

© 2017 by Arka Banerjee. All rights reserved.

COSMOLOGICAL SIGNATURES OF FUNDAMENTAL PHYSICS

BY

ARKA BANERJEE

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2017

Urbana, Illinois

Doctoral Committee:

Assistant Professor Jessie Shelton, Chair
Assistant Professor Neal Dalal, Director of Research
Professor Charles Gammie
Professor Gilbert Holder

Abstract

This thesis deals with the study of the cosmological signatures of certain aspects of fundamental physics, and how cosmological observables can be used to constrain properties of fundamental particles. Over the past decades, increasingly precise measurements in cosmology have become powerful probes of fundamental physics - for example, the inference of dark matter and dark energy from cosmological observations remain the most significant evidence of new physics beyond the Standard Model of particle physics. Another example is the cosmological scalar-to-tensor ratio, which can potentially differentiate different models of inflation and other early Universe theories.

Determining the absolute mass scale of neutrinos is an interesting problem in particle physics, and can shed light on the mass generation mechanism for neutrinos, which, in turn, can tell us about physics beyond the Standard Model. To fully exploit the signatures of massive neutrinos on cosmological observables, one needs to perform accurate simulations. In this thesis, we explore a new method for performing neutrino simulations, which overcome the shortcomings of previous methods which were employed. From these simulations, we identify an observable which is very sensitive to the neutrino mass - the clustering of cosmological voids on large scales. We also forecast how well the neutrino masses and thermal “dark radiation” models can be constrained in future cosmological surveys using their effect on various observables in these surveys, such as the clustering of galaxies, galaxy-galaxy lensing, and cosmic shear.

Cosmological observables can also be used to constrain the properties of Dark Matter itself. While Dark Matter has traditionally been considered a collisionless fluid, there has been recent interest in self-interactions of dark matter. We consider a special form of self-interactions in this thesis - where the interactions are elastic but anisotropic. We develop the formalism and methods to simulate these interactions, and study the signatures of these interactions on the properties of dark matter halos.

Acknowledgments

This work would not have been possible without the help and support of a number of people. First, I'd like to thank my advisor, Neal Dalal, for his support, encouragement, and guidance. Working with him has been a privilege, and I have truly learnt a lot, both academically and personally, from our interactions.

I would like to thank my collaborators, in particular, Bhuvnesh Jain, Andrey Kravtsov, and Jessie Shelton, for their guidance, time, and enthusiasm. I would also like to thank Charles Gammie, Nick Gnedin, Ryan Foley, Mani Chandra, Ben Ryan, and Ricky Chue for stimulating discussions at various stages of my graduate work.

I would like to thank both the Department of Physics, and the Department of Astronomy, for supporting me over the course of my graduate school. I must thank Mel and Wendy from the Physics Office, and Jeri, Judy and Rebecca in the Astronomy Office, for their help with all things official.

Finally, I would like to thank Susmita, for always being there - a constant source of support and companionship, and without whose presence, the last few years would have seemed much longer.

Table of Contents

List of Tables	vi
List of Figures	vii
List of Abbreviations	xiii
Chapter 1 Introduction	1
1.1 Neutrinos	1
1.2 Self-Interacting Dark Matter	6
Chapter 2 Simulating nonlinear cosmological structure formation with massive neutrinos	9
2.1 Introduction	9
2.2 Equations of motion	13
2.2.1 Collisionless Boltzmann equation	13
2.2.2 Poisson equation	14
2.2.3 Obtaining the Boltzmann moment equations	14
2.2.4 SPH equations	16
2.3 Closing the hierarchy	16
2.3.1 Early evolution	17
2.3.2 Late evolution	20
2.3.3 Smoothing in SPH simulations	21
2.4 Integration techniques	22
2.4.1 Particle time integration	22
2.4.2 Hydrodynamics	23
2.4.3 Gravity	25
2.5 Code tests	26
2.5.1 Plummer sphere advection	27
2.5.2 Collision of Plummer spheres	29
2.5.3 Linear growth rate	31
2.5.4 CDM comparison run	35
2.5.5 Simulating Warm Dark Matter cosmologies	40
2.6 Neutrino Simulations	45
2.7 Conclusions	58

Chapter 3	Tests of Neutrino and Dark Radiation Models from Galaxy and CMB	
surveys		60
3.1	Introduction	61
3.2	Method	64
3.2.1	Weak lensing in galaxy surveys	64
3.2.2	CMB lensing	66
3.2.3	Model differentiation	66
3.3	Survey parameters and systematics	67
3.3.1	LSST	67
3.3.2	CMB lensing	71
3.4	Neutrino mass constraints	72
3.4.1	LSST constraints on neutrino mass	72
3.4.2	Comparison to CMB Stage 4 lensing	77
3.4.3	Combination of LSST and CMB S4	79
3.5	Dark radiation	80
3.6	NNaturalness	83
3.7	Summary and Discussion	87
Chapter 4	Structure formation in Anisotropic Self-Interacting Dark Matter cos-	
mologies		92
4.1	Introduction	92
4.2	Formalism for anisotropic SIDM	94
4.3	Toy Problem	97
4.3.1	Setup	97
4.3.2	Deterministic Drag Method	98
4.3.3	Random Angle Method	99
4.3.4	Results	101
4.4	Cosmological simulations with anisotropic SIDM	104
4.4.1	Implementation in L-Gadget	104
4.4.2	Results	105
4.5	Summary and Discussion	108
References		111

List of Tables

3.1	List of redshift bin centers and l_{\max} used for each redshift bin for the LSST survey. .	68
3.2	Forecasts of $1\text{-}\sigma$ constraints on the neutrino mass, and the dark energy equation of state w at LSST for different number of lens and source bins, and different priors. The N_{eff} prior used was 0.03.	73

List of Figures

1.1	Damping of the linear matter power spectrum on small scales due to the presence of massive neutrinos. We plot the ratio of the power spectrum in the massive neutrino cosmology to the power spectrum in a cosmology with only massless neutrinos. Higher neutrino masses imply higher Ω_ν and therefore more damping.	4
2.1	Density profile of an isolated advecting Plummer sphere with $b = 5$. The solid black lines indicate the initial density profile of particles (top panel) and fluid (bottom panel). The dotted lines in the top panel shows the final density profile of particles from 4 runs - a run where particles source gravity (black), a run where the fluid sources gravity and the particles are evolved using the Kick-Drift-Kick method (red), a fluid-sourced gravity run where the particles are evolved using the Drift-Kick-Drift method (blue), and a fluid-sourced gravity run where we estimated the velocity dispersion instead of the bulk velocity (green) - after 20 dynamical times. The dotted lines in the bottom panel shows the fluid profile from the same four runs after 20 dynamical times. While the other methods all agree well with the run in which particles source gravity, the Drift-Kick-Drift method is not suitable for coupling with our hydro scheme.	28
2.2	Density profiles for advecting Plummer spheres after 20 dynamical times for $b = 5$ (black), $b = 4$ (red), $b = 3$ (blue) and $b = 2$ (green). The solid lines represent the profiles from runs where gravity is sourced by particles. The dotted lines are from runs where the fluid sources gravity. The top panel plots the particle density profile, normalized by the maximum value of particle density for a given value of b from the particle based runs. The bottom panel plots the fluid density profiles with the same normalization. The profiles have been staggered for clarity.	30
2.3	Particle (top panel) and fluid (bottom panel) density profiles for an advecting Plummer sphere with $b = 5$, with different number of test particles (N_p) used to sample the initial velocity distribution and estimate bulk velocity. We find that for this size, $N_p \sim 10^5$ is sufficient to follow the dynamics correctly.	31
2.4	Evolution of the density profiles of two colliding Plummer spheres. The solid curves represent the particle density (upper panels) and fluid density (lower panels) for three different runs - gravity from particles (black), gravity from fluid Superbee (SB) flux limiter (red), and another fluid gravity run with Monotonized Central (MC) flux limiter (blue). The dotted curve in each panel shows the initial configuration. The left panels represent the densities when the Plummer spheres have passed through each other once. The right panels represent densities when the two spheres have merged.	32

2.5	Comparison of the growth of the power spectrum of neutrinos ($m_\nu = 0.1$ eV). The top panel plots the comparison of the power spectra from the hybrid simulations (solid lines) to the outputs from CLASS (dot-dashed lines) . We plot the initial power spectra at $z = 199$, as well as at redshifts of $z = 22.53$ and $z = 0$. We also plot the power spectra (dashed lines) at the later redshifts from an N-body treatment of neutrinos, which are dominated by shot noise at small scales. In the bottom panel we plot the ratio of the power spectra from the hybrid simulations to the linear theory results at $z = 199$ (teal), $z = 23.53$ (black) and $z = 0$ (red).	34
2.6	Linear growth of the power spectrum for a WDM particle of mass 200 eV. The top panel plots the power spectrum at initial redshift $z = 99$ and at $z = 60.54$ and $z = 27.57$ from the hybrid simulations and from linear theory. The bottom panel plots the ratio of the power spectra from the simulations and from linear theory at $z = 99$ (teal), $z = 60.54$ (black) and $z = 27.57$ (red). At large scales in the box, the ratio is different from 1 due to sample variance.	35
2.7	Final power spectrum at $z = 0$ for CDM runs. In the top panel, we plot the results a normal N-body run (black) as well as from the particles (red) and fluid (blue) in our hybrid simulations. In the bottom panel we plot the ratio of the particle power spectrum from the hybrid simulations to the power spectrum from the N-body run.	37
2.8	Mass function from the N-body run (black) and the hybrid Eulerian method (red) for a Λ CDM cosmology. The red curve shows a deficit of halos at all mass scales, including near the high end of the mass function. In contrast, the hybrid Lagrangian (SPH) method is formally identical to N-body for CDM cosmologies.	38
2.9	A zoom-in of a halo at $z = 0$ from the CDM comparison simulations - the left panel shows the density field from a standard N-body simulation, while the right panel shows the density field from our hybrid method. In the N-body case, the central object has already merged, whereas in the hybrid method, the two objects are still distinct.	40
2.10	Final power spectrum at $z = 0$ for WDM species with mass 200 eV. In the top panel, we plot the results from a hot N-body run (black), a cold N-body run (green), as well as our hybrid SPH simulation (red). The results from the SPH run and the cold start N-body run are virtually indistinguishable. On large scales in the box, the hot start N-body run agrees with the other two, but there are differences on small scales. In the bottom panel we plot the ratio of the particle power spectrum from the hybrid SPH run to the cold N-body power spectrum (red) to show that these methods match each other down to the smallest scales resolved by this simulation box, even though the initial conditions for the SPH run was the same as that for the hot start N-body run. We also plot the ratio of the power spectra from the hot run to the cold run (black) to show their difference on scales affected by the streaming of particles.	42
2.11	Mass function for WDM species with $m = 200$ eV. We plot results from a hot N-body run (black), a cold N-body run (blue), and our hybrid SPH method (red). The brown dashed line represents M_{damp} defined in Eqn. 2.42. The hot run shows the presence of many halos below M_{damp} , while the other runs do not show these smaller halos.	43

2.12	Mass function for WDM species with $m = 60$ eV. We plot results from a hot N-body run (black), a cold N-body run (blue), and our hybrid SPH method (red). The brown dashed line represents M_{damp} . Below M_{damp} the hot run shows a number of spurious halos. At even smaller scales, both the cold start and the hybrid SPH run start showing “beads on a string” artifacts. The black dash-dotted line is an extrapolation from the flat part of the halo mass function.	44
2.13	A slice through the density field in the simulation boxes at $z = 0$ for two neutrino masses - 0.5 eV and 3 eV. The simulation box side length is $175 h^{-1}$ Mpc. The left panel shows the CDM density field, the middle panel shows the density field for the 0.5 eV neutrino, while the right panel shows the density field for the 3 eV neutrino. For the lighter neutrino, on large scales, the density traces the underlying cosmic web structure laid down by the CDM component, but is much more diffuse on small scales. For the heavier neutrino, the density field is less diffuse and follows the CDM density more closely down to smaller scales.	46
2.14	Relative damping of the matter power spectrum in the presence of a massive neutrino with $m = 0.1$ eV with realistic energy density ($\Omega_\nu \approx 0.0022$), compared to the CDM-only prediction. The linear theory prediction is plotted with the dashed red while the dotted blue curve represents the HALOFIT prediction. The solid black curve is the result of our hybrid simulation, while the solid green curve is from a simulation treating neutrinos as another set of N-body particles.	47
2.15	Absolute value of the bias as a function of scale, averaged from 8 runs for $m_\nu = 0.1$ eV and $f_\nu = 0.1$. The voids were defined using the underdensity in the total matter field. The threshold for void definition was set at -0.7 . We compare the results from our simulations (black) to the other existing methods of treating the neutrinos as a linear fluid (green) and treating neutrinos as a set of particles with a different mass in N-body simulations (red). Our method and the N-body method yield results which match to within error bars, but show a strong scale dependence. The linear method shows an even stronger scale dependence. We also plot (in blue) the bias for voids defined in exactly the same manner from 8 CDM-only runs in for which the final power spectrum matches the final CDM power spectrum for the runs including neutrinos.	49
2.16	Comparison of the ratio of Eulerian biases b_ν and b_{CDM} at $f_\nu = 0.1$, measured in simulations (black circles) vs. predictions from the spherical expansion model (red triangles), as a function of the threshold floor used in the void definition. The bias ratio diminishes for decreasing void thresholds, a trend quantitatively predicted by the spherical expansion model. The error bars represent the errors on the best fit values from the simulations.	51
2.17	Behavior of the void bias from the hybrid simulations when the voids are selected using the CDM field only, and then split on the basis of the enclosed total matter overdensity for $m_\nu = 0.1$ eV and $f_\nu = 0.1$. The black curve represents $b_{\text{tot}}(k)$ for voids selected using the CDM field only with overdensity threshold of -0.7 . The red curve represents $b_{\text{tot}}(k)$ for the subsample of voids whose enclosed total matter overdensity was lower than the median enclosed total matter overdensity in the above sample. Similarly, the green curve shows $b_{\text{tot}}(k)$ for the subsample whose enclosed total matter overdensity is higher than the median.	52

2.18	Effect of f_ν on the scale-dependent bias using our method for a fixed neutrino mass of $m_\nu = 0.1$ eV. We compare the case where $f_\nu = 0.1$ (black) to the case where $f_\nu = 0.05$ (red). The dashed lines show the predictions from the spherical expansion model.	53
2.19	Effect of neutrino mass on the scale dependent bias. We compare the bias when the mass of the neutrino is 0.1 eV (black) to the case when the mass is 0.8 eV (red) with $f_\nu = 0.1$ in both cases. The red dot-dashed line gives the free streaming length for the 0.8eV neutrino. The free streaming length of the lighter neutrino lies outside (to the left of) this plot. The scale dependent bias is evident below the streaming length of each species. The dashed lines show predictions from the spherical expansion model.	54
2.20	Scale dependent bias from the spherical expansion model for inverted hierarchy (black curve) and normal hierarchy (red curve). In the inverted hierarchy, $\sum m_\nu = 0.12$ eV, $f_\nu \sim 0.0087$ and $f_{\text{CDM}} \sim 0.9913$. In the normal hierarchy, $\sum m_\nu = 0.06$ eV, $f_\nu \sim 0.0043$ and $f_{\text{CDM}} \sim 0.9957$. Voids are defined as regions enclosing an under-density -0.7 in both cases. We assume that $b_{\text{CDM}} = -2$ for reference.	57
3.1	Expected source galaxy counts (per square arcmin) as a function of redshift in the LSST survey is plotted with the black curve. The solid red lines indicate the photometric redshifts z_{min} and z_{max} of the 6 lens bins used in our analysis. The dashed blue lines represent the z_{min} and z_{max} for the 6 source bins. The blue lines have been slightly displaced for clarity.	68
3.2	Comparison of the C_l^{gg} and $C_l^{\kappa\kappa}$ power spectra at different redshifts to the shot noise and shape noise levels at those redshifts. The solid lines plot the power spectrum, while the dashed line of the same color plots the shot noise or shape noise for that power spectrum at that redshift expected at LSST. The dot-dashed lines represent the value of l_{max} for each redshift bin. The spectra are sample variance dominated out to $l \sim 400$	69
3.3	Noise level N_l^{dd} in the deflection power spectrum C_l^{dd} from CMB Stage 4 lensing for the assumed survey parameters. The lensing signal is sample variance dominated out to $l \sim 1000$	71
3.4	Marginalized probability distributions for the sum of neutrino masses (left), and dark energy equation of state w (right), from LSST. The blue, red and black curves correspond to $N_L = N_S = 1, 4, 6$ respectively. Increasing the number of redshift helps extract more tomographic information, but this gain saturates as the individual bins become too thin.	73
3.5	1- σ (solid lines) and 2- σ (dashed lines) confidence intervals on 2-d subspace of the parameters $\sum m_\nu$ and w from LSST. The red curves are the results from 4 source bin and 4 lens bin analysis. The black curves represent the results when 6 source bins and 6 lens bins are used.	74
3.6	Degradation of the 1- σ constraints on the parameters as a function of the prior assumed on m_i , the shear calibration of the source bins at LSST. The constraints change very weakly with the value of the prior, especially for w	76

3.7	Marginalized probability distributions for the sum of neutrino masses, and dark energy equation of state w from CMB Stage 4 lensing (red curves) compared to LSST (black curve). The solid red curve corresponds to $f_{\text{sky}} = 0.75$, the dashed red curve corresponds to $f_{\text{sky}} = 0.5$ and the dot-dashed curve corresponds to $f_{\text{sky}} = 0.25$ for CMB Stage 4 experiments. While neutrino mass constraints from the two experiments are similar, w can be constrained more tightly by the LSST survey.	79
3.8	Marginalized probability distributions for the sum of neutrino masses $\sum m_\nu$, and dark energy equation of state w for joint LSST and CMB lensing analysis (green curve) compared to those coming from LSST only (black curve) and CMB lensing only (red curve). The constraints improve somewhat marginally compared to the LSST only results.	81
3.9	Damping in the power spectrum caused by a Standard Model neutrino (black curve) with mass $m_\nu = 0.06$ eV compared to the damping due a fermionic dark radiation (red curve) particle producing $\Delta N_{\text{eff}} = 0.15$ at CMB and having the same mass, $m_{\text{DR}} = 0.06$ eV. The shapes and the amplitudes of the damping are different for the two particles, implying that dark radiation and Standard Model neutrino masses can be measured simultaneously in surveys like LSST.	82
3.10	Ratio of the power spectrum at $z = 0$ for different NNaturalness scenarios (labeled by m_ϕ) compared to the Standard Model normal hierarchy scenario. The presence of extra free streaming species damps the power spectrum on the free streaming scale of the lightest of the extra species. On linear scales ($k > 0.1$ h/Mpc), the damping is below a percent level.	85
4.1	Elastic collision of two dark matter particles in the center of mass frame. For anisotropic SIDM, the scattering angle θ is drawn from a distribution where small values of θ are favored over large values of θ	94
4.2	Comparison of the analytic PDF (black dashed line) and the cut-off PDF (red solid line) for the choice of $\epsilon = 0.005$. The normalization is adjusted to fit within the plot range. We also plot the CDF (blue dot-dashed line) for cut-off PDF. We find that the momentum transfer cross section σ_T computed with the cut-off PDF is within 0.5% of the analytic value for this choice of ϵ	100
4.3	Results of our toy model for anisotropic SIDM interactions. We plot the average parallel velocity of all subhalo particles in the left panel, and the average of the square of the perpendicular velocity in the right panel. We look at three different cases, $N_{\text{subhalo}} = 50$, $N_{\text{subhalo}} = 100$, and $N_{\text{subhalo}} = 200$, and 10 realizations of each case. The 10 realizations are slightly displaced along the x-axis to help with the visualization. The analytic expectations are plotted using the blue lines. The results from the deterministic drag method are represented using black points, and the random angle method using the green dots.	102
4.4	Stacked density profiles for halos of mass $1 - 4 \times 10^{14} h^{-1} M_\odot$ for different implementations of self-interactions. The random angle method changes the density profile more than the deterministic drag method.	106
4.5	Slope calculated from the stacked density profiles of halos with mass $1 - 4 \times 10^{14} h^{-1} M_\odot$ for different implementations of self-interactions. It is clear that the random angle method changes the slope of the density profile more than the deterministic drag method, but there is no effect on the splashback radius of dark matter particles.	107

4.6 Stacked subhalo density profiles for halos of mass $1 - 4 \times 10^{14} h^{-1} M_{\odot}$. Only subhalos with mass greater than $10^{12} h^{-1} M_{\odot}$ were used to obtain this profile. In this case, we find that the deterministic drag method shows a larger effect for the same cross section. This implies a higher subhalo evaporation rate for this method compared to the random angle method. 108

List of Abbreviations

DM	Dark Matter.
CDM	Cold Dark Matter.
WDM	Warm Dark Matter.
SIDM	Self Interacting Dark Matter.
CMB	Cosmic Microwave Background.
SPH	Smoothed Particle Hydrodynamics
PM	Particle-Mesh
PT	Perturbation Theory
CIC	Cloud-in-Cell
DES	Dark Energy Survey
DESI	Dark Energy Spectroscopic Instrument
LSST	Large Synoptic Survey Telescope
WFIRST	Wide Field Infrared Survey Telescope
PDF	Probability Distribution Function
CDF	Cumulative Distribution Function
NDBD	Neutrinoless Double Beta Decay

Chapter 1

Introduction

Starting with Zwicky’s observations of the velocity dispersion of galaxies in the Coma cluster [1], observations on cosmological scales have had profound implications for fundamental physics. It is now well-established that roughly 27% of the energy content of the Universe is in the form of Dark Matter - matter which does not interact through the electromagnetic force, but does participate in the gravitational evolution of the Universe. Even in the simplest scenario that Dark Matter is made up of a single species of fundamental particles, one needs to invoke particles which are beyond the Standard Model of particle physics. Many models have been proposed to extend the Standard Model of physics in a way that produces a Dark Matter candidate with the correct mass and interaction strength to match the current observed density of Dark Matter on cosmological scales.

