. S., Blakeslee, J. P., Brout, D., de Jaeger, T., Dupuy, A., Guinet, D., Howlett, C., et al. (2023). Cosmicflows-4. *The Astrophysical Journal*, 944(1):94.
- Vagnozzi, S., Brinckmann, T., Archidiacono, M., Freese, K., Gerbino, M., Lesgourgues, J., and Sprenger, T. (2018). Bias due to neutrinos must not uncorrect'd go. *Journal of Cosmology and Astroparticle Physics*, 2018(09):001.
- Vargas-Magaña, M., Ho, S., Cuesta, A. J., O'Connell, R., Ross, A. J., Eisenstein, D. J., Percival, W. J., Grieb, J. N., Sánchez, A. G., Tinker, J. L., Tojeiro, R., Beutler, F., Chuang, C.-H., Kitaura, F.-S., Prada, F., Rodríguez-Torres, S. A., Rossi, G., Seo, H.-J., Brownstein, J. R., Olmstead, M., and Thomas, D. (2018). The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: theoretical systematics and Baryon Acoustic Oscillations in the galaxy correlation function. *Monthly Notices of the Royal Astronomical Society*, 477(1):1153–1188.
- Venzor, J., Garcia-Arroyo, G., Pérez-Lorenzana, A., and De-Santiago, J. (2022). Massive neutrino self-interactions with a light mediator in cosmology. *Physical Review D*, 105(12):123539.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020a). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods*, 17(3):261–272.

- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors (2020b). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.
- Vitagliano, E., Tamborra, I., and Raffelt, G. (2020). Grand unified neutrino spectrum at Earth: Sources and spectral components. *Reviews of Modern Physics*, 92(4):045006.
- Wadekar, D. and Scoccimarro, R. (2020). Galaxy power spectrum multipoles covariance in perturbation theory. *Physical Review D*, 102(12):123517.
- Wallisch, B. (2019). *Cosmological probes of light relics*. Springer Nature.
- Watkins, R., Allen, T., Bradford, C. J., Ramon Jr, A., Walker, A., Feldman, H. A., Cionitti, R., Al-Shorman, Y., Kourkchi, E., and Tully, R. B. (2023). Analyzing the Large-Scale Bulk Flow using CosmicFlows4: Increasing Tension with the Standard Cosmological Model. *Monthly Notices of the Royal Astronomical Society*, 524(2):1885–1892.
- Watkins, R. and Feldman, H. A. (2015a). An unbiased estimator of peculiar velocity with Gaussian distributed errors for precision cosmology. *Monthly Notices of the Royal Astronomical Society*, 450(2):1868–1873.
- Watkins, R. and Feldman, H. A. (2015b). Large-scale bulk flows from the Cosmicflows-2 catalogue. *Monthly Notices of the Royal Astronomical Society*, 447(1):132–139.
- Watkins, R., Feldman, H. A., and Hudson, M. J. (2009). Consistently large cosmic flows on scales of $100 h^{-1}$ Mpc: a challenge for the standard Λ CDM cosmology. *Monthly Notices of the Royal Astronomical Society*, 392(2):743–756.
- Weinberg, S. (1989). The cosmological constant problem. *Reviews of modern physics*, 61(1):1.
- Wes McKinney (2010). Data Structures for Statistical Computing in Python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 56 – 61.
- Whitford, A. M. (2024). Using the peculiar velocities of galaxies to test cosmology and particle physics. In *Windows on the Universe: 30th Anniversary of the Rencontres du Vietnam*.
- Whitford, A. M., Howlett, C., and Davis, T. M. (2022). Using peculiar velocity surveys to constrain neutrino masses. *Monthly Notices of the Royal Astronomical Society*, 513(1):345–362.
- Whitford, A. M., Howlett, C., and Davis, T. M. (2023). Evaluating bulk flow estimators for CosmicFlows–4 measurements. *Monthly Notices of the Royal Astronomical Society*, 526(2):3051–3071.