Similarly, other observations, including observations of Type I supernovae in distant galaxies [2], have now confirmed that even larger fraction of the energy of the Universe, roughly 68%, is made of something named as “Dark Energy”. While Dark Matter behaves like ordinary matter in the gravitational context, Dark Energy acts like a fluid with negative pressure on cosmological scales, leading to an accelerating expansion of the Universe. Explaining the nature of Dark Energy is another challenge that cosmological observations have thrown up to fundamental physics.

1.1 Neutrinos

Even though neutrinos are part of the Standard Model of physics, we now know that neutrinos are massive, and that their mass eigenstates are different from their flavor eigenstates [3, 4, 5]. Though we do not yet know the absolute mass scale of the different mass eigenstates, we know that even the maximum allowed mass from terrestrial experiments, ~ 1 eV, is much smaller than the mass of

all other Standard Model particles. The next lightest Standard Model particle, the electron, has a mass of 0.511 MeV, nearly six orders of magnitude larger! Various mass mechanisms, including the famous see-saw mechanism, have been proposed to explain the light mass of the neutrinos, see [6, 7] for reviews. Most of these models assume extensions to the Standard Model, and so, pinning down the exact mass scales of neutrinos could possibly provide us information about new physics beyond the Standard Model.

Ongoing terrestrial experiments are looking to pin down the mass scale more accurately than previous experiments. One approach is experiments like the KATRIN (KARlsruhe TRITium Neutrino) experiment [8], which looks to constrain the mass of the electron anti-neutrino by measuring the energy spectrum of the emitted electron in a tritium beta decay down to a precision of roughly 0.2 eV. Other terrestrial experiments include searches for neutrinoless double beta decays (NDBD), such as the SNO+ experiment [9]. Such beta decays can only occur if the neutrinos are Majorana particles, that is, the neutrino and the anti-neutrino are the same particle. If neutrinos are indeed Majorana particles, then the event rate of neutrinoless double beta decays depend on sum of the neutrino masses. Observations of NDBD events can, therefore, be used to place constraints on the sum of neutrino masses, and therefore the absolute mass scales of the eigenstates.

Massive neutrinos can also produce signatures on cosmological observables. Along with all other fundamental particles, neutrinos existed in thermal equilibrium in the early Universe. Since neutrinos are fermions, their initial distribution is given by a thermal Fermi-Dirac distribution:

$$f_\nu(p) \propto \frac{1}{\exp\left[\frac{E-\mu}{T}\right] + 1}, \quad (1.1)$$

where μ is the chemical potential. Given the low masses of neutrinos, they remained relativistic until redshift of $z \sim 100$ depending on the exact mass. As the Universe cooled, heavier particles dropped out of equilibrium. As time evolves, the Universe also expands, and physical densities of both radiation and matter goes down. Neutrinos interact with other species through the weak interaction, and based on the strength of the weak interaction, and the expansion rate of the Universe, it has been shown that neutrinos remained in equilibrium with photons, electrons and positrons till the time the average temperature of the Universe was roughly ~ 2 MeV. Below this

temperature the number of interactions were no longer sufficient to keep neutrinos in thermal equilibrium. After decoupling, the neutrinos continue to evolve with the thermal Fermi-Dirac distribution at the time of decoupling frozen in, but with redshifting momenta $p \propto a^{-1}$, where a is the scale factor. As long as the neutrinos are relativistic, and the chemical potential is negligible, this distribution is equivalent to a Fermi-Dirac distribution with redshifting temperature $T_\nu \propto a^{-1}$.

Just below the temperature at which the neutrinos decouple, the temperature of the Universe falls below the limit at which photons have sufficient energy to produce electron-positron pairs. All remaining positrons end up annihilating with nearby electrons, but since only the photons are in thermal equilibrium at this point, the temperature of the photons increases with respect to the temperature of the neutrinos. Simple arguments, coming from the conservation of entropy density yields

$$T_\nu = T_\gamma \left(\frac{4}{11} \right)^{1/3}. \quad (1.2)$$

This means that the temperature of the Cosmic Neutrino Background is roughly 0.71 of the temperature of the Cosmic Microwave Background (CMB). Since the neutrinos follow a thermal spectrum, the temperature also determines the number density. Given the temperature, the number density of each of the three species is $n_{\nu_\alpha} \sim 113 \text{ cm}^{-3}$, comparable to the CMB number density. This means that neutrinos are the second most abundant particles in the Universe, after photons.

For massive neutrinos, we can calculate their overall contribution to the energy density of the Universe today, given the number density calculated above. It is given by the following formula [10], in terms of the sum of the masses of the three neutrino species:

$$\Omega_\nu = \frac{\sum m_\nu}{94 h^2 \text{ eV}}. \quad (1.3)$$

Except at very low redshifts, clustering of neutrinos is minimal given their large thermal velocities. Therefore, while they contribute to the overall expansion rate of the universe by a factor proportional to Ω_ν , they do not contribute as sources for the peculiar gravitational potential, which sources the clustering of dark matter and baryons. This manifests itself through the damping of the matter power spectrum on scales smaller than the free streaming scales of the neutrino particles. This is shown in Fig. 1.1, where we plot the ratio of the matter power spectra in three different




Figure 1.1: Damping of the linear matter power spectrum on small scales due to the presence of massive neutrinos. We plot the ratio of the power spectrum in the massive neutrino cosmology to the power spectrum in a cosmology with only massless neutrinos. Higher neutrino masses imply higher Ω_ν and therefore more damping.

cosmologies with massive neutrinos to the power spectrum in a cosmology with massless neutrinos. All power spectra here have been calculated using linear perturbation theory, and is, therefore, more accurate on larger scales. As expected, for more massive neutrinos, Ω_ν is higher, and the amplitude of damping is also larger. This characteristic damping of the power spectrum is one of the strongest signatures of neutrinos on large scale structure of the Universe. However, detecting the amplitude of this damping in cosmological observations is complicated by the fact that it is degenerate with some of the other parameters of the Λ -CDM paradigm, such as the tilt of the primordial power spectrum, n_s , and the optical depth to reionization, τ .

To study the effects of neutrinos on the power spectrum on smaller scales, linear perturbation theory is no longer adequate. We know that on small scales, Cold Dark Matter and baryons cluster very strongly. Accurate simulations of Λ -CDM cosmologies have been performed using N-body techniques for more than three decades now. To form a fully consistent picture of the effect of massive neutrinos, we therefore need to include neutrinos in some way in these cosmological simulations. A number of methods have been suggested for this, but each of these methods has its own drawbacks. For example, extending the N-body technique to include neutrino particles suffers from shot noise issues due to poor sampling of the phase space. In §2, we present a completely

new method for running full, non-linear simulations of cosmologies with massive neutrinos which gets around these problems. Also in §2, we identify a new cosmological observable which is very sensitive to the neutrino mass - the clustering of voids defined using the matter field on large scales. We show that quantity used to quantify this clustering - called *bias* - becomes a scale-dependent quantity. For standard Λ -CDM cosmologies, this bias is expected to be independent of scale at large scales. Therefore, unlike the damping of the power spectrum on small scales, this effect of neutrinos on large scale structures is much harder to produce within standard cosmologies, and could provide a new method for constraining the neutrino mass.

Massive neutrinos also leave their imprints on observables which will be measured with increasingly greater accuracy in surveys like LSST [11] and CMB Stage 4 experiments [12]. In the case of LSST, which is a wide photometric galaxy and lensing survey covering nearly half the sky, one of the major science goals is to measure various two-point correlations such as the galaxy auto-clustering, galaxy-galaxy lensing, and cosmic shear. Galaxy autocorrelations, as the name suggests, measures the clustering of distant galaxies on the sky in various redshift bins. As light from distant galaxies passes through intervening matter before reaching our detectors, the shapes of the galaxies are slightly distorted through the effect of gravitational lensing. These changes are usually small, and falls in the *weak lensing* regime. Galaxy-galaxy lensing refers the measurement of this shape distortion of distant source galaxies around a foreground lens galaxy. In other words, this measures the two-point correlation of galaxy position and the weak-lensing shear. Similarly, cosmic shear measures the shape distortions of source galaxies about random points on the sky. This is therefore, the two-point correlation of the weak lensing shear with itself.

All of these correlations, along with the lensing of the CMB due to intervening matter between the last scattering surface and us, depend on the shape and amplitude of the matter power spectrum at different redshifts. Massive neutrinos affect both these properties of the matter power spectrum. Neutrinos damp the power spectrum on small scales, thereby changing the shape. The energy content of the neutrinos Ω_ν affects the expansion rate of the Universe, as well as the growth rate of perturbations, thereby changing the amplitude of the power spectrum at different redshifts. In §3, we forecast how well LSST and CMB Stage 4 experiments will be able to constrain the neutrino mass by exploiting the observables mentioned above. Even for the scenario where the sum of

the neutrino masses is at the minimum value allowed by neutrino oscillation experiments, these upcoming experiments can possibly detect it at a statistical significance of roughly 2σ .

In §3, we also explore whether these surveys can also constrain the properties of other light particles that have been predicted by various extensions of particle physics beyond the Standard Model. One example is thermal dark radiation, which is a generic prediction of a number of different models. These extra light species have similar effects on the matter power spectrum - that is, they affect the shape and the growth, but since they have different masses and temperatures as compared to the Standard Model neutrinos, they produce slightly different signatures on the various observables at LSST.

1.2 Self-Interacting Dark Matter

On scales larger than a few Megaparsecs, the standard non-interacting Cold Dark Matter scenario is an excellent fit to all current observations, see [13, 14] for example. This scenario assumes that the interactions of dark matter with baryons, as well as interactions of dark matter with itself has negligible effect on the evolution and structure formation in the Universe. The effect of Dark Matter is only mediated through gravity - the overall energy density in Dark Matter participates in determining the time evolution of the scale factor of the Universe, while density fluctuations in Dark Matter sources the growth of structure through the peculiar gravitational potential.

However, on small scales, predictions of Λ -CDM cosmologies, where Cold Dark Matter has only gravitational interactions, seem to be in contradiction with a number of observations. One of the most studied of these problems is the missing satellite problem, and the related too-big-to-fail problem [15, 16, 17, 18]. This refers to the observation that we find far fewer satellite galaxies for the Milky Way than is predicted for a halo with the halo mass of the Milky Way from simulations. Another well-known problem is the core-cusp problem [19, 20]. This refers to the fact that dark matter halos in simulations have cuspy inner profiles whereas observations find that the density profiles near the centers of galaxies are cored - that is, much flatter than the theoretical predictions.

Motivated by these small scale problems of the non-interacting Dark Matter scenario, [21] put forward the idea that Dark Matter with strong self interactions could potentially solve these conflicts with observations, and this idea was further explored in [22], with the help of simulations.

These studies specifically looked only at short-range elastic interactions of Dark Matter particles with themselves. Since then, there has been growing interest in exploring self-interactions of Dark Matter, as well as exploring more complex models of the dark sector with multiple species and interactions.

While short range elastic interactions have been explored in detail with high-resolution simulations (see [23, 24, 25, 26, 27], for example), recent observations have spurred interest in a slightly different form of self-interactions for Dark Matter. The first of these observations is that the “splashback radius” of halos, defined as the boundary of the multi-streaming, virialized region, detected in observations [28, 29] is significantly different from theoretical predictions [30, 31]. The splashback radius in these studies was defined by looking at the stacked density profiles of satellite galaxies around clusters. Since satellite galaxies resides inside subhalos, these observation could be explained if Dark Matter subhalos feel an extra *drag* force as they fall into the dark matter halo of the host.

The second set of observations is the separation of the light profile from the lensing mass profiles for galaxies falling into cluster [32, 33]. These offsets are difficult to explain if both stars and Dark Matter are essentially collisionless - in that scenario, both would follow essentially the same trajectory during their infall phase. Once again, if Dark Matter has extra interactions, this is no longer true, and it is possible for the stars and Dark Matter to have somewhat different trajectories.

Analytical studies and toy model implementations [34, 28] suggest if Dark Matter has interactions which are elastic, but frequent, and with an anisotropic differential cross-section, it would produce observational signatures consistent with the above. In this picture, Dark Matter particles suffer frequent collisions, but most collisions have a small scattering angle, which corresponds to a small momentum transfer per collision. Averaged over the full trajectory, this form of self-interaction produces a drag effect on subhalos as they travel through the density profile of the host halo. These interactions can also lead to subhalo evaporation, similar to what happens in the case of isotropic interactions. However, the evaporation rate is much slower than in the isotropic case. It is therefore, very interesting to study the full signatures of such interactions with the help of cosmological simulations.

In §4, we study the effects of anisotropic self-interactions of Dark Matter on the profiles of halos. We set down the formalism for describing this sort of self-interactions, and then implement them in full cosmological simulations. We present the results from these simulations, and discuss future extensions and improvements to the presented work.

Chapter 2

Simulating nonlinear cosmological structure formation with massive neutrinos

We present a new method for simulating cosmologies that contain massive particles with thermal free streaming motion, such as massive neutrinos or warm/hot dark matter. This method combines particle and fluid descriptions of the thermal species to eliminate the shot noise known to plague conventional N-body simulations. We describe this method in detail, along with results for a number of test cases to validate our method, and check its range of applicability. Using this method, we demonstrate that massive neutrinos can produce a significant scale-dependence in the large-scale biasing of deep voids in the matter field. We show that this scale-dependence may be quantitatively understood using an extremely simple spherical expansion model which reproduces the behavior of the void bias for different neutrino parameters. ¹

2.1 Introduction

One of the central tenets of the standard cosmological model is that structure observed in the present universe formed via gravitational evolution of initially linear density perturbations which arose in the primordial universe [35, 10]. The initially linear perturbations responsible for producing anisotropies in the cosmic microwave background [36] are expected to eventually develop into nonlinear cosmological structures like halos and filaments at low redshift [37]. Predicting the nonlinear evolution of cosmic structure has been a numerically challenging problem for many years. The method of choice for computing structure formation has been N-body simulation. The accuracy and efficiency of this method have been well established for standard Λ CDM cosmologies [38, 39]. These simulations have been used to study many aspects of structure formation, including the mass

¹This chapter was previously published in the Journal of Cosmology and Astroparticle Physics as A. Banerjee and N. Dalal, *Simulating nonlinear cosmological structure formation with massive neutrinos*, Journal of Cosmology and Astroparticle Physics 11 (Nov., 2016) 015, [1606.06167]. This chapter matches the published version.

functions [40, 41, 42] and profiles [43, 44, 45, 17, 46, 18, 47] of halos and voids, and their large scale clustering [e.g. 48, 49, 41, 50].

Although N-body simulations have been enormously successful in describing the evolution of Λ CDM cosmologies, they have not fared as well for studying certain other cosmologies, especially universes containing massive particles with large thermal velocities in the initial phase space distribution of the species. The large thermal velocities of these particles generate spurious structures in the density field due to discreteness effects inherent in N-body methods, as the particles stream large distances in random directions. At scales below this streaming scale, the density distribution of the simulation particles follows a Poisson shot noise distribution, instead of the correct, physical density. The power spectrum at large wavenumber would then be given by $P(k) \sim \bar{n}^{-1}$, where \bar{n} is the mean density of particles in the simulation. For typical values of \bar{n} , the shot noise power spectrum can completely dominate the physical power spectrum (which is damped on these scales by the streaming motions), leading to spurious structures forming everywhere in the simulation volume. Therefore, improvements in numerical methods are essential to develop a reliable and consistent method for studying structure formation in such cosmologies.

An example of a cosmology with fast moving massive particles is our own universe, which is known to contain massive neutrinos. Neutrinos are among the most abundant particles in the universe, comparable to photons in terms of their number density [10]. Oscillation experiments have clearly established that at least two neutrino states are massive, and these experiments have also placed tight bounds on the mass differences between the three mass eigenstates of neutrinos [3, 4, 5, 51, 52]. Even though the absolute masses are not yet known, the mass splittings imply that at least one of the species is as heavy as ~ 0.06 eV. Given the current cosmic neutrino background temperature $T_\nu \sim 1.68 \times 10^{-4}$ eV [10], this means that this species is highly non-relativistic today, and can therefore gravitationally cluster. The effects of massive neutrinos on large-scale structure are well-understood on large scales and at early times, when fluctuations are in the linear regime [53, 54, 55]. However, at late times and smaller scales, density fluctuations become non-linear, rendering linear perturbation theory calculations invalid.

Different groups have put forward different approaches to attacking this problem. [56] and [57] have proposed treating the neutrinos as a linear fluid on a grid coupled to the fully non-linear,

N-body evolution of the cold dark matter (CDM) particles. Similarly, [58] suggested solving the linearized Boltzmann equations for neutrinos coupled to the the N-body evolution of CDM particles. By their very nature, these quasi-linear methods break down when the overdensities in the neutrino fluid approach $\mathcal{O}(1)$. At late times, massive neutrinos can become cold enough to be captured into the deepest potential wells in the simulations - the largest halos. The overdensities of neutrinos in these halos can be significantly larger than $\mathcal{O}(1)$, limiting the validity of linearized methods at late times.

Another method which has been proposed is to treat the neutrinos as a different species of particles with different mass than CDM particles in a normal N-body simulation at all redshifts [59]. A related method is to use the linear treatment for neutrinos at early redshifts, but then to switch to N-body treatment once perturbations in the neutrino fluid become non-linear [60]. While traditional N-body methods suffer from the shot noise issue mentioned above, these authors argue that since the neutrinos constitute a small part of the energy budget ($\Omega_\nu \ll 1$), their effects are always subdominant to CDM, and the shot noise effect is not strong enough to significantly alter structure formation. On the other hand, since shot noise can be significant compared to the real physical clustering of neutrinos in the simulation, this means that calculations of neutrino effects on the power spectrum can suffer from large *fractional* errors, even if the absolute errors on the power spectrum are small thanks to the small mass fraction in neutrinos. Given that the point of performing such simulations is the precise calculation of neutrino effects, the fractional error on neutrino effects may be a more relevant metric than the absolute error on the total power spectrum. A number of interesting results have been found using these methods [61, 62, 63, 64, 65, 66, 67, 68]. However, with future cosmological surveys expected to reduce error bars on multiple observables significantly, an improvement in the accuracy of neutrino simulations may now be warranted.

Similarly, besides neutrinos, dark matter particles can also have significant random thermal velocities, depending on the DM temperature at freezeout. It has been suggested that Warm Dark Matter [69] can alleviate certain small-scale problems present for CDM universes, like the core-cusp problem of halo profiles [19, 20] and the missing satellite problem [15, 16, 17, 18]. If DM particles have a finite thermal velocity dispersion, then those particles will randomly stream a finite distance, and this random streaming acts to suppress structure on scales below the free-streaming length.

The streaming length of WDM particles is much smaller than the streaming length of neutrinos [69, 70], and if the simulations are initialized in a way that the particles have random thermal velocities, the scales affected by shot noise would be smaller than in neutrino simulations. We refer to these sort of simulations as “hot start” simulations. However, since the WDM is the dominant component in these simulations, any amount of shot noise coming from the thermal velocities is sufficient to seed the formation of spurious structures. “Hot start” N-body simulations are therefore not an accurate method for studying such WDM cosmologies.

To get around this problem, the “cold start” method has typically been used in the literature [71, 72]. In this method, the random thermal velocities of particles are *not* included in the simulation initial conditions. To account for the damping of the power spectrum, this method initializes using the linear power spectrum for the WDM species at redshift $z = 0$, scaled back to the starting redshift of the simulation using CDM growth factors. This method therefore necessarily does not capture the spatial dependence and time dependence of the growth of structure, but it *does* eliminate artifacts arising from random thermal velocities. Nonetheless, the cold start method suffers from its own artifacts, like the “beads on a string” phenomenon [72]. [73] have proposed an interesting method to eliminate these “beads on a string” artifacts and other spurious structures from WDM simulations. Their method employs a tetrahedral tessellation of the 6-dimensional phase space of simulation particles to follow the evolution of the WDM densities. However, it is not yet clear if this method is accurate inside collapsed and virialized regions such as halos or subhalos, whose abundance will likely provide stringent constraints on WDM models in upcoming years [e.g. 74, 75]. Alternatively, [76] have proposed that adaptive softening of the gravitational force can help to suppress spurious structure found in cold start simulations of WDM cosmologies.

In this chapter, we present a novel method to simulate cosmologies with hot particles which is valid at all redshifts, in both linear and non-linear regimes. This method makes use of both particle techniques from N-body simulations as well as hydrodynamic techniques from fluid simulations. This chapter is arranged as follows. In §2.2, we derive the relevant equations of motion for hot species in an expanding universe. We describe the implementation of our method at early and late redshifts in §2.3. In §2.4 we discuss the time integration techniques for the particles in our simulations, as well the hydrodynamic scheme we implement. In §2.5, we discuss a number of tests,

in both the linear as well as the non-linear regime to validate our code. We use this method to show a novel effect in void biasing with neutrinos in §2.6. Finally, we list our main conclusions and directions for future work in §2.7.

2.2 Equations of motion

In this section, we review the equations of motion governing hot species in an expanding universe. We take moments of the Boltzmann equation to derive effective fluid equations that allow us to evolve the hot species using hydrodynamic methods.

2.2.1 Collisionless Boltzmann equation

For the applications we are interested in, we consider the WDM and neutrinos to be essentially collisionless, and our starting point will be the collisionless Boltzmann equation

$$\frac{df}{dt} = \frac{\partial f}{\partial t} + \frac{\partial f}{\partial x^i} \frac{dx^i}{dt} + \frac{\partial f}{\partial p^i} \frac{dp^i}{dt} = 0 \quad (2.1)$$

where $f(\mathbf{x}, \mathbf{p})$ is the phase space distribution function of the dynamical species in the simulation volume. Working in Newtonian gauge with the Newtonian potentials Φ and Ψ , which are typically of order 10^{-5} in units where $c = 1$, we have up to first order in the potentials [10]

$$\frac{dx^i}{dt} = \frac{p^i}{aE} (1 - \Phi + \Psi) . \quad (2.2)$$

In equation 2.2, $E = \sqrt{p^i p^i + m^2}$ and a is the scale factor. Similarly,

$$\frac{dp^i}{dt} = -p^i \frac{\partial \Phi}{\partial t} - p^i \frac{\dot{a}}{a} - \frac{E}{a} \frac{\partial \Psi}{\partial x^i} \quad (2.3)$$

to first order in the potentials. Here \dot{a} means a derivative with respect to time, and not a conformal time derivative. Substituting 2.2 and 2.3 into 2.1 we have

$$\begin{aligned} \frac{\partial f}{\partial t} + \frac{\partial f}{\partial x^i} \left[\frac{p^i}{aE} (1 - \Phi + \Psi) \right] \\ + \frac{\partial f}{\partial p^i} \left[-p^i \frac{\partial \Phi}{\partial t} - p^i \frac{\dot{a}}{a} - \frac{E}{a} \frac{\partial \Psi}{\partial x^i} \right] = 0. \end{aligned} \quad (2.4)$$

2.2.2 Poisson equation

The metric potentials Φ and Ψ are related to the matter content via the Poisson equation. Since we are interested in epochs during matter domination ($z \lesssim 300$) when the overall anisotropic stress is small, we can assume $\Psi = -\Phi$. In Newtonian gauge, we have

$$\nabla^2 \Psi = 4\pi G a^2 \left[\bar{\rho} \delta - 3 \frac{\dot{a}}{a} (1 + w) \partial_i v^i \right] \quad (2.5)$$

where v_i is the local peculiar velocity, $\delta = \delta\rho/\bar{\rho}$ is the local overdensity, and $w = \bar{P}/\bar{\rho}$ is the background equation of state. All of these fluid quantities will be defined formally in terms of moments of the distribution function below. For simulation boxes where the volume is much smaller than the Hubble volume, the second term in the brackets on the right hand side of eqn. 2.5 can be neglected.

2.2.3 Obtaining the Boltzmann moment equations

We can use the Boltzmann equation 2.1 to derive fluid equations for collisionless particles by integrating over various moments. First, we derive the continuity equation in the usual way, by multiplying 2.4 by E and integrating over momentum. We define the real-space density in terms of the phase space density in the following manner,

$$\rho(\mathbf{x}) = \int E f(\mathbf{x}, \mathbf{p}) d^3 \mathbf{p} \quad (2.6)$$

Since we are interested in the relative density contrast $\delta = (\rho - \bar{\rho})/\bar{\rho}$, we can cast the continuity equation into an equation for the time evolution of δ :

$$\begin{aligned} \dot{\delta} = -\frac{1}{a} \frac{\partial [(1 + 2\Psi)\Pi^i]}{\partial x^i} & - 3\dot{\Phi}(1 + \delta)(1 + W) \\ & - 3\frac{\dot{a}}{a}(1 + \delta)(W - w) . \end{aligned} \quad (2.7)$$

where $w = \bar{P}/\bar{\rho}$ is the background equation of state and we have defined

$$\Pi^i(\mathbf{x}) = \frac{\int p^i f(\mathbf{x}, \mathbf{p}) d^3\mathbf{p}}{\bar{\rho}} \quad (2.8)$$

and

$$W(\mathbf{x}) = \frac{\int \frac{p^i p^i}{3E} f(\mathbf{x}, \mathbf{p}) d^3\mathbf{p}}{\rho(\mathbf{x})} . \quad (2.9)$$

If the species is non-relativistic, the bulk velocity can be simply defined as

$$v^i(\mathbf{x}) = \frac{\Pi^i(\mathbf{x})}{1 + \delta(\mathbf{x})} \quad (2.10)$$

In eqn. 2.7, the last two terms on the right hand side are typically small compared to the first term. Since we are interested in matter domination regimes, $\dot{\Phi}$ is small compared to the spatial derivatives of Φ , and we can neglect the second term on the RHS without a loss of accuracy. If the particles are relativistic, the last term is small or zero because the local sound speed and the background sound speeds are the same, and $(W - w)$ vanishes. When particles are non-relativistic, W and w are individually small ($\sim 10^{-6}$) and even in non-linear regimes the last term remains much smaller than the first term.

Next we multiply eqn. 2.4 by p^i and integrate to get the Euler equation for the fluid

$$\begin{aligned} \dot{\Pi}^i = -(1 - 3w) \frac{\dot{a}}{a} \Pi^i & - \frac{1}{a} \frac{\partial [(1 + \delta)W^{ij}]}{\partial x^j} \\ & - \frac{1 + \delta}{a} \frac{\partial \Psi}{\partial x^i} \end{aligned} \quad (2.11)$$

with

$$W^{ij}(\mathbf{x}) = \frac{\int \frac{p^i p^j}{E} f(\mathbf{x}, \mathbf{p}) d^3 \mathbf{p}}{\rho(\mathbf{x})} \quad (2.12)$$

Note that, in principle, we can continue this procedure of generating equations using higher and higher moments of the Boltzmann equation. Each equation will be coupled to the next - this is apparent by looking at the structure of the two equations we have derived, 2.7 and 2.11. In the next section, we comment on how we close this infinite hierarchy of equations.

2.2.4 SPH equations

As we will show in §2.5, there are some situations in which it is advantageous to use a Lagrangian description of the fluid rather than an Eulerian description. We use a Smoothed Particle Hydrodynamics (SPH) approach for these problems, following the procedures in [77]. For non-relativistic collisionless particles, the equations of motion can be written as

$$\frac{dx^i}{dt} = \frac{v^i}{a} (1 - \Phi + \Psi) \quad (2.13)$$

and

$$\begin{aligned} \frac{dv^i}{dt} = & - v^i \frac{\partial \Phi}{\partial t} - v^i \frac{\dot{a}}{a} - \frac{1}{a} \frac{\partial \Psi}{\partial x^i} \\ & - \frac{1}{(1 + \delta) a} \frac{\partial ((1 + \delta) W^{ij})}{\partial x^j}. \end{aligned} \quad (2.14)$$

2.3 Closing the hierarchy

To study the evolution of the collisionless fluid with the above equations, we need some way to close the Boltzmann hierarchy. For collisional fluids, one can use an equation of state to relate the energy density and the pressure to close the system of equations. However, in the collisionless cases that we will be interested in (neutrinos and WDM) there is no simple equation of state, and so an alternative closure method is required.

To motivate our method to close the Boltzmann hierarchy, it will be useful to consider the approach used in particle-mesh (PM) N-body simulations [78]. In N-body simulations, Lagrangian

particles are evolved under the influence of their collective gravitational field. Those particles represent a Monte Carlo sampling of the distribution function, and in PM simulations, those particles are used to estimate the density field $\delta(\mathbf{x})$ that enters Eqn. 2.5. In effect, PM simulations use particles to close the Boltzmann hierarchy at its zeroth moment. In principle, however, we could use those particles to estimate other quantities that enter the fluid equations. For example, we could use particles to estimate a bulk fluid velocity, and then use that estimated velocity to evolve the density field using the continuity equation 2.7. Alternatively, we could use particles to estimate a stress tensor entering the Euler equation 2.11 that would truncate the hierarchy at its 2nd moment. Indeed we can estimate an arbitrary moment of the distribution function from particles, and truncate the Boltzmann hierarchy accordingly.

Therefore, the method we use is the following. We represent the collisionless fluid (e.g. neutrinos or WDM) simultaneously using fluid quantities on a grid and using test particles as well. We evolve the grid fluid using fluid equations like 2.7 and 2.11, and we truncate the Boltzmann hierarchy of fluid equations using moments of the distribution function estimated from the test particles, e.g. Eqn. 2.8 or 2.12. The test particles evolve under the gravitational field estimated from the fluid, i.e. Eqns. 2.2 and 2.3. Compared to traditional PM simulations, our approach involves solving more equations than the Poisson equation (i.e. fluid equations), and involves estimating higher moments of the particle distribution function (i.e. 3 components of the bulk velocity, or 6 components of the stress tensor, rather than a single scalar density field). This would appear to be considerably more expensive than traditional PM codes, but as we argue below, the benefits of using this approach in certain situations can outweigh the added costs.

2.3.1 Early evolution

We initialize our simulations using perturbation theory. We use Eulerian PT to initialize fluid quantities on the grid, and Lagrangian PT to initialize the test particles. In addition to the LPT velocities, the test particles are also given random thermal velocities drawn from the distribution function of the species we are interested in.

At early times, the particles can have large thermal velocities. These random thermal velocities can produce shot noise in any quantity we try to estimate from the particles, in the same way that

the shot noise in particle positions generates noise in the density field computed from the particles, as discussed in the introduction. One difference between the shot noise in particle velocities, compared to the noise in particle positions, is that the velocity noise diminishes over time as the universe expands. This means that shot noise in fluid quantities like the bulk velocity or velocity dispersion becomes small at low redshift, in contrast to the shot noise in the estimated density field. This illustrates one reason why it can be advantageous to estimate quantities other than the density from the particles.

Nevertheless, at early times the shot noise in velocities is large. In principle, this can be suppressed by increasing the number of particles in the simulation, but in many cases of interest, the required number of particles is orders of magnitude too large to be feasible. Therefore, using more particles in the simulation is not a practical solution in most situations of interest.

Another way to effectively increase the number of particles used in estimating fluid quantities is to spatially smooth those quantities. Spatially smoothing the fluid quantities is equivalent to estimating those quantities at a point using a larger volume, and hence more particles. The obvious reason why simulations normally do not spatially smooth over large volumes to suppress shot noise is that smoothing erases any small-scale structure in the estimated quantities. For species like cold dark matter, there is structure on all spatial scales, and so spatially smoothing would incorrectly eliminate real physical structure in the DM distribution. However, for the hot species of interest to us, the high temperature implies that there is a Jeans scale $k_J \sim aH/c_s$ below which small-scale structure is actively suppressed. Arguably, therefore, spatially smoothing is not necessarily invalid as long as the smoothing length is always safely below the Jeans scale.

On the other hand, the power spectrum does not vanish at $k > k_J$. Since we would like to accurately evolve the power spectrum on all scales in the simulation, including scales below the Jeans scale, this restricts what fluid quantities we can spatially smooth. For example, we cannot use the smoothing technique to estimate the velocity field, because this will extinguish the growth of structure on all scales below the smoothing scale. This becomes apparent if one looks at the linear continuity equation in Fourier space:

$$\dot{\delta}_k = -\frac{i\mathbf{k}\cdot\mathbf{v}_k}{a} \quad (2.15)$$

Smoothing sets \mathbf{v}_k to 0 for $k > k_{\text{smoothing}}$, which gives the unphysical result $\dot{\delta}_k = 0$.

However, from our tests, we find that we *can* spatially smooth the velocity dispersion W^{ij} at early times without sacrificing accuracy. To see this, note that when the thermal velocities are large enough to require spatial smoothing to suppress shot noise, then the Jeans scale is large, and so structure in the hot species remains linear. Under the fluid approximation, fluctuations in the velocity dispersion are second order, however, since they involve two perturbed velocities. Therefore, spatially smoothing the velocity dispersion only drops second order fluctuations in the fluid approximation, and does not affect the linear evolution of the velocity field in the Euler equation 2.11. This breaks down when structure in the hot species becomes nonlinear, but in order for structure to become nonlinear, the thermal velocities must be small, eliminating the need for spatial smoothing to suppress shot noise. Therefore, at all redshifts, we can estimate fluid quantities without significant thermal shot noise.

In order for this method to work, we must pick a sensible smoothing length. If the smoothing length is too small, shot noise will corrupt the evolution of fluid quantities, whereas if the smoothing length is too large (exceeding the Jeans scale), then spatial smoothing artificially removes real physical structure in the simulation. We set the smoothing length using the following argument. The quantity we are estimating from the particles is the velocity dispersion, and we require that the error in our estimate of the dispersion to be small compared to velocities generated by gravity. We estimate the dispersion at every point on the grid by evaluating

$$W^{ij}(\mathbf{x}) = \frac{\sum_{\mathbf{x}} \frac{p^i p^j}{E}}{\sum_{\mathbf{x}} E} - \frac{\sum_{\mathbf{x}} p^i}{\sum_{\mathbf{x}} E} \frac{\sum_{\mathbf{x}} p^j}{\sum_{\mathbf{x}} E} \quad (2.16)$$

where $\sum_{\mathbf{x}}$ stands for the sum over all particles at position \mathbf{x} . Note that for non-relativistic particles, E for every particle is approximately the mass, and so the denominators in the above expression count the total mass of particles at a given point. We use a cloud in cell (CIC) scheme to evaluate the different sums. The average thermal velocity dispersion is given by the equation of state w . Therefore, if N particles have been used to estimate the velocity dispersion, the error in the estimate will be $\Delta = w/\sqrt{N}$. The error is going to be small if Δ is small compared to the average velocity dispersion sourced by gravity $\langle v_{esc}^2 \rangle$. The latter can be estimated by evaluating the average of the

magnitude of the gravitational potential in the box, $|\Phi_{rms}|$. If we set $\Delta = \epsilon|\Phi_{rms}|$ for some error tolerance $\epsilon \ll 1$, then the number of particles that are needed to achieve the error tolerance is

$$N = \frac{w^2}{\epsilon^2|\Phi_{rms}|^2}. \quad (2.17)$$

When structure is linear, we know that n , the average number of particles per pixel, is a good estimate of the actual number of particles per pixel (modulo shot noise). This means that we have

$$N = nV_s \approx nL_s^3 \quad (2.18)$$

where $V_s \propto L_s^3$ is the smoothing volume required. This gives us an expression for the smoothing length

$$L_s = \left(\frac{w^2}{\epsilon^2|\Phi_{rms}|^2 n} \right)^{1/3}. \quad (2.19)$$

We can adjust ϵ and n to ensure that our smoothing length never exceeds the local Jeans scale calculated in linear theory. We also need to adjust our parameters so that we are no longer smoothing when non-linear structures like halos start forming in the simulation volume. At early times, the particles are hot (large w) and the smoothing length is large. As time progresses, the particles become colder, which means w decreases. At the same time, structure starts forming in the box and $|\Phi_{rms}|$ increases with time. These effects together mean that the smoothing length is a rapidly decreasing function of time.