- Whitford, A. M., Rivera-Morales, H., Howlett, C., Vargas-Magaña, M., Fromenteau, S., Davis, T. M., Pérez-Fernández, A., de Mattia, A., Ahlen, S., Bianchi, D., et al. (2025). Constraining the phase shift of relativistic species in DESI BAOs. *Monthly Notices of the Royal Astronomical Society*, 538(3):1980–2000.
- Willick, J. A., Courteau, S., Faber, S., Burstein, D., Dekel, A., and Strauss, M. A. (1997). Homogeneous velocity-distance data for peculiar velocity analysis. III. The Mark III catalog of galaxy peculiar velocities. *The Astrophysical Journal Supplement Series*, 109(2):333.
- Wright, A. H., Stölzner, B., Asgari, M., Bilicki, M., Giblin, B., Heymans, C., Hildebrandt, H., Hoekstra, H., Joachimi, B., Kuijken, K., et al. (2025). KiDS-Legacy: Cosmological constraints from cosmic shear with the complete Kilo-Degree Survey. *arXiv preprint arXiv:2503.19441*.
- Yèche, C., Palanque-Delabrouille, N., Baur, J., and Des Bourboux, H. D. M. (2017). Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100. *Journal of Cosmology and Astroparticle Physics*, 2017(06):047.
- York, D. G., Adelman, J., Anderson Jr, J. E., Anderson, S. F., Annis, J., Bahcall, N. A., Bakken, J., Barkhouser, R., Bastian, S., Berman, E., et al. (2000). The sloan digital sky survey: Technical summary. *The Astronomical Journal*, 120(3):1579–1587.
- Yoshii, Y. and Peterson, B. A. (1995). Interpretation of the faint galaxy number counts in the k band. *Astrophysical Journal, Part 1 (ISSN 0004-637X)*, vol. 444, no. 1, p. 15-20, 444:15–20.
- Yuan, S., Zhang, H., Ross, A. J., Donald-McCann, J., Hadzhiyska, B., Wechsler, R. H., Zheng, Z., Alam, S., Gonzalez-Perez, V., Aguilar, J. N., et al. (2024). The DESI one-per cent survey: exploring the halo occupation distribution of luminous red galaxies and quasi-stellar objects with AbacusSummit. *Monthly Notices of the Royal Astronomical Society*, 530(1):947–965.
- Zaroubi, S., Hoffman, Y., Fisher, K., and Lahav, O. (1995). Wiener reconstruction of the large scale structure. *Astrophysical Journal*, 449:446–459.
- Zheng, Y., Zhang, P., and Jing, Y. (2015). Determination of the large scale volume weighted halo velocity bias in simulations. *Physical Review D*, 91(12):123512.
- Zhou, R., Dey, B., Newman, J. A., Eisenstein, D. J., Dawson, K., Bailey, S., Berti, A., Guy, J., Lan, T.-W., Zou, H., et al. (2023). Target selection and validation of DESI luminous red galaxies. *The Astronomical Journal*, 165(2):58.
- Zwicky, F. (1933). Die rotverschiebung von extragalaktischen nebeln. *Helvetica Physica Acta*, 6:110–127.
- Zwicky, F. (1937). On the Masses of Nebulae and of Clusters of Nebulae. *The Astrophysical Journal*, 86:217.

Zyla, P., Barnett, R., Beringer, J., Dahl, O., Dwyer, D., Groom, D., Lin, C.-J., Lugovsky, K., Pianori, E., et al. (2020). Review of particle physics. *Progress of Theoretical and Experimental Physics*, 2020(8):083C01.

Appendix A

Appendices

A.1 Linear theory velocity covariance matrix

The linear theory covariance matrix of galaxy PVs G_{mn} in a survey can be calculated given information about their coordinates $\{\theta, \phi, r(z)\}$ (here $r(z)$ denotes the dimensionful comoving distance to the galaxy) as follows; the matrix element corresponding to the galaxies labelled by m and n is given by

$$G_{mn} = \frac{a^2 H(a)^2 f^2}{2\pi^2} \int_{V_k} P_{mm}(k) W_{mn}(k) dk, \quad (\text{A.1})$$

where $P_{mm}(k)$ is the matter density power spectrum, $W_{mn}(k)$ is a function that depends on the angle between the galaxies m and n and f is the growth rate of matter in the universe, $f = \frac{d \ln D(a)}{da}$. a is the scalefactor and $D(a)$ is the linear growth factor. $W_{mn}(k)$ can be written as (Ma et al., 2011)