2.3.2 Late evolution

As the simulation proceeds, the smoothing length reduces below a pixel at some redshift. The random motions of particles at redshifts after this time do not produce levels of shot noise which will affect the evolution of the power spectrum. Once the shot noise becomes negligible, we do not need to spatially smooth the fluid quantities. At subsequent times, our fluid approach may appear unnecessary, given the computational costs associated with our method compared to traditional N-body methods. Unfortunately, we cannot switch from fluid evolution to N-body even when the smoothing length is less than 1 pixel. The reason is that the particle density field has shot noise in it even at late time, arising from the motion of the particles at higher redshifts. Using the

density field for further evolution would mean that structure induced by this shot noise would grow gravitationally, and would start forming spurious halos at lower redshifts. This means that even at later times, the density field of the particles should not be used directly to source gravity.

However, we can safely switch from estimating the velocity dispersion and evolving the fluid density and bulk velocity fields, to estimating the bulk velocity field and only evolving the density field using the continuity equation 2.7. This is safe at late times because we are not spatially smoothing the estimated velocity field. Switching from estimating dispersion to estimating velocity produces a considerable speed-up in the simulation, since fewer quantities are being estimated (3 velocities vs. 6 dispersions) and fewer fluid equations need to be evolved (only continuity, not Euler). To estimate bulk velocities on the grid from the positions and velocities of the particles, we once again use a CIC interpolation scheme,

$$u^i(\mathbf{x}) = \frac{\sum_{\mathbf{x}} p^i}{\sum_{\mathbf{x}} E}. \quad (2.20)$$

As noted above, switching to velocity estimation speeds up the code significantly. We have verified that we obtain consistent simulation results using the faster velocity estimation and the slower dispersion estimation at late times. Therefore, in all of our Eulerian simulations below, we will switch to velocity estimation once the smoothing length shrinks to below one pixel.

2.3.3 Smoothing in SPH simulations

We will show below in §2.5 that in certain cases, it can be advantageous to use a Lagrangian formulation of the fluid equations rather than an Eulerian description. The Lagrangian description we will use in those cases is smoothed particle hydrodynamics (SPH). In those simulations, we use the following technique. We use two sets of particles - the first set evolves following the SPH equations of motion Eqs. 2.13, 2.14, and a second set of test particles which are evolved using Eqs. 2.2, 2.3. Like in the Eulerian case, the test particles are given thermal velocities drawn from their Fermi-Dirac distributions. Once again, these particles are not used in any of the density estimates, but only used to measure the velocity dispersion on a grid, as defined in Eqn. 2.16. Such an estimate of the velocity dispersion will be plagued by shot noise at early times, and needs to be smoothed, and this is done following the same prescription for the smoothing length discussed

in §2.3.1. Given the smoothed velocity dispersions on a grid, we interpolate from the grid to the positions of individual SPH particles to assign velocity dispersions to them. Once the thermal dispersions have been assigned, we use standard cubic spline interpolations from [77] to measure the density and the velocity dispersion gradient required in Eq. 2.14. For example, the density the position of the a -th particle is given by [77]

$$\rho(\mathbf{r}_a) = \sum_b m_b W(\mathbf{r}_a - \mathbf{r}_b, h) \quad (2.21)$$

where the sum runs over all other particles, W is the interpolation kernel, and h is the spline smoothing length, which we choose to be one grid cell. Notice that this spline smoothing length is different from the smoothing length we defined in Eq. 2.19.

Once the smoothing length from Eq. 2.19 falls below our force resolution on the grid, we start treating the SPH particles as standard N-body particles for the rest of the evolution, as the effects of their thermal velocities beyond that point is below the resolution of the simulation. In practice, this means that we use Eqs. 2.2, 2.3 for time evolution.

Our technique does not make full use of the capabilities of SPH, as we use an intermediate grid to find and smooth the velocity dispersions from the test particles. This automatically limits the resolution of the SPH technique, but since our objective was only to suppress the shot noise from thermal velocities and given that our gravitational force resolution is also limited by the same grid, the above method is sufficient for the purposes of our simulations.

2.4 Integration techniques

2.4.1 Particle time integration

For the test particles we use in the simulations, as well as for CDM particles in neutrino simulations, we use the standard Kick-Drift-Kick leapfrog time integration [79]:

$$\mathbf{v}^{n+\frac{1}{2}} = \mathbf{v}^n + \frac{\Delta t}{2} \mathbf{f}^n \quad (2.22)$$

$$\mathbf{x}^{n+1} = \mathbf{x}^n + \Delta t \mathbf{v}^{n+\frac{1}{2}} \quad (2.23)$$

$$\mathbf{v}^{n+1} = \mathbf{v}^{n+\frac{1}{2}} + \frac{\Delta t}{2} \mathbf{f}^{n+1} \quad (2.24)$$

where \mathbf{x}^i , \mathbf{v}^i are the particle positions and velocities at time step i , and \mathbf{f}^i are the forces at those timesteps. For the Kick-Drift-Kick method, the Poisson equation is solved after the particle or Drift update - with the updated positions of the particles. This method is formally second order accurate, apart from being symplectic in nature [79].

The Drift-Kick-Drift method [79]

$$\mathbf{x}^{n+\frac{1}{2}} = \mathbf{x}^n + \frac{\Delta t}{2} \mathbf{v}^n \quad (2.25)$$

$$\mathbf{v}^{n+1} = \mathbf{v}^n + \Delta t \mathbf{f}^{n+\frac{1}{2}} \quad (2.26)$$

$$\mathbf{x}^{n+1} = \mathbf{x}^{n+\frac{1}{2}} + \frac{\Delta t}{2} \mathbf{v}^{n+1} \quad (2.27)$$

is also second order accurate and symplectic, but we use the Kick-Drift-Kick method because the latter uses the forces (and hence the potential) at full time steps, whereas the former uses the potential at half time steps. The test particles are going to be coupled to a fluid, whose own time integration scheme gives the potential at full time steps, and so using Kick-Drift-Kick is essential for the fluid and the particles to remain coupled to each other.

2.4.2 Hydrodynamics

As we saw earlier, we will be solving the continuity and Euler equations at early times when the smoothing length defined by Eq. 2.19 is larger than a grid cell, and most structure in the box is linear. At late times, when the smoothing length falls below a grid cell and highly non-linear structures start forming in the box, we will be solving only the continuity equation. Both equations are hyperbolic partial differential equations with source terms.

We use an operator splitting method to split any source term present in the equations from the hyperbolic advection part. We also use directional splitting [80] so that the advection in 3 dimensions is reduced to 3 1-dimensional advectons. To solve the 1-dimensional advection problem on the grid, we employ a finite-volume scheme which is piecewise linear, and hence second order accurate in space.

The advection equation

$$\frac{\partial q}{\partial t} = -\frac{\partial(q u)}{\partial x} \quad (2.28)$$

is discretized so that the update equation for q_i^n , the value of q at cell center i at timestep n can be written as

$$\frac{q_i^{n+1} - q_i^n}{\Delta t} = - \frac{f_{i+\frac{1}{2}}^{n+\frac{1}{2}} - f_{i-\frac{1}{2}}^{n+\frac{1}{2}}}{\Delta x} \quad (2.29)$$

where $f_{i\pm\frac{1}{2}}^{n+\frac{1}{2}}$ are the fluxes at the cell edges at timestep $n + \frac{1}{2}$ constructed from the data at timestep n . The fluxes are constructed in the following manner:

$$\begin{aligned} f_{i+\frac{1}{2}}^{n+\frac{1}{2}} &= 0.5 u_{i+\frac{1}{2}}^n \left(\left(1 + \theta_{i+\frac{1}{2}}\right) q_i^n + \left(1 - \theta_{i+\frac{1}{2}}\right) q_{i+1}^n \right) \\ &+ 0.5 \left| u_{i+\frac{1}{2}}^n \right| \left(1 - |c_{i+\frac{1}{2}}|\right) \phi(r_{i+\frac{1}{2}}^n) (q_{i+1}^n - q_i^n) \end{aligned} \quad (2.30)$$

$$\begin{aligned} f_{i-\frac{1}{2}}^{n+\frac{1}{2}} &= 0.5 u_{i-\frac{1}{2}}^n \left(\left(1 + \theta_{i-\frac{1}{2}}\right) q_{i-1}^n + \left(1 - \theta_{i-\frac{1}{2}}\right) q_i^n \right) \\ &+ 0.5 \left| u_{i-\frac{1}{2}}^n \right| \left(1 - |c_{i-\frac{1}{2}}|\right) \phi(r_{i-\frac{1}{2}}^n) (q_i^n - q_{i-1}^n) \end{aligned} \quad (2.31)$$

where $c_i = u_i \Delta t / \Delta x$. We define $\theta_{i\pm\frac{1}{2}} = 1$ for $u_{i\pm\frac{1}{2}} > 0$ and $\theta_{i\pm\frac{1}{2}} = -1$ for $u_{i\pm\frac{1}{2}} < 0$. We further define

$$r_{i-\frac{1}{2}}^n = \begin{cases} \frac{q_{i-1}^n - q_{i-2}^n}{q_i^n - q_{i-1}^n} & \text{if } u_{i-\frac{1}{2}} > 0 \\ \frac{q_{i+1}^n - q_i^n}{q_i^n - q_{i-1}^n} & \text{if } u_{i-\frac{1}{2}} < 0 \end{cases} \quad (2.32)$$

and similarly for $r_{i+\frac{1}{2}}^n$. $\phi(r)$ is the flux limiter function which is required so that the method is Total Variation Diminishing [81], by converting to a first order method near extrema in the profile of q . Even though there are no real shocks in the collisionless fluids that we will be dealing with, once non-linear structures start forming in the box, there are sharp density gradients, which lead to spurious oscillations if the hydro scheme we use is not TVD in nature. Though the TVD property is essential for the stability of the code, it also means that there will be artificial diffusion near the extrema - sharper the change in gradient, larger the diffusion. In our code tests and cosmological

simulations we use the Superbee flux limiter [82], defined as

$$\phi(r) = \begin{cases} 0 & \text{if } r < 0 \\ \min(2r, 1) & \text{if } 0 < r < 1 \\ \min(2, r) & \text{if } 1 < r \end{cases} \quad (2.33)$$

This turns out to be the least diffusive flux limiter among the ones we tested.

While the piecewise linear method is formally correct to second order in time in smooth regions, it switches to first order time accuracy near saddle points and extrema. In cosmological simulations, especially the ones involving WDM, saddle points appear throughout the box as structures form and move under the influence of gravity. This means that, typically, the solution will only be first order accurate in large parts of the box. To make the scheme at least second order accurate in time throughout the box, we use a second order explicit Runge-Kutta time integration scheme. We use data at timestep n for the predictor step to get the predicted data at time $n + 1$, and then use this information in the corrector step to get the final solution at time $n + 1$. As we mentioned in the previous subsection, we calculate fluid quantities like the density at every full timestep, and therefore, to keep the particles and the fluid coupled, we use the Kick-Drift-Kick method which requires forces and potentials at every full time step, rather than the Drift-Kick-Drift method.

2.4.3 Gravity

We use Fast Fourier Transforms (FFT) to determine the potential on the grid from the densities of the neutrino fluid and WDM or CDM particles, just as in PM simulations [78]. For CDM, we use cloud-in-cell (CIC) interpolation to obtain the grid density from the positions of the particles. In Fourier space, the Poisson equation is given by

$$\tilde{\phi}(\mathbf{k}) = G(\mathbf{k})\tilde{\delta}(\mathbf{k}) \quad (2.34)$$

where $G(\mathbf{k})$ is the Green function for the Poisson equation. On the grid, the discrete version of $G(\mathbf{k})$ becomes

$$G(k_x, k_y, k_z) = -C \left[\sin^2 \left(\frac{k_x}{2} \right) + \sin^2 \left(\frac{k_y}{2} \right) + \sin^2 \left(\frac{k_z}{2} \right) \right]^{-1} \quad (2.35)$$

where C is a constant independent of scale and $\{k_x, k_y, k_z\}$ are the wavenumbers on the cubic grid. Once we solve the Poisson equation using the FFT method, we calculate forces on the grid using a four point stencil

$$\begin{aligned} f_x(i, j, k) &= -\frac{\partial \phi_{(i,j,k)}}{\partial x} \\ &= -\frac{1}{12\Delta x} \left[8\phi_{(i+1,j,k)} - 8\phi_{(i-1,j,k)} \right. \\ &\quad \left. -\phi_{(i+2,j,k)} + \phi_{(i-2,j,k)} \right] \end{aligned} \quad (2.36)$$

where i, j and k label the coordinates of the grid points. These forces are then used to update the fluid when we solve the Euler equation. For updating the velocities of the particles, we again use a CIC interpolation to interpolate the forces from the grid to the positions of individual particles.

2.5 Code tests

In this section, we present various tests of the new method described above. The first two tests are designed to check the accuracy of this method for the dynamics of virialized objects. The third test is devised to check accuracy at early times, when traditional N-body simulations can produce large errors. We then compare the results of this method to N-body results for a Λ CDM universe, where we know the latter yields accurate results. In our final test, we run simulations of Warm Dark Matter (WDM) cosmologies to check whether this method is able to eliminate the spurious halos that are known to plague both “hot start” and “cold start” N-body simulations of WDM cosmologies.

2.5.1 Plummer sphere advection

The Plummer sphere has an isotropic mass distribution with a radial density profile given by

$$\rho(r) = \left(\frac{3M}{4\pi b^3} \right) \left(1 + \frac{r^2}{b^2} \right)^{-5/2} \quad (2.37)$$

where b is the Plummer radius, which sets the size of the virialized region. The potential for the Plummer sphere is

$$\Phi(r) = -\frac{GM}{(r^2 + b^2)^{1/2}} \quad (2.38)$$

For this density profile, the phase space distribution function $f(\mathbf{r}, \mathbf{v})$ is [83]

$$f(\mathbf{r}, \mathbf{v}) d\mathbf{r} d\mathbf{v} \propto (-E(r, v))^{7/2} r^2 v^2 dr dv \quad (2.39)$$

where $E = \Phi(r) + \frac{1}{2}mv^2$. At a given radius r , the probability of finding a particle with absolute velocity v is given by

$$\tilde{f}(v)dv \propto (-E(r, v))^{7/2} v^2 dv = \left(-\Phi(r) - \frac{1}{2}mv^2 \right) v^2 dv \quad (2.40)$$

Since the potential is known analytically at all values of r , the velocity distribution is known everywhere.

In our tests, we generate particles with density profile following 2.37 for $b = 5$ in grid units. For every particle we generate a random velocity whose magnitude is drawn from the probability distribution 2.40, and whose direction is drawn isotropically. To test advection, we give each particle an additional constant velocity. This constant velocity will have the effect of shifting the Plummer sphere in space maintaining the shape of the density profile. We also initialize the fluid on the grid with the same density profile.

We evolve the particles and the fluid using the methods mentioned in §2.4. In our tests, we compare two different types of runs. In the first case, gravity is sourced by the particles, and so the fluid density acts as a tracer, evolving passively due to the motion of the particles. This provides a good check for our method - by comparing the particle and fluid density profiles at different times, we can see if the two descriptions of the same underlying dynamics do remain closely coupled to




Figure 2.1: Density profile of an isolated advecting Plummer sphere with $b = 5$. The solid black lines indicate the initial density profile of particles (top panel) and fluid (bottom panel). The dotted lines in the top panel shows the final density profile of particles from 4 runs - a run where particles source gravity (black), a run where the fluid sources gravity and the particles are evolved using the Kick-Drift-Kick method (red), a fluid-sourced gravity run where the particles are evolved using the Drift-Kick-Drift method (blue), and a fluid-sourced gravity run where we estimated the velocity dispersion instead of the bulk velocity (green) - after 20 dynamical times. The dotted lines in the bottom panel shows the fluid profile from the same four runs after 20 dynamical times. While the other methods all agree well with the run in which particles source gravity, the Drift-Kick-Drift method is not suitable for coupling with our hydro scheme.

each other.

In the second type of runs, we use the fluid itself to source the gravity, which is the method we will use in our actual cosmological simulations. In this case, the particles act as tracer particles which are used only to estimate either the bulk velocity or the velocity dispersion on the grid. Again we check if the two density profiles - one from the particles and the other from the fluid - match each other at different times.

We compare these runs in Fig 2.1. We represent the particle-based gravity run by the black dotted lines. A fluid gravity run based on estimating bulk velocity from particles, and using Kick-Drift-Kick to update the particles is shown with the red dotted line. A similar run, but using

Drift-Kick-Drift for the particles is plotted with the blue dotted line. Finally a fluid gravity run which the test particles (KDK evolution) to estimate the velocity dispersion is shown with the green dotted line. We see that the Kick-Drift-Kick method provides a better coupling to our hydro method than the Drift-Kick-Drift method. We also find that the run which used velocity dispersion estimation and the run which used bulk velocity estimation agree very well. Therefore, whenever we do not need to smooth quantities to suppress shot noise, we will use the faster bulk velocity estimation method.

We next investigate how the results from these tests are affected by the resolution of the advecting Plummer sphere. In the previous example with $b = 5$, the Plummer sphere was resolved by roughly ten grid elements in each dimension, meaning it was well resolved. In cosmological simulations, depending upon the shape of the power spectrum, small virialized objects which are not well resolved may form. Therefore, we redo our test for $b = 4$, $b = 3$, and $b = 2$. We see from Fig. 2.2 that as we reduce the value of b , and therefore the resolution, the density profiles from the particle gravity runs after 20 dynamical times start to diverge from the fluid gravity runs. For the case where $b = 2$, the difference in the density profile of particles between the two types of runs (where particles source gravity and where fluid sources gravity) is as much as 15%. These differences grow over time, and will cause artificial damping of small scale structures.

We also test how the number of particles we use to estimate the bulk velocity or velocity dispersion on the grid affects the advection. We find that for $b = 5$, the results do not change much as long as we use more than $\sim 10^5$ particles as seen in Fig. 2.3.

2.5.2 Collision of Plummer spheres

To further test our code for dynamics of virialized objects, we next consider what happens when two individual Plummer spheres are made to advect through each other. If the two Plummer spheres form an isolated system, the mutual gravitational attraction would mean that the Plummer spheres would slosh through each other before finally merging into one bound object. The dynamics of this system is analogous to the ubiquitous merger of structures one finds in typical cosmological simulations.

In our test, we have two Plummer spheres with the same Plummer radius ($b = 5$ in grid units)




Figure 2.2: Density profiles for advecting Plummer spheres after 20 dynamical times for $b = 5$ (black), $b = 4$ (red), $b = 3$ (blue) and $b = 2$ (green). The solid lines represent the profiles from runs where gravity is sourced by particles. The dotted lines are from runs where the fluid sources gravity. The top panel plots the particle density profile, normalized by the maximum value of particle density for a given value of b from the particle based runs. The bottom panel plots the fluid density profiles with the same normalization. The profiles have been staggered for clarity.

but different masses. We use the same initialization technique as in the previous test - modified to take into account the different masses of the two spheres. We also give bulk velocities to the two Plummer spheres so that they move toward each other head on.

Once again, we compare what happens when gravity is sourced by the particles and when it is sourced by the fluid. For the latter case, we also compare how the results of this test is affected by our choice of the flux limiter for our hydro scheme. We show the comparison in Fig.2.4. We see that the particle profiles (top panels) and fluid profiles (bottom panels) of the run with the Superbee (SB) flux limiter (red curve) remain closer to the particle based run (black curve), than the run in which we used the Monotonized Central (MC) flux limiter [84] (blue curves), especially at late times. This is understandable because the MC flux limiter is known to be more diffusive than the Superbee.




Figure 2.3: Particle (top panel) and fluid (bottom panel) density profiles for an advecting Plummer sphere with $b = 5$, with different number of test particles (N_p) used to sample the initial velocity distribution and estimate bulk velocity. We find that for this size, $N_p \sim 10^5$ is sufficient to follow the dynamics correctly.

2.5.3 Linear growth rate

In cosmology, the high-redshift growth of the power spectrum of all species, including light particles like neutrinos and heavy particles like WDM, can be calculated in linear theory. We use the linear Boltzmann code CLASS [85, 86, 87] to do these calculations. The power spectrum and transfer functions from CLASS at the starting redshift, z_{start} , are used to initialize our simulations. The evolution of the power spectrum from the simulation boxes can then be compared at later redshifts to the outputs from CLASS at those redshifts.

In simulations involving light neutrinos, linear perturbation theory can be used to describe the evolution of the neutrino power spectrum for most of the evolution, except maybe at very late times. This means that for most of the simulation, the smoothing length defined in Eqn. 2.19 is larger than a pixel, and we need to estimate the velocity dispersion from the test particles. Comparing the growth of the power spectrum at these times against CLASS provides a test for the code when




Figure 2.4: Evolution of the density profiles of two colliding Plummer spheres. The solid curves represent the particle density (upper panels) and fluid density (lower panels) for three different runs - gravity from particles (black), gravity from fluid Superbee (SB) flux limiter (red), and another fluid gravity run with Monotonized Central (MC) flux limiter (blue). The dotted curve in each panel shows the initial configuration. The left panels represent the densities when the Plummer spheres have passed through each other once. The right panels represent densities when the two spheres have merged.

it is solving both Eqns. 2.7 and 2.11. Note that at these early times, the growth of the neutrino power spectrum is scale dependent. On scales larger than the free streaming scale of the neutrinos, the neutrino power spectrum grows at the same rate as the CDM power spectrum. Below the free streaming scale, however, different scales can grow at different rates. To compute the growth factor accurately on all scales, it is essential to determine the correct “effective sound speed” for the neutrinos. This is exactly what we estimate from the velocity dispersion. The effective sound speed depends on redshift, meaning that a comparison of the power spectrum at different times provides a test for the accuracy of our estimates of the velocity dispersion.

We check for the linear growth rate for neutrinos with a neutrino species with mass $m_\nu = 0.1 \text{ eV}$. The starting redshift for these simulations was $z = 199$. We compare the results from our simulations to the growth predicted by CLASS for the same cosmology at redshifts $z = 22.53$ and $z = 0$ in Fig. 2.5. We see that the power spectra from our simulations match the linear theory

predictions quite well at all scales in the simulation box at $z = 22.53$. Leaving aside the large scales, which are affected by sample variance, there are, however, differences of about 10% even on the smaller scales. This arises from the fact we use the particles directly to get the velocity dispersion, while in CLASS, an approximation is used to evaluate the sound speed when the fluid approximation is turned on [87]. We have checked that if we use the same approximation in our code, instead of the velocity dispersion tensor measured from the particles, we match the CLASS results to within a few percent. However, this small difference in the neutrino power spectrum will have minimal effect on the observable matter power spectrum. We note that all our comparisons with CLASS were with the fluid approximation turned on, which means that the effective sound speed of the neutrinos is treated to be scale independent. While this is done explicitly in CLASS when the fluid approximation is turned on, in our simulations, the large smoothing length at early times for light neutrino species means that the velocity dispersion is scale independent over much of the simulation box - effectively using the fluid approximation.

To illustrate the effect of shot noise in neutrino simulations, we also plot in Fig. 2.5 the neutrino power spectrum at the later redshifts. These spectra were obtained from a simulation which treats the neutrinos as N-body particles. The power spectrum of neutrinos from this simulation is dominated by shot noise for most scales in the box - this is true even at $z = 0$.

While using a high redshift z_{start} helps eliminate transient effects coming from the initial conditions, one concern about starting the above simulations at $z_{\text{start}} = 199$ is that the size of the simulation box is comparable to the horizon size at that redshift. However, as shown in [88], by modifying the positions of particles, one can still use N-body simulations to accurately evolve the large scale modes. For the neutrino component, we could similarly account for super-horizon scales by including all the terms in Eqs. 2.7 and 2.11. While we have not incorporated these into our simulations, we have checked with lower resolution simulations that using a lower z_{start} - which ensures that the simulation box is always smaller than the horizon - does not appreciably change the results shown in Fig. 2.5. In the simulations shown in §2.6, we always start at $z = 49$, which for the box sizes considered, always lie within the horizon.

In Fig. 2.6, we show an example of a WDM simulation. In this case, while the power spectrum is initially damped on small scales, non-linear structure does develop at high k as time progresses.




Figure 2.5: Comparison of the growth of the power spectrum of neutrinos ($m_\nu = 0.1 \text{ eV}$). The top panel plots the comparison of the power spectra from the hybrid simulations (solid lines) to the outputs from CLASS (dot-dashed lines). We plot the initial power spectra at $z = 199$, as well as at redshifts of $z = 22.53$ and $z = 0$. We also plot the power spectra (dashed lines) at the later redshifts from an N-body treatment of neutrinos, which are dominated by shot noise at small scales. In the bottom panel we plot the ratio of the power spectra from the hybrid simulations to the linear theory results at $z = 199$ (teal), $z = 23.53$ (black) and $z = 0$ (red).

The power spectrum on these scales will then disagree with the linear CLASS power spectrum. However, if the simulation volume is large enough, then the largest scales in the box will still be well described by linear theory, and can be matched to the outputs from CLASS. Unlike the neutrino simulations, where the energy density and gravitational potential are dominated by CDM particles, the growth of the matter power spectrum in the WDM simulations is governed by the fluid, and testing the growth rate in these simulations provides a stronger test of the coupling between the fluid equations and the gravitational potential. Fig. 2.6 shows the growth of the power spectrum from a 200eV WDM particle compared to the outputs of CLASS at two different redshifts, $z = 60.54$ and $z = 27.57$. At these early times, our method is able to reproduce the linear results




Figure 2.6: Linear growth of the power spectrum for a WDM particle of mass 200 eV. The top panel plots the power spectrum at initial redshift $z = 99$ and at $z = 60.54$ and $z = 27.57$ from the hybrid simulations and from linear theory. The bottom panel plots the ratio of the power spectra from the simulations and from linear theory at $z = 99$ (teal), $z = 60.54$ (black) and $z = 27.57$ (red). At large scales in the box, the ratio is different from 1 due to sample variance.

down to almost the smallest scales in the box.

2.5.4 CDM comparison run

Even though our method is meant to be used in cosmologies which include particles whose thermal velocities cannot be ignored, we can test the validity of our method by using it for a standard CDM cosmology run, and comparing the results to those of an N-body simulation. Since N-body simulations are known to be accurate for CDM-only cosmologies, this comparison allows us to test our method starting at high redshifts when linear theory is valid to late times and low redshifts where we can test the non-linear evolution of our method.

For this comparison, we do not need to smooth quantities, since there are no random particle

velocities which will lead to thermal shot noise. We use the test particles to measure the bulk velocities on a grid at all redshifts, and use these velocities to solve the continuity equation 2.7 to evolve the densities forward. Both simulations are done on a $140 h^{-1}\text{Mpc}$ box with $\Omega_\Lambda = 0.7$ and $\Omega_{\text{CDM}} = 0.3$. The Hubble constant H_0 is taken to be 70 km/s/Mpc . Both sets of simulations used 512^3 particles with a 512^3 grid to calculate the gravitational potential. Our method also uses a 512^3 grid to store densities and bulk velocities. Initial conditions are generated at redshift $z = 99$ using CLASS.

We begin by comparing the growth of the power spectrum on scales which remain linear - the largest scales in the simulation box. At early times, the growth at these scales matches the growth from standard N-body simulations. This agreement persists even to times when halos start forming in the box. However, at very late times, the agreement breaks down, and the largest scales grow at a rate slower than in N-body simulations. This can be seen clearly from the lower panel in Fig. 2.7, where the ratio of the power spectrum from our hybrid simulation to the power spectrum from the N body calculation (both at $z = 0$) is less than 1 at even the largest scales in the box.

At late times in these CDM simulations, structures form at all scales and at all locations. Even deep inside voids, there are small fluctuations in the density field. Because of these small variations, the density field is not smooth over most of box, and there are large numbers of local extrema and saddle points along any of the axes of the simulation box. As mentioned in §2.4.2, our hydrodynamic scheme needs to switch over to a spatially first order scheme whenever it encounters a saddle point - unphysical oscillations set in when this condition is not satisfied. Due to the numerous saddle points which develop late in these CDM simulations, our method is forced to solve the governing equations of motion in a spatially first order manner over most of the pixels in the box. First order methods are known to be highly diffusive - even up to scales comparable to the entire simulation box. Because of this, the power at large scales is damped, and the growth rate of these scales becomes unphysically slow. We found that most of the pixels at which the hydrodynamic method is forced to go first order lie inside voids and regions where $|\delta| \sim 1$, rather than inside halos or regions where $|\delta| \gg 1$. Note that this problem is not as severe for WDM cosmologies in which the streaming scale is resolved. However, there is still numerical diffusion on small scales compared to an N-body treatment, and we will discuss below in §2.5.5 how this can

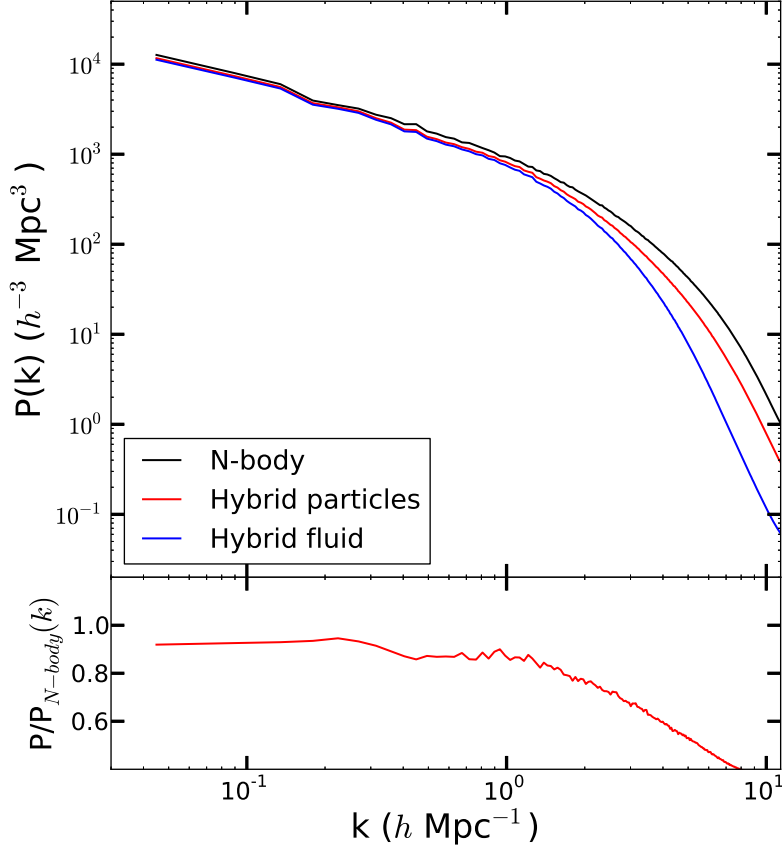


Figure 2.7: Final power spectrum at $z = 0$ for CDM runs. In the top panel, we plot the results a normal N-body run (black) as well as from the particles (red) and fluid (blue) in our hybrid simulations. In the bottom panel we plot the ratio of the particle power spectrum from the hybrid simulations to the power spectrum from the N-body run.

be remedied using an SPH-like approach at early times.

To further illustrate the effects of artificial diffusion, we next compare the mass functions of halos from our simulations to N-body halo mass functions. To find halos in our simulations, we use the Rockstar halo finder [89] as well as a spherical overdensity halo finder. For the latter method, we define a halo as spherical regions of radius R_{halo} around overdensity peaks within which the average enclosed matter overdensity is greater than 200. The associated halo mass is then defined as

$$M_{\text{halo}} = \frac{4}{3}\pi R_{\text{halo}}^3 \times (200 \Omega_m \rho_{\text{crit}}) \quad (2.41)$$

where ρ_{crit} is the critical density of the universe.

In our simulations we can use either the fluid density or the test particle density to define halos.




Figure 2.8: Mass function from the N-body run (black) and the hybrid Eulerian method (red) for a Λ CDM cosmology. The red curve shows a deficit of halos at all mass scales, including near the high end of the mass function. In contrast, the hybrid Lagrangian (SPH) method is formally identical to N-body for CDM cosmologies.

Since there is no thermal shot noise in the test particle density field in these CDM simulations, there is no problem using the test particle densities in our halo finder. However, we can also do the same even in WDM simulations where there is shot noise from the thermal velocities of particles. This is because throughout the simulation, we never use the test particle density to source gravity, hence the shot noise in the density field is not allowed to grow gravitationally. Therefore the shot noise fluctuations in the density field of the test particles arise at an early time and become frozen once the particles cool down. These initial fluctuations are much smaller than overdensities of $\delta \sim 200$ that are needed for a halo to be detected. Therefore, the halos that we detect at late times in the density field of test particles are real halos and not shot noise artifacts. This would not have been true if we had used the test particle density instead of the fluid density in the Poisson equation - in that case, the small shot noise fluctuations would have grown gravitationally over time and produced spurious halos.

Comparing mass functions, we find that if the halos defined using the test particle densities

rather than fluid densities, the results match those from the N-body simulations more closely. This is expected from the tests we performed in §2.5.1 and §2.5.2 - the fluid profile near the center is flattened by artificial diffusion. The particle profile is not as affected on scales which are well-resolved on the grid. However, there are still differences between the mass functions from our simulations compared to the N-body mass functions. Also, there are differences on small scales between the N-body power spectrum and the power spectrum from our simulations. This is once again due to artificial diffusion on small scales. For small halos with steep profiles, the flattening of the fluid profile means a large fraction of the mass is moved away from the peak. This also ends up affecting the test particles which respond to the gravitational field of the fluid.

This loss of smaller scale power due to diffusion leads to fewer smaller halos found in our simulation box than in the corresponding N-body boxes. However, at a given redshift, this can also cause differences in the masses of large, fast-accreting halos. This is because the lack of small scale power alters the dynamics of the halos and also the time at which mergers happen - the mergers in our simulations lag behind the corresponding merger in the N-body simulation, as can be seen in Fig.2.9. The fast accreting halos at a given redshift have a rapid change in their mass because of successive mergers. Because these mergers have not yet happened in the our simulations, the halos are at a lower mass. This shows up in the mass function as a lack of halos at the largest masses, as can be seen in Fig. 2.8.

Our findings show that in the CDM case, numerical diffusion in our grid-based fluid method affects all useful quantities measured from the simulation boxes - even at the largest scales. This motivates us to use a more Lagrangian approach for simulations in which there is structure on all scales in the box. Specifically, we adopt the SPH-like approach that we outlined in §2.3.3 for WDM simulations where the fluid that we are simulating is the main source of the gravitational potential. As the results in this section have shown, for such simulations any numerical diffusion becomes important. In contrast, for neutrino simulations, where the gravitational potentials are highly dominated by the CDM, we will continue to use grid-based Eulerian methods, as our implementation of SPH is computationally more expensive than our implementation of grid-based hydrodynamics.




Figure 2.9: A zoom-in of a halo at $z = 0$ from the CDM comparison simulations - the left panel shows the density field from a standard N-body simulation, while the right panel shows the density field from our hybrid method. In the N-body case, the central object has already merged, whereas in the hybrid method, the two objects are still distinct.

2.5.5 Simulating Warm Dark Matter cosmologies

In our final test, we simulate WDM cosmologies to see if our fluid methods can eliminate the artifacts generated for these cosmologies when simulated using N-body techniques. For N-body simulations which include the random thermal velocities of the WDM particles - or a hot start - these artifacts are spurious halos seeded by shot noise in the density field. For cold start simulations which do not include the thermal velocities, the artifacts are “beads on a string” halos studied extensively by [72].

For WDM, there is a characteristic free streaming scale below which the linear power spectrum is damped. Above this scale, the behavior of WDM is the same as CDM, so we will concentrate on simulations in which the free streaming scale is resolved. To estimate the scale of this damping, we take the linear power spectrum at $z = 0$ for CDM and the WDM particle we are interested in. We find the R_{damp} such that the CDM power spectrum convolved with $|W(kR_{\text{damp}})|^2$ gives the WDM power spectrum [70], where $W(kR)$ is the Fourier transform of the top-hat window function with radius R . We also define a damping mass scale

$$M_{\text{damp}} = \frac{4}{3}\pi R_{\text{damp}}^3 \Omega_m \rho_{\text{crit}} \quad (2.42)$$

where ρ_{crit} is the critical density of the universe.

For these WDM simulations, we use the hybrid SPH method laid out in Sec. 2.3.3 to suppress the effects of numerical diffusion. It is important to do so, because we are interested in determining the halo mass function below the damping scale down to the smallest scales in the box, and these scales are most affected by diffusion.

To validate our hybrid SPH method, we first run a simulation for a WDM particle with mass $m = 200 \text{ eV}$, for which the damping scale is $3.3h^{-1} \text{ Mpc}$. We do three runs for this cosmology, a hot start N-body run, a cold start N-body run, and a run using our hybrid SPH method. The initial thermal velocities are included in both the hot start run and our SPH method. These simulations were done on 512^3 grids with 1024^3 particles, and for the background cosmology we used $\Omega_\Lambda = 0.7$, $\Omega_m = 0.3$ and $H_0 = 70 \text{ km/s/Mpc}$. The size of the simulation volume was $(140h^{-1}\text{Mpc})^3$.

We first compare the final power spectra from all three simulations at $z = 0$, as illustrated in Fig. 2.10. We see that at large scales, the power spectra from all three runs agree with each other. However, at small scales, while the cold start N-body run and our hybrid SPH run give the same results, the power spectrum from the hot run starts to deviate from the others. This is expected in the hot run as the shot noise due to random streaming of the simulation particles leads to the presence of extra power at small scales, and this noise grows over the course of the simulation. However, we note that even though we included these thermal velocities in our hybrid SPH method, the final result agrees very well with that from the cold start run, down to the smallest scales that we resolve in this simulation box. This can be seen clearly from the bottom panel of Fig. 2.10, where we plot the ratio of the two power spectra, and find that the ratio is very close to 1 at all scales.