$$W_{mn}(k) = \frac{1}{3} \cos(\alpha) (j_0(kA) - 2j_2(kA)) + \frac{1}{A^2} j_2(kA) r_m r_n \sin^2(\alpha), \quad (\text{A.2})$$

where j_l is the spherical bessel function, α is the angle between \mathbf{r}_m and \mathbf{r}_n , the position vectors for the galaxies. A is defined as $A = \sqrt{r_m^2 + r_n^2 - 2r_m r_n \cos(\alpha)}$.

A.2 Biases in galaxy surveys

Malmquist bias occurs when the distance to a galaxy is estimated from a distance-indicator such as the Fundamental Plane or the Tully-Fisher relation; to avoid it, this bias requires information about the line-of-sight density distribution of galaxies (Strauss and Willick, 1995). The bias here arises due to the fact that the selection function of any data generally means galaxies are more likely to sit at greater distances where there is more volume in the survey, than at lower distances. Due to this, there will be some galaxies with measured properties allowing them to lie on the Fundamental Plane or Tully-Fisher relation by ‘scattering in’, particularly if the galaxies lie towards the edge of the galaxy survey volume. These galaxies will have their true distances underestimated.

Selection bias occurs when, for example, you have used the measured properties of a selection of galaxies to calibrate a relation such as the Tully-Fisher relation or Fundamental Plane, but you have a

biased view of the distribution of these properties due to selection effects such as a magnitude limit (Strauss and Willick, 1995). This would result in a biased calibration of a scaling relation. In this case, the issue is caused by the fact that you are more likely to see only the brighter objects when collecting data, and the lower probability of seeing fainter objects if they are more distant must be accounted for mathematically.

A.3 Validating the anisotropic pipeline by comparison to Baumann et al (2019)

We show the fits produced by following the isotropic BAO pipeline by Baumann et al. (2019) with β_ϕ allowed to vary. Fits were produced in configuration space, and compared to the results from the *desilike* pipeline (following an anisotropic fitting scheme) and an extended version of the approach of Baumann et al. (2019) for an anisotropic fitting scheme. The pipeline of Baumann et al. (2019) in configuration space uses a minimizer to find the best fit. The errors on the fits are evaluated by calculating the χ^2 on a grid in $\{\alpha, \beta_\phi\}$ parameters to map out the likelihood surface. The anisotropic extension of this scheme calculates the best fit in a similar way but extends the parameter space to $\{\alpha, \beta_\phi, \alpha_{AP}\}$. The approach by Baumann et al. (2019) uses the nonlinear model for the correlation function described in Vargas-Magaña et al. (2018). However in their work, the Fourier-space fits use the same modelling as this work, following the methodology of Beutler et al. (2017). The results of the fits from each code are compared in Figure A.1.