We also compare the halo mass functions from the three runs in Fig. 2.11. We expect the effect of the streaming length of the WDM particle to show up at masses below M_{damp} represented in the plot by the dashed brown line. Since the power spectrum was initially damped on these scales, we expect very few halos with these low masses to form in the box. That seems to agree with what we see from both the cold start run and the hybrid SPH run - the cumulative mass function flattens off at mass scales smaller than M_{damp} . However, the hot run does show the presence of many small halos. These are the spurious halos which were seeded by the shot noise arising from the thermal




Figure 2.10: Final power spectrum at $z = 0$ for WDM species with mass 200 eV. In the top panel, we plot the results from a hot N-body run (black), a cold N-body run (green), as well as our hybrid SPH simulation (red). The results from the SPH run and the cold start N-body run are virtually indistinguishable. On large scales in the box, the hot start N-body run agrees with the other two, but there are differences on small scales. In the bottom panel we plot the ratio of the particle power spectrum from the hybrid SPH run to the cold N-body power spectrum (red) to show that these methods match each other down to the smallest scales resolved by this simulation box, even though the initial conditions for the SPH run was the same as that for the hot start N-body run. We also plot the ratio of the power spectra from the hot run to the cold run (black) to show their difference on scales affected by the streaming of particles.

motions at early times.

These results suggest that our hybrid SPH method is effective in eliminating the effects of shot noise. Even though we started off by taking into account the thermal velocities of the particles from the Fermi-Dirac distribution, our final results agree very well with simulations which do not include these thermal velocities, and are therefore immune to this form of shot noise.

However, as discussed earlier, the cold start simulations for WDM show “beads on a string” artifacts [72], produced when structures collapse along grid lines in the simulations. We check if our




Figure 2.11: Mass function for WDM species with $m = 200$ eV. We plot results from a hot N-body run (black), a cold N-body run (blue), and our hybrid SPH method (red). The brown dashed line represents M_{damp} defined in Eqn. 2.42. The hot run shows the presence of many halos below M_{damp} , while the other runs do not show these smaller halos.

method can be used to get rid of these artifacts, which once again shows up in the mass functions at scales much smaller than the damping scale. For this, we simulate a lighter WDM particle, $M = 60$ eV, for which $R_{\text{damp}} \sim 10.2 h^{-1} \text{Mpc}$. This allows us to resolve scales much smaller than the streaming scale in our simulation box. We choose our background cosmology similar to that used in [72]: $\Omega_m = 1$, $H_0 = 70 \text{ km/s/Mpc}$ and $A_s = 4.6 \times 10^9$, where A_s is the amplitude of the primordial power spectrum. The size of the simulation volume was $(70 h^{-1} \text{Mpc})^3$. As before, we run both hot start and cold start N-body runs to compare to our simulations.

We compare the mass function from the three runs in Fig. 2.12. Once again we see that the hot run produces many spurious halos below M_{damp} . In the cold run and the hybrid SPH run, we see that the cumulative mass function flattens just below the damping scale, meaning that very few halos are produced in this mass range. At scales much smaller than the damping scale, the cold start run shows an up-turn in the mass function. This upturn is due to the halos collapsing along grid lines in simulations because of very low power on small scales. The black dash-dotted




Figure 2.12: Mass function for WDM species with $m = 60$ eV. We plot results from a hot N-body run (black), a cold N-body run (blue), and our hybrid SPH method (red). The brown dashed line represents M_{damp} . Below M_{damp} the hot run shows a number of spurious halos. At even smaller scales, both the cold start and the hybrid SPH run start showing “beads on a string” artifacts. The black dash-dotted line is an extrapolation from the flat part of the halo mass function.

line is an extrapolation of the flat part of the mass function to show how these “beads on a string” halos affect the mass function. We see the same feature even in the hybrid SPH run, where the mass function follows that from the cold start N-body closely down to the scales where the artificial halos start showing up.

From this test we conclude that while our hybrid method is highly effective in eliminating shot noise in WDM simulations where the initial thermal velocities of particles are taken into account, they cannot get rid of the “beads on a string” artifacts seen in cold start N-body simulations of WDM cosmologies. Since this hybrid method is computationally much more expensive than N-body methods, and since applying our method to WDM simulations yields no advantage over the traditional cold-start N-body simulations, we do not find a persuasive reason to use this in further studies of WDM cosmologies.

2.6 Neutrino Simulations

In this section, we discuss examples where our simulations are applied to cosmologies containing massive neutrino species. We will mostly focus on relatively low-resolution simulations of large volumes that contain a large fraction of the neutrino free streaming scale; in future work we will study nonlinear structure formation with neutrinos at higher resolution. Because of the low spatial resolution, most halos will be unresolved. However, cosmic voids can be studied using these simulations, and as we show below, massive neutrinos can have interesting effects on the large-scale clustering of highly empty voids. Indeed, from simple physical arguments we can expect that neutrinos will have a more significant impact on voids than on halos. While some neutrinos, mostly coming from the low momentum tail of the initial distribution function, do get captured by massive halos at late times, the average enclosed overdensity Δ_ν of neutrinos is still much smaller than the average enclosed overdensity of CDM, $\Delta_{\text{CDM}} \sim 200$. Any effects from neutrinos would be expected to be of the order of $f_\nu \Delta_\nu / (f_{\text{CDM}} \Delta_{\text{CDM}}) \lesssim 1\%$. This effect has been verified already [59, 90, 91]. Inside deep voids, on the other hand, the CDM overdensities are ~ -1 , while the neutrino overdensities are not as negative due to their thermal dispersion. This means that the overall mass of the neutrinos in the void can be comparable to the total mass of CDM. Therefore we might expect any effects due to neutrinos to be magnified for voids, compared to halos. Note that throughout this section, we define voids using a threshold on the total matter density, including both CDM and neutrinos. This is in contrast to voids defined using only the CDM densities, as used by [92] to study the profiles of voids in the presence of neutrinos.

In our simulations, we mostly concentrate on a single massive neutrino species with mass $m_\nu = 0.1 \text{ eV}$. Neutrinos of this mass would give energy density $\Omega_\nu < 10^{-3}$, producing relatively subtle effects on nonlinear structure formation. To illustrate this, we plot the ratio of the matter power spectrum at $z = 0$ from two simulations run with the above choice of parameters compared a CDM-only simulation in Fig. 2.14. In one of the simulations, we use our hybrid method to evolve the neutrinos, whereas in the other we treat the neutrinos as an extra set of N-body particles. Even though the neutrino power spectra are very different due to the presence of shot noise in the N-body treatment, the matter power spectrum obtained from the two simulations are very similar. Since our purpose here is to demonstrate and highlight certain effects, for the void studies, we will use




Figure 2.13: A slice through the density field in the simulation boxes at $z = 0$ for two neutrino masses - 0.5 eV and 3 eV. The simulation box side length is $175 h^{-1}$ Mpc. The left panel shows the CDM density field, the middle panel shows the density field for the 0.5 eV neutrino, while the right panel shows the density field for the 3 eV neutrino. For the lighter neutrino, on large scales, the density traces the underlying cosmic web structure laid down by the CDM component, but is much more diffuse on small scales. For the heavier neutrino, the density field is less diffuse and follows the CDM density more closely down to smaller scales.

parameters which help best illustrate these. We run our simulations for a box size of $700 h^{-1}$ Mpc with $\Omega_\Lambda = 0.7$ and $\Omega_{\text{CDM}} = 0.27$, along with $m_\nu = 0.1$ eV. To amplify neutrino effects in our simulations, we use an unphysically large neutrino number density, to give $\Omega_\nu = m_\nu n_\nu / \rho_{\text{crit}} = 0.03$. The shape of the neutrino power spectrum depends on the mass of the neutrino and the free streaming scale, while the amplitude of the effect of massive neutrinos on observables is, in general, proportional to Ω_ν . By increasing the number density of neutrinos, our goal is to amplify the effect of massive neutrinos on various observables while keeping the actual mass within the allowed parameter range. Holding the mass fixed also ensures a realistic free streaming length, which in turn ensures that the scales over which we see the effects of neutrinos in the simulations roughly match those expected in observations. If we had instead increased the mass to increase f_ν , we would push the free streaming scale down, and neutrino effects would only show up at very small scales. We stress that our choice described above is only for illustrative purposes; later on we will show results for realistic energy and number densities. For this Ω_ν , we have

$$f_\nu = \frac{\Omega_\nu}{\Omega_{\text{CDM}} + \Omega_\nu} = 0.1. \quad (2.43)$$

The Hubble constant H_0 is taken to be 70 km/sec/Mpc. We use 512^3 CDM particles and 1024^3




Figure 2.14: Relative damping of the matter power spectrum in the presence of a massive neutrino with $m = 0.1$ eV with realistic energy density ($\Omega_\nu \approx 0.0022$), compared to the CDM-only prediction. The linear theory prediction is plotted with the dashed red while the dotted blue curve represents the HALOFIT prediction. The solid black curve is the result of our hybrid simulation, while the solid green curve is from a simulation treating neutrinos as another set of N-body particles.

neutrino test particles with a 512^3 grid for hydrodynamic quantities. Initial conditions for both species are generated using CLASS at redshift $z = 49$.

Non-linear structures like voids are known to be biased tracers of the underlying matter field. The overdensity in the void field is denoted by δ_{void} . If the universe contains only one species, say CDM, we can write

$$\delta_{\text{void}}(k) = b(k)\delta_{\text{CDM}}(k) \quad (2.44)$$

where $b(k)$ is the scale-dependent bias factor. On large linear scales, $b(k)$ is independent of k for these CDM only simulations [93]. Using Eqn. 2.44, this linear (large-scale) bias factor is given by

$$b = \frac{P_{\text{void,CDM}}(k)}{P_{\text{CDM}}(k)} \quad (2.45)$$

where $P_{\text{void,CDM}}(k) = \langle \delta_{\text{void}}^*(k)\delta_{\text{CDM}}(k) \rangle$ is the cross spectrum between the void overdensity field and the underlying CDM overdensity field and $P_{\text{CDM}}(k)$ is the CDM auto-power spectrum.

If a new species is added to the matter content of the universe, biasing is no longer as simple. In this case, the void overdensity depends on both the CDM and neutrino overdensities, δ_{CDM} and δ_{ν} . Eqn. 2.44 should then be replaced by

$$\delta_{\text{void}}(k) = b_{\text{CDM}}(k)\delta_{\text{CDM}}(k) + b_{\nu}(k)\delta_{\nu}(k) \quad (2.46)$$

In terms of the total underlying matter field

$$\delta_{\text{tot}} = f_{\text{CDM}}\delta_{\text{CDM}} + f_{\nu}\delta_{\nu} \quad (2.47)$$

we can write

$$\delta_{\text{void}}(k) = b_{\text{tot}}(k)\delta_{\text{tot}}(k) \quad (2.48)$$

Even on large scales where the perturbations are linear, the neutrino power spectrum and the CDM power spectrum can be different from one another, implying that $b_{\text{tot}}(k)$ will not be scale independent. This scale dependent bias in large-scale structure can be an extremely powerful signature because it does not arise in standard cosmologies. Producing scale dependent bias on




Figure 2.15: Absolute value of the bias as a function of scale, averaged from 8 runs for $m_\nu = 0.1$ eV and $f_\nu = 0.1$. The voids were defined using the underdensity in the total matter field. The threshold for void definition was set at -0.7 . We compare the results from our simulations (black) to the other existing methods of treating the neutrinos as a linear fluid (green) and treating neutrinos as a set of particles with a different mass in N-body simulations (red). Our method and the N-body method yield results which match to within error bars, but show a strong scale dependence. The linear method shows an even stronger scale dependence. We also plot (in blue) the bias for voids defined in exactly the same manner from 8 CDM-only runs in for which the final power spectrum matches the final CDM power spectrum for the runs including neutrinos.

linear scales generally requires violating locality of the formation of biased tracers, for example due to non-gaussianity [94] or large scale modifications to gravity [95].

In the case of massive neutrinos, the scale dependence of $b_{\text{tot}}(k)$ depends on the difference of the two power spectra as a function of scale, and hence the mass of the neutrino species. For heavier neutrinos, the free streaming scale may lie well inside the simulation box. We know that on scales well above the free streaming scale, perturbations to both neutrinos and CDM evolve similarly, and so their power spectra on large scales will be identical at late times. In this case $b_{\text{tot}}(k)$ will indeed be scale independent. However for lighter neutrinos with very long free streaming lengths, the power spectra at linear scales will be very different for the neutrinos, compared to CDM. In this case, $b_{\text{tot}}(k)$ can be strongly scale dependent.

In our simulations we define the total matter overdensity as in Eqn. 2.47 and then use a spherical overdensity void finder. In our void finder, we first smooth the total density field on various scales on a grid, and then identify the points corresponding to the deepest underdensities in the smoothed density field. We use a 3-dimensional real space top hat smoothing function. If the smoothed underdensities at these points are below a given cutoff, usually -0.7 in our case, and do not fall within already identified voids, we use these points to define void centers. The smoothing scale is used to derive the size of the void. In this method, void centers cannot exist inside the volumes of other voids, but the volumes of two voids can overlap. Once the voids have been identified, we select the largest voids in the box, so that we are not affected by exclusion effects. Also, using large voids which are well-resolved in terms of the grid size we use to smooth the density field and identify void centers, ensures that the error we make in identifying the void center has minimal effect on our results. This error comes from the fact that in our method, the void centers can only lie on grid points. For small objects like halos in our simulations, which are comparable to the grid size, mis-identifying the center would lead to large errors on calculated quantities like clustering.

We run a total 8 sets of simulations with different realizations of the initial conditions sampled from the initial power spectrum. We calculate the linear bias in each of these 8 realizations, and then average over them to reduce the sample variance.

Plotting the bias defined in Eqn. 2.48 in Fig. 2.15, we find that it is indeed strongly scale dependent. To make sure that this scale dependent linear bias is not an artifact of our simulation method, we perform simulations of the same cosmology using two other methods. In one, we treat the neutrinos as an additional species with a different mass in an N-body code. In the other, we treat the neutrino fluid in a linear approximation scheme. We then use the same criterion we have defined above to find voids and calculate the large scale bias. We find that the scale dependent bias we see in our method matches the result from the N-body method to within our error bars, which denote the scatter in the bias from different realizations. As can be seen, treating the neutrinos in a linear approximation leads to an even stronger scale dependence in the bias.

To illustrate how strong this effect is, we also run a set of 8 simulations with CDM only such that the final power spectra closely matches the final power spectra in the simulations with neutrinos. We then use the same void finder on the CDM-only runs to find voids and calculate the linear bias.




Figure 2.16: Comparison of the ratio of Eulerian biases b_ν and b_{CDM} at $f_\nu = 0.1$, measured in simulations (black circles) vs. predictions from the spherical expansion model (red triangles), as a function of the threshold floor used in the void definition. The bias ratio diminishes for decreasing void thresholds, a trend quantitatively predicted by the spherical expansion model. The error bars represent the errors on the best fit values from the simulations.

This is then plotted on the same figure, and shows that it is extremely flat on linear scales.

To study the nature of the scale dependence, we perform a χ^2 fit of the large scale bias from our simulations using Eqn. 2.46 with $b_{\text{CDM}}(k)$ and $b_\nu(k)$ assumed to be scale independent. For the simulation parameters described above, we find a best-fit value of $b_{\text{CDM}} = -1.91$ and $b_\nu = -3.69$ with $\chi^2/\text{d.o.f.}$ value of 0.488. We do a similar fit for the large scale bias obtained from simulations which assumed the neutrinos to be a linear fluid, which showed a stronger scale dependence. In this case we get best-fit values of $b_{\text{CDM}} = -1.94$ and $b_\nu = -3.95$ with $\chi^2/\text{d.o.f.}$ value of 0.587. We see b_ν shows a larger difference in the two cases than b_{CDM} . Also, the absolute value of b_ν is larger in the linear approximation, meaning that the voids we have selected are rarer in the linear approximation simulations. In the linear approximation, we truncate the evolution equations at the first order in perturbed quantities like the overdensities and peculiar velocities, and so creating deep voids becomes more difficult.

To obtain an analytic understanding of these results, we use the spherical expansion model [e.g.




Figure 2.17: Behavior of the void bias from the hybrid simulations when the voids are selected using the CDM field only, and then split on the basis of the enclosed total matter overdensity for $m_\nu = 0.1$ eV and $f_\nu = 0.1$. The black curve represents $b_{\text{tot}}(k)$ for voids selected using the CDM field only with overdensity threshold of -0.7 . The red curve represents $b_{\text{tot}}(k)$ for the subsample of voids whose enclosed total matter overdensity was lower than the median enclosed total matter overdensity in the above sample. Similarly, the green curve shows $b_{\text{tot}}(k)$ for the subsample whose enclosed total matter overdensity is higher than the median.

96, 97], modified to include the effects of massive neutrinos. In our treatment so far, specifically in Eqs. 2.44 and 2.46, we have used the Eulerian definition for bias. Eulerian bias relates the abundance of tracers at a point in space and at a given redshift to the smoothed density field at the same point and the same redshift. Eulerian bias, therefore takes into account the motion of objects like voids in a fixed coordinate system, making it a natural and easy quantity to measure in simulations and observations. In simulations, for example, the simulation box provides the fixed coordinate system in which halos and voids form and advect. Lagrangian bias, on the other hand, relates the abundance of tracers like halos and voids to the over and under-densities in the initial density field. This is the natural quantity to measure when using a Lagrangian system of coordinates, where the object of interest does not move as the system evolves. This is exactly the




Figure 2.18: Effect of f_ν on the scale-dependent bias using our method for a fixed neutrino mass of $m_\nu = 0.1$ eV. We compare the case where $f_\nu = 0.1$ (black) to the case where $f_\nu = 0.05$ (red). The dashed lines show the predictions from the spherical expansion model.

case in the spherical expansion model, where the void center sits at the origin at all times. We therefore use the spherical expansion model to derive ratios between the Lagrangian bias factors b_ν^L and b_{CDM}^L , and then convert to the Eulerian biases b_ν and b_{CDM} to compare with the simulation results.

As a starting point, we assume that the ratio of the Lagrangian void biases b_ν^L and b_{CDM}^L is given by

$$\frac{b_\nu^L}{b_{\text{CDM}}^L} = \frac{f_\nu \frac{d\Delta_{\text{nl}}}{d\delta_{\text{lin}}}|_\nu}{f_{\text{CDM}} \frac{d\Delta_{\text{nl}}}{d\delta_{\text{lin}}}|_{\text{CDM}}} \quad (2.49)$$

where Δ_{nl} is the actual average nonlinear density of each species enclosed in the voids, while δ_{lin} is the prediction for the average enclosed overdensity from linear theory only. Since the neutrino perturbations are still expected to be small (and therefore close to expectations from linear perturbation theory) in the voids that we studied, we assume that $\frac{d\Delta_{\text{nl}}}{d\delta_{\text{lin}}}|_\nu = 1$.

In the spherical expansion model, we assume that neutrinos form a smooth component, and therefore expand with the background. The CDM component can be thought of as an underdense




Figure 2.19: Effect of neutrino mass on the scale dependent bias. We compare the bias when the mass of the neutrino is 0.1 eV (black) to the case when the mass is 0.8 eV (red) with $f_\nu = 0.1$ in both cases. The red dot-dashed line gives the free streaming length for the 0.8 eV neutrino. The free streaming length of the lighter neutrino lies outside (to the left of) this plot. The scale dependent bias is evident below the streaming length of each species. The dashed lines show predictions from the spherical expansion model.

universe, whose evolution is governed by the Friedmann equation. The time evolution of the scale factor a in the unperturbed background universe follows

$$\left(\frac{da}{dt}\right) = H_0 \left[\frac{1}{a}\right]^{1/2} \quad (2.50)$$

where H_0 is the present Hubble parameter and we have assumed $\Omega_m = 1$. The perturbed universe, whose scale factor we denote as χ evolves as

$$\left(\frac{d\chi}{dt}\right) = H_0 \left[\frac{f_{\text{CDM}}}{\chi} + \frac{5}{3}f_{\text{CDM}}|\delta| + \frac{\chi^2(1-f_{\text{CDM}})}{a^3}\right]^{1/2} \quad (2.51)$$

where we have used the fact $f_\nu = 1 - f_{\text{CDM}}$ is a non-clustering component of the total matter

density. The nonlinear overdensity of the CDM component at a time t is

$$\Delta_{\text{nl}}(t) = \left(\frac{a(t)}{\chi(t)} \right)^3 - 1 \quad (2.52)$$

Eqns. 2.51 and 2.50 can be solved to find the average enclosed nonlinear overdensity Δ_{nl} in Eqn. 2.52 as a function of the linear overdensity δ . This can then be substituted back into the denominator of the right-hand side of Eqn. 2.49. Once we obtain the Lagrangian biases, we need a prescription to convert to Eulerian biases to compare with the simulation results. This mapping is not straightforward in the presence of a free streaming species, i.e. the neutrinos [90]. As an approximation we assume that the advection of the voids is dominated by the advection of the CDM component, neglecting the effect of the advection of neutrinos. This approximation is equivalent to setting $b_\nu = b_\nu^L$ and $b_{\text{CDM}} = b_{\text{CDM}}^L + 1$, where the value of b_{CDM} is calculated using the best fit to the simulation results. We then compare the ratio of the Eulerian biases from this calculation to the ratio we find in simulations for different void threshold definitions. This comparison is shown in Fig. 2.16, which demonstrates that this simple model seems to quantitatively explain the trend seen in the simulations, for void definitions using overdensity thresholds in the range -0.5 to -0.8 . Below threshold of -0.8 , there are too few objects in the simulations box to get reliable statistics on the behavior of the bias. We also do not go above -0.5 in the void definition so that the voids we select always have a Lagrangian bias of $\lesssim -2$.

Since most galaxy surveys define voids using galaxy counts, which depend on the underlying CDM field, rather than the total matter field, we investigate the behavior of the bias when we select voids in the simulation using the CDM field only. Once again, we used an enclosed overdensity threshold of -0.7 . We plot $b_{\text{tot}}(k)$ for this sample of voids with the black curve in Fig. 2.17 and find that it shows very little scale dependence. From this sample of voids selected using the CDM densities only, we split into two subsamples based on the enclosed total matter overdensity - above and below the median value in the sample. When we plot the behavior of $b_{\text{tot}}(k)$ for the two subsamples, (green and red curves in Fig. 2.17) we find a very strong scale dependence in the biasing. This strategy of identifying voids using the CDM field and then splitting the sample using the total matter field could be used in surveys like DES where both galaxy counts and lensing data is available to search for this scale dependent bias. Here we have not taken into account the noise

in the lensing signal, which will, of course, be present in the lensing data from actual surveys. This noise will serve to weaken the differences in the behavior of the biases of the two samples, and in future work we will investigate the effect of lensing noise in washing out the signal shown in Fig. 2.17.

We also investigated how this scale-dependent linear bias changes as we vary f_ν defined in Eqn. 2.43. Once again we ran 8 realizations with the same cosmological parameters as above, except for Ω_{CDM} and Ω_ν , for which we used the values 0.285 and 0.015 respectively. This is equivalent to $f_\nu = 0.05$. As seen in Fig. 2.18, this leads to two effects - the voids defined similarly in both sets of simulations have a lower overall amplitude of the bias for the $f_\nu = 0.05$ case, and secondly, the scale dependence also decreases. This lower bias is a consequence of the lower mass fraction in neutrinos: with smaller f_ν , it is easier to make more voids of the size we consider, and therefore the magnitude of the bias of such objects decreases.

Next, we investigated how the neutrino free streaming scale affects the scale dependence that we observe in the void bias. To do this, we ran a set of 8 simulations for a $m = 0.8\text{eV}$ neutrino with $f_\nu = 0.1$. We keep the other cosmological parameters to be the same as earlier. We use such a heavy neutrino species to illustrate how the shape of the void bias is different above and below the free streaming length of the neutrino. For the $m = 0.8\text{eV}$, the free streaming scale is about $250h^{-1}\text{Mpc}$ - this scale is represented by the dotted red line in Fig. 2.19. At scales larger than the free streaming scale (smaller k), the CDM power spectrum and the neutrino power spectrum are the same, and there is very little scale dependence in the void bias. On the other hand, at scales below the free streaming scale where the CDM and the neutrino power spectra are significantly different, the bias starts showing a clear scale dependence. This is in contrast to the $m = 0.1\text{eV}$ neutrino, for which the free streaming scale is larger than our simulation box and therefore shows a scale dependent void bias at all the scales that we investigate. We also see that the scale dependence is steeper for the lighter neutrino particle - this is because the power spectrum for the lighter neutrino is more damped (and therefore more different from the CDM power spectrum) at any given scale.

While we plan to do a comprehensive study of this scale dependent void bias for realistic ranges of neutrino masses, we can use the spherical expansion model calculation to predict the approximate size of the bias effect for realistic neutrino number densities and minimal masses in




Figure 2.20: Scale dependent bias from the spherical expansion model for inverted hierarchy (black curve) and normal hierarchy (red curve). In the inverted hierarchy, $\sum m_\nu = 0.12$ eV, $f_\nu \sim 0.0087$ and $f_{\text{CDM}} \sim 0.9913$. In the normal hierarchy, $\sum m_\nu = 0.06$ eV, $f_\nu \sim 0.0043$ and $f_{\text{CDM}} \sim 0.9957$. Voids are defined as regions enclosing an underdensity -0.7 in both cases. We assume that $b_{\text{CDM}} = -2$ for reference.

the normal hierarchy and the inverted hierarchy. For the inverted hierarchy case, we assume the minimal sum of the neutrino masses, $\sum m_\nu \approx 0.12$ eV, so that $\Omega_\nu \approx 0.0026$. With the choice of $\Omega_\Lambda = 0.7$, we have $f_{\text{CDM}} \approx 0.9913$ and $f_\nu \approx 0.0087$. In the normal hierarchy, the minimal sum of the neutrino masses is $\sum m_\nu \approx 0.06$ eV, which gives $\Omega_\nu \approx 0.0013$. In this case, $f_{\text{CDM}} \approx 0.9957$ and $f_\nu \approx 0.0043$. We use the spherical expansion model to compute the ratio of the biases b_ν and b_{CDM} , defining voids to be enclosing a total underdensity of -0.7 . We then used linear transfer functions and power spectra at $z = 0$, generated using CLASS, to see how strongly b_{tot} varies as a function of scale on these large scales. We plot the result in Fig. 2.20 assuming that the voids have $b_{\text{CDM}} = -2$ for reference. From our $f_\nu = 0.1$ simulations, we found that voids with radius about $6 h^{-1}\text{Mpc}$ to $8 h^{-1}\text{Mpc}$ with an enclosed overdensity threshold of -0.7 had Eulerian biases close to -2 .

To quickly evaluate the ratio of the biases b_ν^L/b_{CDM}^L predicted by the spherical expansion model

over the range of realistic values of the neutrino masses as a function of f_ν and the threshold underdensity for void definition, Δ_v , we provide below a fitting function in these variables. We find that

$$\frac{b_\nu^L}{b_{\text{CDM}}^L} \approx A f_\nu^b |\Delta|^c \exp(d|\Delta|) \quad (2.53)$$

with the best fit parameters $A = 5.37 \times 10^{-6}$, $b = 1.07$, $c = -7.57$ and $d = 16.77$. This fitting formula is accurate to a few percent over the range of $0.001 < f_\nu < 0.025$ and $0.5 < |\Delta| < 0.8$.

2.7 Conclusions

In this chapter we have presented a novel, hybrid method for performing cosmological simulations in which particles have finite thermal velocities, using hydrodynamics along with standard N-body techniques used in CDM simulations. We have tested our method extensively in both the linear regime, as well as on non-linear problems.

For cosmologies with massive neutrinos, our novel method appears to accurately evolve cosmic structures at all redshifts and all scales, including both linear and nonlinear regimes. This is in contrast to traditional N-body methods, which are known to produce significant errors in the clustering of neutrinos even at low redshift (e.g. Fig. 2.5). However, for warm dark matter (WDM) cosmologies, we find that the Eulerian version of our method produces too much numerical diffusion to be useful. Instead, we find that using Lagrangian approaches to solving the fluid equations, like smoothed particle hydrodynamics, allows us to circumvent the problem of diffusion on small scales.

We have also presented a novel effect in cosmologies with massive neutrinos - the strong scale dependence in the bias of voids, even on linear scales. These voids were defined using the overall matter fields (CDM and neutrinos), rather than just the CDM fields. We note that even for voids defined using CDM field only, there is scale dependence in the large scale bias, but that effect is very weak, similar to the dependence seen in the halo bias in [62, 90]. For mass-defined halos, the scale dependence of the bias found from our hybrid method matches the results from simulations in which neutrinos were treated as N-body particles to within the error bars. However, linear theory simulations over-predict the strength of the scale dependence. For an inverted mass hierarchy of neutrinos, with $f_\nu \sim 0.01$, we find about 5% scale dependence of the void bias. For the normal

hierarchy, the effect is about 2%. A detailed study scanning all of the allowed parameter space in neutrino mass is required to predict how strongly this effect might show up in actual observations. Observationally, voids are normally identified using galaxy counts in surveys [e.g. 98, 99], rather than identification from the lensing shear field. Our results indicate that it may be worthwhile to attempt to detect voids in the mass field directly from lensing, rather than from galaxies, since mass-selected voids should exhibit the neutrino-dependent bias effect discussed above. The other strategy could be to first identify voids using the galaxy counts, and then subdividing the sample based on the mass field, which can be inferred through the gravitational lensing signal. As we have shown, the two subsamples of voids should exhibit scale dependence in their bias.

The simulations presented here used a fixed grid and therefore had low spatial resolution. In the future, we plan to combine our fluid method with adaptive mesh refinement simulation codes, which will allow us to study dark matter halos and galaxies. A particularly interesting area to study is redshift space distortions. The scale dependent growth factor in cosmologies with massive neutrinos is in principle observable in surveys measuring these redshift space distortions. While this effect has been studied using methods like perturbation theory [100] or N-body techniques [101], it will be interesting to see how the results from our hybrid simulations compare with these other approaches.

Chapter 3

Tests of Neutrino and Dark Radiation Models from Galaxy and CMB surveys

We analyze the ability of galaxy and CMB lensing surveys to constrain massive neutrinos and new models of dark radiation. We present a Fisher forecast analysis for neutrino mass constraints with the LSST galaxy survey and the CMB S4 survey. A joint analysis of the three galaxy and shear 2-point functions from LSST, along with key systematics parameters and Planck priors, can constrain the neutrino masses to $\sum m_\nu = 0.041$ eV at 1- σ level, comparable to constraints expected from Stage 4 CMB lensing. If low redshift information from upcoming spectroscopic surveys like DESI is included, the constraint becomes $\sum m_\nu = 0.032$ eV. These constraints are derived having marginalized over the number of relativistic species (N_{eff}), which is somewhat degenerate with the neutrino mass. We also explore the gain by combining LSST and CMB S4, that is, using the five relevant auto- and cross-correlations of the two datasets. We conclude that advances in modeling the nonlinear regime and the measurements of other parameters are required to ensure a neutrino mass detection. Using the same datasets, we explore the ability of LSST-era surveys to test “nonstandard” models with dark radiation. We find that if evidence for dark radiation is found from N_{eff} measurements, the mass of the dark radiation candidate can be measured at a 1- σ level of 0.162 eV for fermionic dark radiation, and 0.137 eV for bosonic dark radiation, for $\Delta N_{\text{eff}} = 0.15$.

We also find that the NNaturalness model of Arkani-Hamed et al [102], with extra light degrees of freedom, has a sub-percent effect on the power spectrum: even more ambitious surveys than the ones considered here will be needed to test such models. ¹

¹This chapter has been submitted to the Journal of Cosmology and Astroparticle Physics as A. Banerjee, B. Jain, N. Dalal and J. Shelton, *Tests of Neutrino and Dark Radiation Models from Galaxy and CMB surveys*. It is available online as [arXiv:1612.07126]

3.1 Introduction

Over the last few decades, flavor oscillation experiments e.g. [3, 4, 5] have established that neutrinos in the Standard Model are massive, and have three mass eigenstates. These experiments have accurately measured two of the mass splittings as well as the mixing angle between the various eigenstates. However, these experiments do not directly measure the absolute masses of the neutrinos, or the mass hierarchy of the three species.

The presence of massive neutrinos has non-trivial consequences for cosmology. The fraction of the total energy density of the universe that is contributed by neutrinos is proportional to the sum of the masses of the three neutrino species [10]. For most realistic neutrino masses, all three mass eigenstates were relativistic prior to last scattering of the Cosmic Microwave Background (CMB), but at least two mass eigenstates are non-relativistic today. These neutrinos contribute to the overall matter density at late times, but with little clustering below their free streaming scale. For light neutrinos, the free streaming scale can be comparable to the Hubble radius. This is unlike the Cold Dark Matter (CDM) component which clusters strongly on all scales at low redshifts to form halos and filaments. The clustering of neutrinos has the effect of damping the growth of the matter power spectrum on small scales when compared to a CDM-only case [103].

Since the damping of the power spectrum depends on the neutrino energy density, and therefore on the sum of the neutrino masses, an accurate measurement of the matter power spectrum can provide a strong constraint on the sum of the neutrino masses. Combined with the two mass splittings that have been measured from terrestrial experiments, this would allow for an accurate determination of the masses of individual eigenstates and the hierarchy.

Apart from neutrinos in the Standard Model, many theories of physics beyond the Standard Model include light, weakly interacting degrees of freedom in their particle content. While the details of these particles' interactions may be very different, as long as the interaction rates are small on cosmological time scales, the effect of these light particles on the observable matter power spectrum will be similar to that of the Standard Model neutrinos. These extra light species will damp the power spectrum on scales below their free streaming scale by an amount proportional to the energy density in that species.

Recently many models of dark sectors with varied particle content and interactions have been

proposed, [21, 104, 105, 106, 107, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117, 118] for example. Many of these models predict the existence of some form of thermal dark radiation, both bosonic and fermionic. Pseudo-Goldstone bosons are the most common example of the former, and are realized in many extensions of the Standard Model. Among fermionic dark radiation candidates, sterile neutrinos are perhaps the best-motivated and extensively studied. In all these cases, the presence of extra dark radiation in the universe would lead to a non-zero ΔN_{eff} in measurements of the CMB, i.e. a change to the effective number of relativistic degrees of freedom at CMB last scattering. However, since these particles are expected to be relativistic at the epoch of CMB, a measurement of ΔN_{eff} alone does not measure the masses of these particles. As long as the masses of these particles are light enough that they are fully non-relativistic today, the damping of the matter power spectrum is directly proportional to the mass of the dark radiation particle, and measurements of this damping on small scales can potentially constrain the mass, both for fermions and for bosons [119].

One specific example of such a model with extra light particles is the “NNaturalness” mechanism [102] proposed to solve the hierarchy problem. This model introduces N non-interacting copies of the Standard Model field content, with the Standard Model being identified as the copy with the lowest non-zero Higgs vacuum expectation value. One of the predictions of this model is that massive neutrinos from sectors close to our Standard Model could have energy densities not too much smaller than the energy density of the Standard Model neutrinos. Apart from a signature on N_{eff} at CMB, there will also be extra damping of the low redshift matter power spectrum compared to the Standard Model, which is potentially observable.

Accurate measurement of the late time matter power spectrum, therefore, will be extremely important for constraining the neutrino mass in the Standard Model, as well as for constraining light degrees of freedom from more exotic models. One probe of late-time clustering of mass is lensing of the CMB by intervening matter. The *Planck* experiment along with ground based experiments such as SPT [120] and ACT [121], have already produced lensing maps of the CMB, and future planned experiments such as the Simons Observatory and CMB Stage 4 will be able to this more accurately and down to smaller scales [12, 122].

Another powerful method for determining the late time matter power spectrum is weak lensing

measurements in large galaxy photometric surveys, like the ongoing Dark Energy Survey [123], Subaru HSC survey [124] and KiDS [125], and the upcoming surveys by the LSST [11], Euclid [126] and WFIRST [127] missions.. While measurements of galaxy-galaxy lensing and galaxy autocorrelations from these surveys do not individually measure matter auto-correlations, due to degeneracy with the unknown galaxy bias, the combination of the two on linear scales (where the bias is expected to be deterministic) eliminates the bias uncertainty [128]. Inclusion of cosmic shear measurements from these surveys further enhances the ability to pin down the underlying matter power spectrum. Using weak lensing from current and upcoming CMB and photometric galaxy surveys, as well as information about low redshifts from ongoing and future spectroscopic surveys, numerous authors have investigated the bounds placed by these measurements on the neutrino mass e.g. [129, 130, 131, 132, 133, 134, 135, 136, 137, 138, 139, 140, 141, 142, 143, 144, 145, 146, 12, 147, 148, 149, 150, 151, 152, 153, 154, 57, 155, 156, 157, 158, 159]. With the inclusion of smaller scales in successive generations of cosmological experiments, as well as lower experimental noise, these bounds have gotten progressively tighter, to the point that current data from the *Planck* experiment offers stronger bounds on the sum of neutrino masses than any terrestrial experiment, albeit with assumptions about the background cosmology.

Here, we present a Fisher forecast analysis of the ability of a survey like LSST to constrain the sum of the neutrino masses, as well as its ability to constrain other beyond Standard Model light degrees of freedom using weak lensing measurements. In §3.2, we present our method for calculating constraints on cosmological parameters, as well for differentiating between different models producing changes in the angular power spectrum C_l . In §3.3, we overview the survey parameters for LSST and CMB Stage 4 lensing experiment, as well as parameterize the systematic uncertainties in these surveys. In §3.4, we present the constraints on the sum of neutrino masses from LSST and compare it to those that will be obtained from CMB Stage 4 lensing. We also present constraints from the joint analysis of the two. In §3.5, we discuss how observables at LSST can provide constraints on dark radiation. In §3.6, we discuss the prospects of detecting models of NNaturalness using galaxy clustering and lensing at LSST. Finally, in §3.7, we summarize our findings, and discuss future avenues of study.