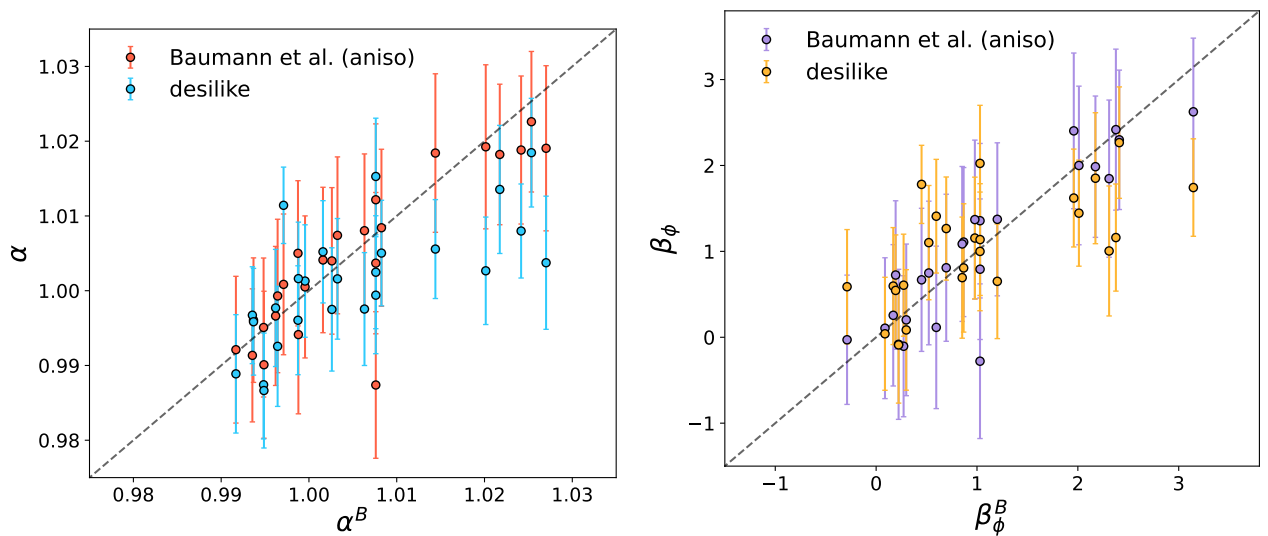


Figure A.1: Results of fitting α (left panel) and β_ϕ (right panel) for 25 mock realisations using the Pipeline of Baumann et al. (2019), compared to an extended version of the Baumann pipeline for anisotropic fitting, and additionally comparing to the *desilike* pipeline.

A.4 Sampling methods in cosmology

Markov Chain Monte Carlo (MCMC)

Markov Chain Monte Carlo (MCMC) refers to a group of algorithms that allow one to estimate a probability distribution for a parameter (or more usually, many parameters). For example, the *Metropolis Hastings algorithm* is a type of MCMC algorithm. In MCMC algorithms, one samples a probability distribution by proposing new steps (positions in a multi-dimensional parameter space) in a chain of steps in a manner that depends only on the previous step; this is called the *Markov Property*. After proposing many steps, the chain of steps (which are saved along the way) gives one an estimate of the probability distribution of the parameters of interest.

In the Metropolis Hastings algorithm, the approach to proposing steps normally involves computing the likelihood of some randomly generated point; depending on whether the likelihood is greater in the new proposed position or the current position, the new step will or won't be accepted.

Nested sampling

Nested sampling refers to a separate set of algorithms that vary from MCMC algorithms, and can be more useful for distributions with multi-modal likelihoods (there may be more than one maximum in the probability distribution of each parameter). Unlike MCMC, which estimates a quantity proportional to the likelihood (or the posterior distribution) of interest, nested sampling estimates the posterior and also the Bayesian Evidence. This is essentially the marginalized likelihood (obtained by integrating the conditional probability of observing the data given some model over all free parameters of the model) and allows one to estimate how much a particular model is favoured by the data; this can be useful to know.

Nested sampling algorithms generally start by taking randomly selected samples from the parameter space, and then identify the points with the highest likelihood in the parameter space. With these points, one then estimates a boundary around the set of points with this likelihood (a contour). In next step, new points are sampled from within this contour, and this step is done continuously to find the next contour with higher likelihood. In this way, the posterior is estimated in 'nested' shells, thus allowing one to integrate the prior volume over iso-contours of approximately constant likelihood.

Importance sampling

Importance sampling is a method that allows one to combine constraints on a parameter (or set of parameters) from various probability distributions. For a set of different samples, one normally takes the least constraining set of samples, and then assigns higher importance (higher weight) to values in this set based on the sets of more constraining samples available. It allows one to reduce the variance in an estimate of a distribution.

A.5 Additional comparisons from Barry

In Figure A.2 we show an additional fit for the mock mean of the first-generation control-variate mocks for ELGs with $z = 1.1$, to compare the power spectrum and correlation function fits for the two broadband methodologies, and also to compare the results to the standard BAO fitting procedure.

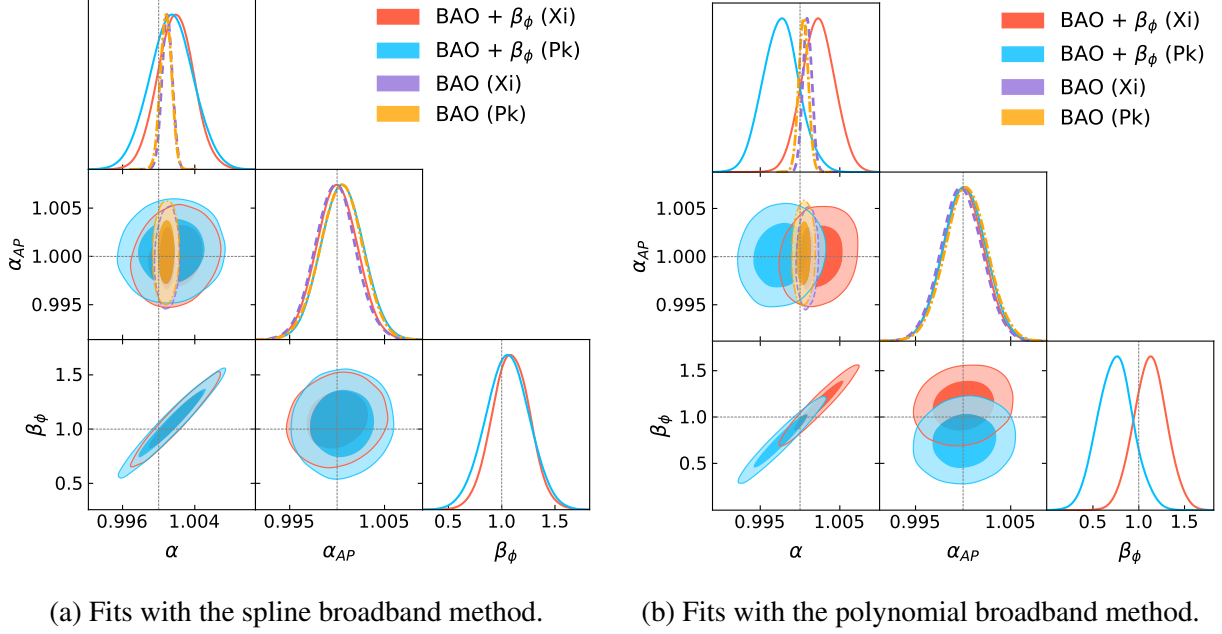


Figure A.2: Comparison of the power spectrum and correlation function, for standard BAO fits with CV ELGs (first-generation snapshots at $z = 1.1$) and for fits when the phase shift is varied. Here however, we show the difference when the spline broadband methodology is used.

A.6 Fisher information and forecasts

The Fisher information for a parameter x evaluated at $x = \theta$ from some observable X is given by negative of the expectation of variance of the *score*, which is defined as the derivative of the log-likelihood with respect to θ . This gives

$$F(x)|_{x=\theta} = - \left\langle \left(\frac{d \log \mathcal{L}(X|x)}{dx} \right)^2 \right\rangle_{|x=\theta}. \quad (\text{A.3})$$

For multiple parameters, $\mathbf{x} = \{x_0, x_1 \dots x_n\}$ this can be extended into a matrix format to obtain the Fisher matrix F_{ij} of size $n \times n$. One can write more generally

$$\mathbf{F}_{ij}|_{\mathbf{x}=\theta} = - \left\langle \frac{d^2 \log \mathcal{L}(X|x)}{dx_i dx_j} \right\rangle_{|x=\theta}. \quad (\text{A.4})$$

The Cramér-Rao Bound says that the variance on the parameter x_i is bounded by the Fisher information,

$$\sigma_i^2 \geq (\mathbf{F})_{ii}^{-1}. \quad (\text{A.5})$$

For a multivariate-Gaussian distribution of our parameters \mathbf{x} with observable(s) with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, we can write down the Fisher Information as

$$F_{ij} = \frac{d\boldsymbol{\mu}}{d\theta_i} \boldsymbol{\Sigma}^{-1} \frac{d\boldsymbol{\mu}}{d\theta_j} + \frac{1}{2} \text{tr} \left(\boldsymbol{\Sigma}^{-1} \frac{d\boldsymbol{\Sigma}}{d\theta_i} \boldsymbol{\Sigma}^{-1} \frac{d\boldsymbol{\Sigma}}{d\theta_j} \right). \quad (\text{A.6})$$

A.7 The effect of N_{eff} on the phase shift

The effect of N_{eff} on the phase shift amplitude is demonstrated in Figure A.3; for neutrinos with universal self-interactions, a larger value of N_{eff} increases the impact of the phase shift induced by self-interactions, relative to a model with no self-interactions. The goal here is to show that it is possible to separate the impacts of N_{eff} and G_{eff} on the phase shift. However, I note that in general for a fixed value of G_{eff} , the epoch at which the neutrinos decouple may slightly vary when N_{eff} is varied. Therefore, the approach in Montefalcone et al. (2025b) was to study the phase shift as a function of the decoupling redshift of the neutrinos rather than G_{eff} , although both approaches are valid.

In Figure A.4 we show the impact of N_{eff} on the CMB temperature power spectrum which has a measurable phase shift (both in the CMB and BAO signal) independent of G_{eff} , due to neutrino free-streaming. This has been previously studied Bashinsky and Seljak (2004); Green and Ridgway (2020); Baumann et al. (2018) and measured in Follin et al. (2015); Baumann et al. (2019); Whitford et al. (2025); Montefalcone et al. (2025c).

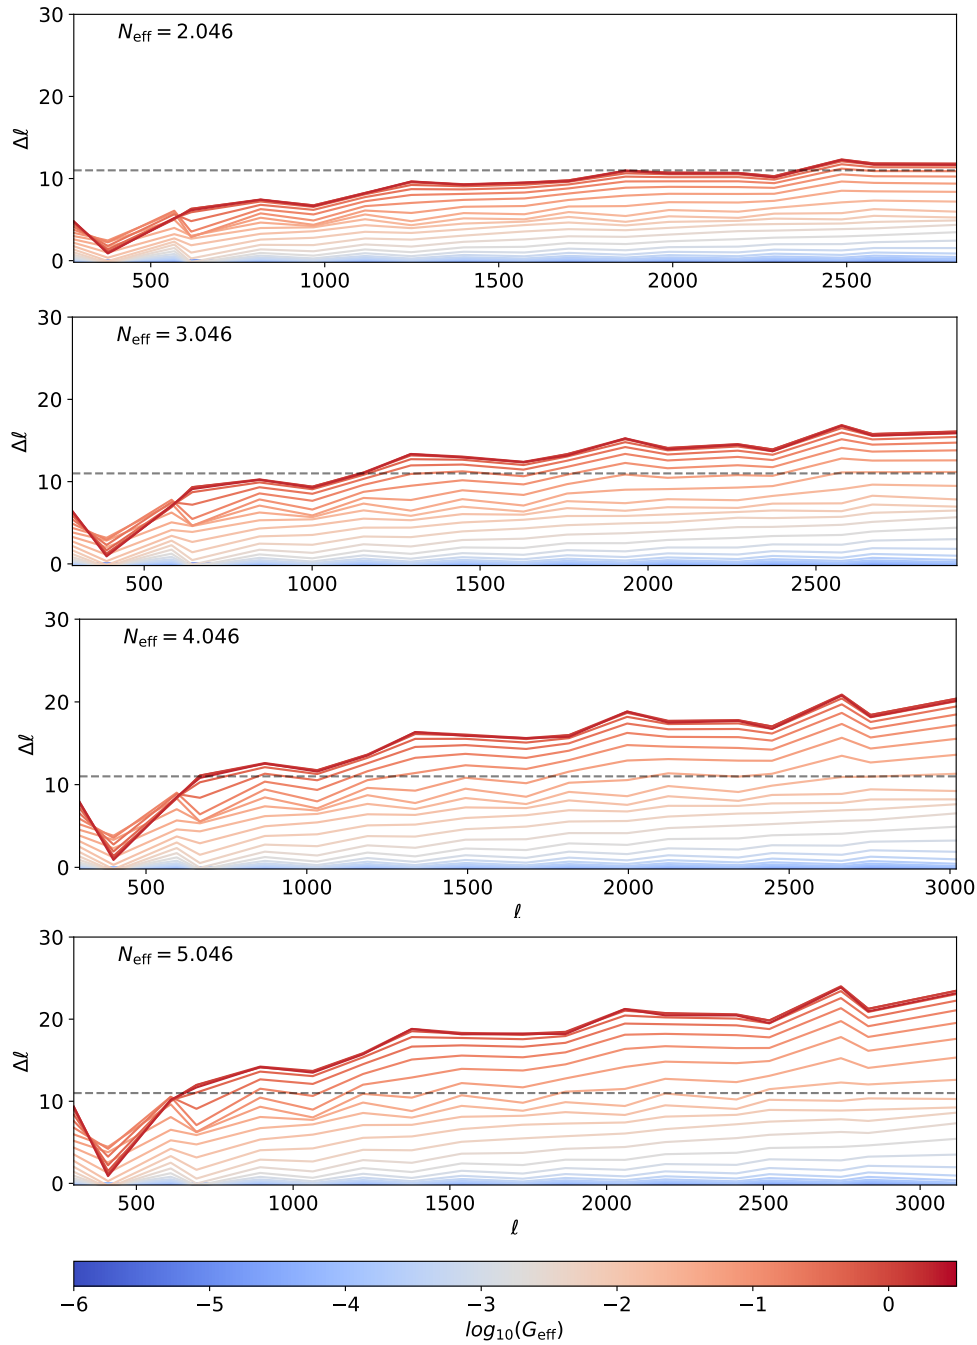


Figure A.3: The phase shift $\Delta\ell$, demonstrated in the CMB temperature-temperature power spectrum, when varying $\log_{10}(G_{\text{eff}})$. The various panels show the phase shift for different fixed values of N_{eff} , to demonstrate the impact that additional light relics have on the amplitude of the phase shift, and to show that it is possible to separate the dependence of the phase shift signal on N_{eff} and $\log_{10}(G_{\text{eff}})$.

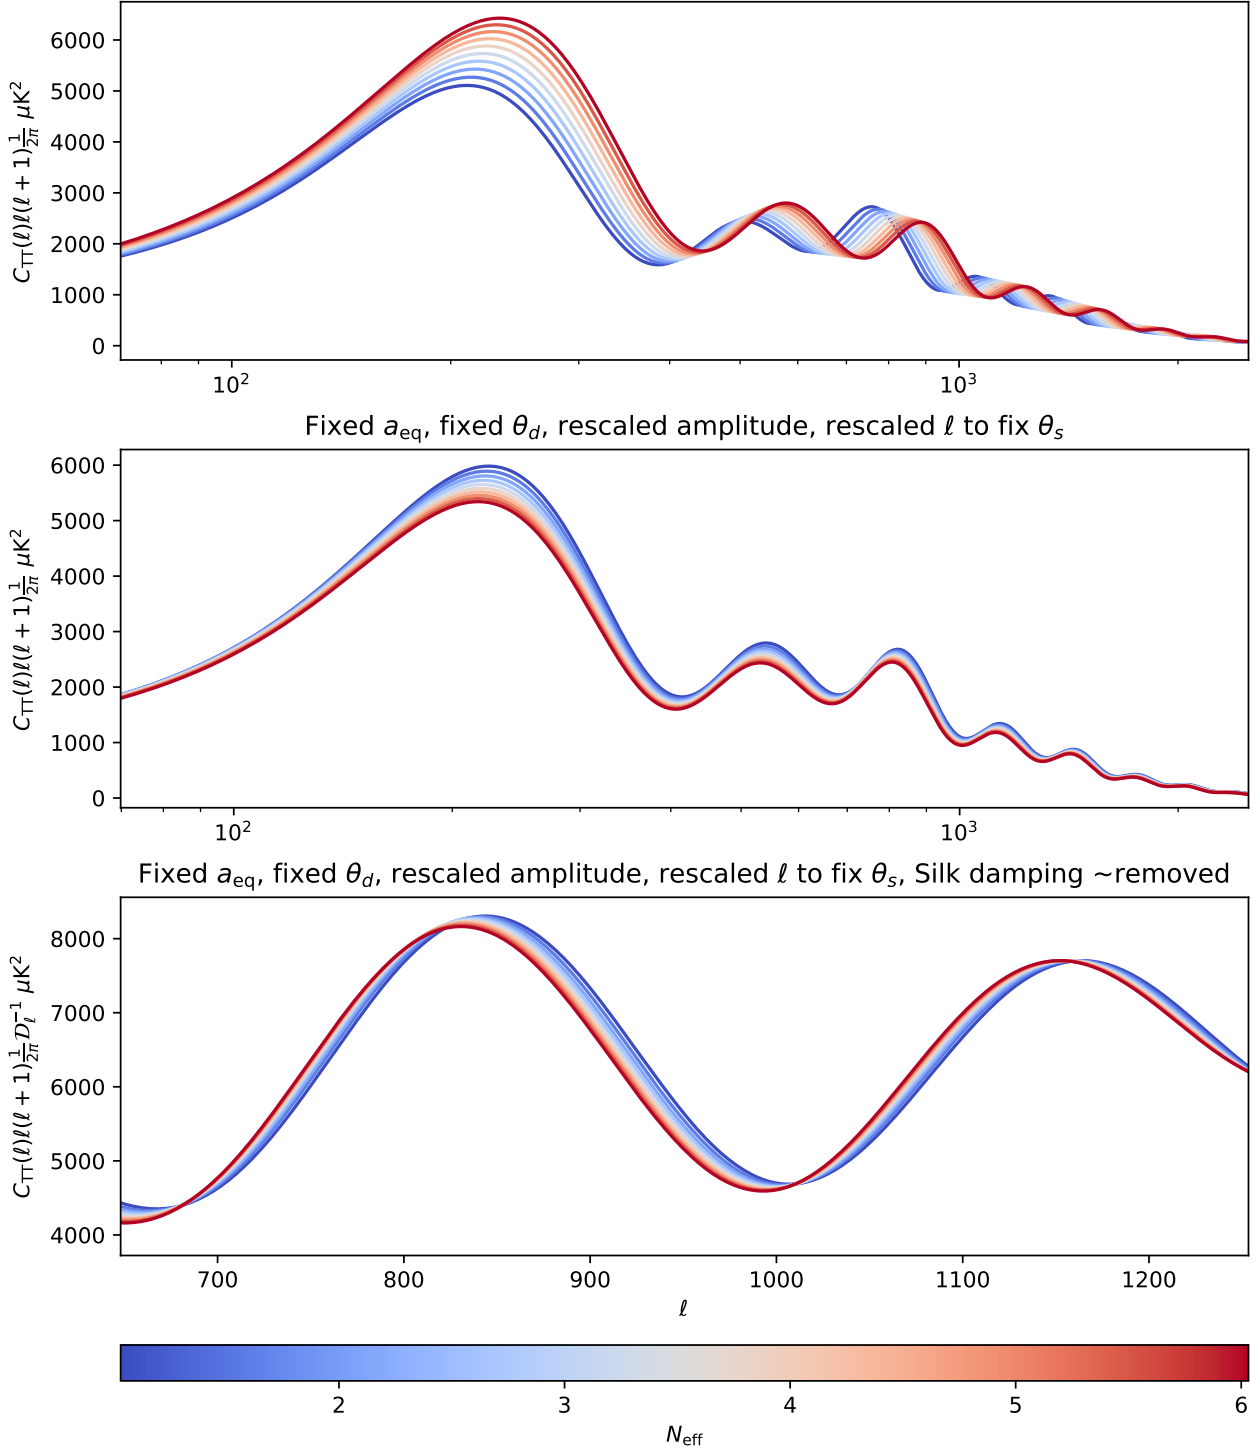


Figure A.4: The neutrino-induced phase shift in the CMB. The top panel is the CMB with changing N_{eff} . In the middle panel, the phase shift has been isolated by removing the impact of changes to a_{eq} , θ_s , θ_D and changes to the amplitude of the CMB. The lowest panel shows the same but has been undamped (approximately) by multiplying by $\sim e^{0.25(\ell\theta_D)^{1.35}}$ following the approach taken in Follin et al. (2015); Montefalcone et al. (2025c) to approximately remove the impact of silk damping; this exponential damping is cosmology dependent. This shows the isolated phase shift for standard model neutrinos.