3.2 Method

3.2.1 Weak lensing in galaxy surveys

Galaxy surveys like LSST provide maps of galaxy distributions and shear on the sky, which can be used to construct the different 2-point correlation functions that will be used in our analysis: the galaxy-galaxy autocorrelation function C_l^{gg} , the galaxy-convergence cross spectrum $C_l^{g\kappa}$ and the convergence autospectrum $C_l^{\kappa\kappa}$. Here we make use of the simple relation between shear spectra and convergence spectra. Using the Limber approximation these spectra are constructed from the underlying 3-dimensional power spectra [160]:

$$C_l^{x_i x_j} = \int dz \frac{H}{D_A^2} W_i(z) W_j(z) P^{s_i s_j}(k = l/D_A; z), \quad (3.1)$$

where x_i stand for g or κ , while s_i stand for the underlying 3 dimensional source fields. D_A is the angular diameter distance, and W_i are the weighting functions in redshift space. For the galaxy number fluctuations, the three dimensional source field is the fluctuations in the 3-dimensional number density:

$$s(\mathbf{r}; z) = \frac{\delta n_V}{\bar{n}_V}. \quad (3.2)$$

The weighting function for galaxy fluctuations is given by

$$W_g(z) = \frac{D_A^2 \bar{n}_V}{H \bar{n}_A}, \quad (3.3)$$

where the normalization factor \bar{n}_A is chosen so that $\int W_g(z) dz = 1$. For the shear field the 3-dimensional source field is the fluctuation of the matter density:

$$s(\mathbf{r}; z) = \frac{\delta \rho_m}{\rho_m}, \quad (3.4)$$

and the weighting function is

$$W_\kappa(z) = \frac{3}{2} \Omega_m \frac{H_0}{H} \frac{H_0 D_{OL}}{a} \int_z^\infty dz' \frac{D_{LS}}{D_{OS}} W_g(z'), \quad (3.5)$$

where D_{OL} stands for the angular diameter distance to the lens, D_{OS} is the angular diameter distance to the source, and D_{LS} is the distance between the lens and the source. We use the publicly available CAMB code [161] to generate the various power spectra that go into our analysis.

To construct the covariance matrix for the C_l , we assume that the different l are uncorrelated and so the covariance matrix is diagonal in l . We also assume that up to l_{\max} , we are roughly in the linear regime, where Gaussian statistics are valid, and all n -point functions can be broken down into products of various two point functions given by the C_l . To take into account the shape noise in the shear spectra and the shot noise in the galaxy spectra we define

$$\tilde{C}_l^{x_i x_j} = C_l^{x_i x_j} + N_l^{x_i x_j}, \quad (3.6)$$

where x stands for g or κ and i is used to label different redshift bins. The noise terms are given by

$$N_l^{g_i g_j} = \frac{\delta_{ij}}{\bar{n}_i}, \quad (3.7)$$

$$N_l^{\kappa_i \kappa_j} = \frac{\delta_{ij} \gamma_{\text{rms}}^2}{\bar{n}_i}, \quad (3.8)$$

$$N_l^{g_i \kappa_j} = 0. \quad (3.9)$$

Here \bar{n}_i represent the number counts of lens and source galaxies in each redshift bin in units of sr^{-1} . For the shape noise contribution, we use a value of $\sqrt{\gamma_{\text{rms}}^2} = 0.22$.

In the linear regime all covariances of the power spectra can be written as products of the power spectra themselves. Therefore the different elements of the covariance matrix \mathcal{C}_l can be written as:

$$[\mathcal{C}_l]^{ij,kl} \equiv \tilde{C}_l^{x_i x_k} \tilde{C}_l^{x_j x_l} + \tilde{C}_l^{x_i x_l} \tilde{C}_l^{x_j x_k}. \quad (3.10)$$

Using this covariance matrix, we construct different components of the LSST Fisher matrix

$$F_{\alpha\beta}^{\text{LSST}} = f_{\text{sky}} \sum_l (2l+1) \sum_{ijkl} \frac{\partial C_l^{x_i x_j}}{\partial p_\alpha} [\mathcal{C}_l]_{ij,kl}^{-1} \frac{\partial C_l^{x_k x_l}}{\partial p_\beta}, \quad (3.11)$$

where p_α represent the model parameters, and f_{sky} is the fraction of the sky covered in the survey,

which we take to be 0.5 for LSST.

To use information on the parameters from *Planck*, we use the covariance matrices available in the *Planck* Legacy Archive for the baseline model. We ensure that we do not include *Planck* constraints coming from CMB lensing within the *Planck* experiment. The total Fisher matrix is obtained by adding together the resultant *Planck* Fisher matrix to the Fisher matrix we derive for the LSST experiment

$$F = F^{\text{LSST}} + F^{\text{Planck}}. \quad (3.12)$$

The constraint on parameter α is then given by

$$\sigma(p_\alpha) = \sqrt{(F^{-1})_{\alpha\alpha}}. \quad (3.13)$$

3.2.2 CMB lensing

For the CMB lensing analysis, we consider the lensing power spectrum C_l^{dd} where \mathbf{d} , the deflection field, is the gradient of the lensing potential $\mathbf{d} = \nabla\phi$. The noise level N_l^{dd} for a given set of experimental sensitivities can be estimated following [162], and was calculated using QUICKLENS [163]. The covariance matrix for the CMB lensing power spectrum is then

$$\mathcal{C}_l = 2 \left(C_l^{dd} + N_l^{dd} \right)^2, \quad (3.14)$$

and the elements of the Fisher matrix are given by

$$F_{\alpha\beta}^{\text{CMB}} = f_{\text{sky}} \sum (2l+1) \frac{\partial C_l^{dd}}{\partial p_\alpha} \mathcal{C}_l^{-1} \frac{\partial C_l^{dd}}{\partial p_\beta}. \quad (3.15)$$

Once again, we use *Planck* priors, where lensing information from *Planck* is not used. This is done by adding together the two Fisher matrices, as in §3.2.1.

3.2.3 Model differentiation

To calculate the statistical significance of distinguishing between two models which are not connected by a parameter which can be varied smoothly, we use the following procedure [74]. Let the difference in the various predicted power spectra from the two models be denoted by $\delta C_l^{x_i x_j}$. The

χ^2 difference between the two models is then given by

$$\chi^2 = f_{\text{sky}} \sum_l (2l+1) \delta C_l^{x_i x_j} [\mathcal{C}]_{ij,kl}^{-1} \delta C_l^{x_k x_l}. \quad (3.16)$$

A part of this χ^2 difference can be taken into account by varying the continuous parameters in the fiducial model. The change of the observable power spectra given a change in the parameters of the fiducial model is $\delta C_l^{x_i x_j} = (\partial C_l^{x_i x_j} / \partial p_\alpha) \delta p_\alpha$. Taking this variation into account, the new χ^2 is given by

$$\chi^2 = f_{\text{sky}} \sum_l (2l+1) \left[\delta C_l^{x_i x_j} + \delta p_\alpha \frac{\partial C_l^{x_i x_j}}{\partial p_\alpha} \right] [\mathcal{C}]_{ij,kl}^{-1} \left[\delta C_l^{x_k x_l} + \delta p_\beta \frac{\partial C_l^{x_k x_l}}{\partial p_\beta} \right]. \quad (3.17)$$

By setting $\partial \chi^2 / \partial p_\alpha = 0$ in the above equation, we calculate the δp_α which needs to be made about the fiducial set of parameters p_α to obtain the minimum χ^2 :

$$\delta p_\alpha = - (F)_{\alpha\beta}^{-1} \frac{\partial C_l^{x_i x_j}}{\partial p_\beta} [\mathcal{C}]_{ij,kl}^{-1} \delta C_l^{x_k x_l}, \quad (3.18)$$

where F is the Fisher matrix of the fiducial model. The final χ^2 is then calculated using (3.17).

3.3 Survey parameters and systematics

3.3.1 LSST

To get the number counts of sources in an LSST like survey, we use a redshift distribution of source galaxies given by the following form:

$$\frac{dn}{dz} \propto z^{1.2} \exp\left(-\frac{z}{0.5}\right), \quad (3.19)$$

with a total number density $n_{\text{source}} = 30 \text{ arcmin}^{-2}$. For lens galaxies, we will use redMaGiC type of galaxies [164] which have a redshift distribution $dn/dz \propto \chi(z)^2 / H(z)$, and total number density of lenses $n_{\text{lens}} = 0.25 \text{ arcmin}^{-2}$.

In our calculations, we consider different redshift binnings - we start from 1 lens redshift bin and 1 source redshift bin, and go up to 6 lens redshift bins and 6 source redshift bins. We show the




Figure 3.1: Expected source galaxy counts (per square arcmin) as a function of redshift in the LSST survey is plotted with the black curve. The solid red lines indicate the photometric redshifts z_{\min} and z_{\max} of the 6 lens bins used in our analysis. The dashed blue lines represent the z_{\min} and z_{\max} for the 6 source bins. The blue lines have been slightly displaced for clarity.

	z_{center}	l_{max}		z_{center}	l_{max}
Lenses	0.40	210	Sources	0.85	200
	0.55	240		1.00	220
	0.70	330		1.15	250
	0.85	440		1.30	280
	1.00	570		1.45	320
	1.15	720		1.60	370

Table 3.1: List of redshift bin centers and l_{max} used for each redshift bin for the LSST survey.

bin centers and widths for this last case in Fig. 3.1. For each lens redshift bin center, we calculate k_{max} such that $\frac{k^3 P(k)}{2\pi^2} \Big|_{k_{\text{max}}} \approx 0.2$. Using this, we obtain the highest multipole $l_{\text{max}} = k_{\text{max}} \chi$, which we use in our analysis. This is done to ensure that that even for the highest multipoles in each bin, we are in the regime where perturbations can still be treated as being roughly linear, and assumptions of the independence of different l modes and linear biasing are valid. As expected, l_{max} increases with redshift, so we can go out to smaller scales at higher redshifts. This allows for tighter constraints on the different cosmological parameters, which affect the shape of the power spectrum, along with the overall amplitude. We tabulate the l_{max} for each of our bins in Table 3.1.

We also note that we use the linear matter power spectrum for all our calculations. Since




Figure 3.2: Comparison of the C_l^{gg} and $C_l^{\kappa\kappa}$ power spectra at different redshifts to the shot noise and shape noise levels at those redshifts. The solid lines plot the power spectrum, while the dashed line of the same color plots the shot noise or shape noise for that power spectrum at that redshift expected at LSST. The dot-dashed lines represent the value of l_{\max} for each redshift bin. The spectra are sample variance dominated out to $l \sim 400$.

we restrict ourselves to the scales for which $\frac{k^3 P(k)}{2\pi^2} \Big|_{k_{\max}} = 0.2$, using the linear power spectrum, as opposed to the full nonlinear power spectrum is a valid choice, except for some scales near k_{\max} and consequently l_{\max} . The differences are expected to be small, and using the linear power spectrum will always yield a conservative estimate on the constraints.

In Fig. 3.2, we plot the power spectra and the associated noise levels for the assumed number counts of lenses and sources in some of the redshift bins that we use in our calculations. In the left panel of Fig. 3.2, we plot the galaxy-galaxy autospectra C_l^{gg} for four of the lens redshift bins, with the solid lines. We represent the shot noise level in each bin using the dashed lines of the same color. Similarly, in the right panel, we plot the shear autospectra $C_l^{\kappa\kappa}$ for four of the source redshift using solid lines. We also indicate the shape noise for these bins using dashed lines of the same color.

As mentioned, the lens galaxy counts we use throughout our analysis are chosen to approximate the distribution of redMaGiC galaxies [164]. The galaxy-galaxy autocorrelations for the lens bins are also calculated using the same subsample of galaxies. Since these galaxies form only a fraction of all the galaxies in those redshift bins, the level of shot noise in our measurements, as expressed in Eq. 3.7, will be higher than the case where all galaxies are used. However, as Fig. 3.2 shows, shot noise is sub-dominant compared to the signal covariance over the range of scales we consider,

meaning that the increased shot noise from the redMaGiC subset should not degrade the errors significantly. The advantage of using this subset is that the redshifts of redMaGiC galaxies may be determined to a very high level of accuracy using photometry alone, allowing us to effectively ignore lens redshift errors as a source of systematics. For the lens redshift bins, therefore, we only consider the galaxy bias of each bin b_i as sources of systematics, and include these as nuisance parameters in our Fisher matrix analysis.

For the source galaxies, however, we use the entire galaxy population from those redshift bins. While using these high number densities helps reduce the shape noise, we need to consider multiple sources of systematics [165]. Amongst these, the most significant systematic that we need to account for is the photometric redshift uncertainties of the source galaxies, as these are not as well measured as those in the redMaGiC sample. Another important systematic that we need to consider is the shear calibration. Here we allow for a multiplicative error arising from calibration errors in the shear and photo-z. We do not attempt to model additive, scale dependent uncertainties in the shear or clustering as a reasonable analytical model or even level of uncertainty is not available. To account for these sources of systematic errors, we introduce the nuisance parameters, m_i , one for each source bin. We assume that these effects can be parameterized by allowing for an overall rescaling of the shear measurements: $\kappa_i \rightarrow \kappa_i(1 + m_i)$. We will consider this parameterization for the source uncertainties throughout the study. Note that the relation between our m_i and photo-z bias is not linear, and is redshift dependent. But it is a reasonably good approximation to capture both shear and photo-z bias into a multiplicative parameter for our forecasting purposes.

Since we are using a single nuisance parameter m_i , per bin, to account for both shear calibration and redshift uncertainties, as the calibration of the shear in LSST gets better, the m_i will mostly encode our uncertainty on the photometric redshifts of the source galaxies.

We note that our analysis assumes that other sources of systematic errors are subdominant to statistical errors and the systematic uncertainty due to calibration errors that we have modeled. These include intrinsic alignments of galaxies [166, 167], and additive errors in the shear measurements due to residuals from the Point Spread Function (PSF) correction [168]. For intrinsic alignments (in particular the GI alignments), a recent study [169] has shown that the systematic uncertainty is well below statistical errors for an analysis similar to ours (see also [170, 171, 167, 172]).




Figure 3.3: Noise level N_l^{dd} in the deflection power spectrum C_l^{dd} from CMB Stage 4 lensing for the assumed survey parameters. The lensing signal is sample variance dominated out to $l \sim 1000$.

Additive errors in the shear are not generally modeled in forecast studies as they are exceedingly difficult to anticipate and various mitigation strategies are employed to deal with them in the measurement and analysis.

3.3.2 CMB lensing

While a number of specifications have been proposed for the survey parameters of CMB Stage 4 experiments, here, we will use the following specifications. We consider fractional sky coverages $f_{\text{sky}} = 0.75$, $f_{\text{sky}} = 0.5$ and $f_{\text{sky}} = 0.25$. The beam size is assumed to be $1'$ in all cases and we assume an overall experimental sensitivity of $0.58 \mu\text{K}\text{-arcmin}$. In Fig. 3.3, we plot the deflection power spectrum C_l^{dd} , and the noise level N_l^{dd} on this observable for the survey parameters that we assume.

For CMB lensing we include information from multipoles up to $l_{\text{max}} = 3000$, while using a lower cutoff $l_{\text{min}} = 30$. Since the lensing kernel for CMB peaks at around $z = 2$, it is justified for us to go up to this high l_{max} while still using assumptions of linearity for the power spectra and their covariances.

3.4 Neutrino mass constraints

We assume a fiducial model with $(\tau, n_s, \ln[10^{10} A_s], \sum m_\nu, N_{\text{eff}}, \Omega_m, \Omega_b, \Omega_\Lambda, h, w, b_i, m_i) = (0.066, 0.967, 3.15, 0.06, 3.046, 0.3, 0.05, 0.7, 0.7, -1, 1.5, 0)$ where b_i stands for the bias in different redshift bins - with one bias parameter per redshift bin. While there may be additional signatures in the galaxy bias [90, 91], we do not attempt to model them or use the additional information they may provide - our forecasts are therefore conservative. The bias parameters will be marginalized over when deriving constraints on various cosmological parameters like the sum of the neutrino masses $\sum m_\nu$ and the dark energy equation of state w . In our analysis, we do not assume any priors on the bias parameters - the constraints on these parameters come from data only. Apart from the bias parameters, we also marginalize over the shear uncertainties m_i defined in § 3.3, and whose fiducial value we assume to be 0. Our fiducial constraints on $\sum m_\nu$ and w are derived without assuming priors on the nuisance parameters m_i . We discuss the effect of placing priors on these parameters later in this section. We note that in all our calculations, we have assumed $\Omega_k = 0$, that is, curvature is neglected. Further, we have assumed that the equation of state of dark energy w is time invariant.

3.4.1 LSST constraints on neutrino mass

We use the formalism described in § 3.2.1 to obtain constraints on cosmological parameters for different number of lens and source bins at LSST. We find that for bins of fixed width, increasing the number of bins improves the constraints on the neutrino mass. This happens for two reasons - by increasing the number of bins, we increase the redshift coverage of the sources and lenses. Secondly, using more bins allows for the measurement of higher number of cross-correlations. This improvement in the constraints is summarized in Table 3.2, and is illustrated in Fig. 3.4, where we plot the marginalized probability distributions for $\sum m_\nu$ and w for different number of bins used in the analysis. In Fig 3.5 we plot the 1- σ and 2- σ confidence intervals for pair of the parameters $(\sum m_\nu, w)$ for two choices of source and lens bin numbers - 4 and 6. While the gain in going from one source and one lens bin to 4 of each gives a large improvement in the constraints, the constraints start to saturate as we add more bins. This is because the finer binning leads to smaller galaxy counts in each bin, raising the shot noise level. Using 6 lens and 6 source bins yields a

	$\sigma(\sum m_\nu)$ (eV)	$\sigma(w)$
$N_S = 1, N_L = 1$	0.093	0.069
$N_S = 4, N_L = 4$	0.052	0.028
$N_S = 6, N_L = 6$	0.041	0.020
$N_S = 6, N_L = 6$ (+ DESI[146])	0.032	0.017
$N_S = 6, N_L = 6$ (+ DESI[146] + N_{eff} prior))	0.028	0.016

Table 3.2: Forecasts of $1\text{-}\sigma$ constraints on the neutrino mass, and the dark energy equation of state w at LSST for different number of lens and source bins, and different priors. The N_{eff} prior used was 0.03.




Figure 3.4: Marginalized probability distributions for the sum of neutrino masses (left), and dark energy equation of state w (right), from LSST. The blue, red and black curves correspond to $N_L = N_S = 1, 4, 6$ respectively. Increasing the number of redshift helps extract more tomographic information, but this gain saturates as the individual bins become too thin.

constraint on the sum of the neutrino masses $\sigma(\sum m_\nu) = 0.041 \text{ eV}$. From the same analysis, we obtain a constraint of 0.020 on the dark energy equation of state w .

Since we do not assume any priors on the biases of the lens galaxies, we check how well the biases are constrained by data. We find that the bias parameters in each lens redshift bin is constrained at about 2%, with the bias parameters for the low redshift lens bins being slightly better constrained than the bias parameters for the higher redshift bins. This happens because of two opposing effects. The lower redshift bins have a higher signal to noise ratios in their galaxy clustering power spectra. On the other hand, we go out to a higher l_{max} for the higher redshift, increasing the sensitivity of those redshift bins to the bias parameters. These two effects roughly cancel each other out to provide similar constraints on the bias parameters for all redshift bins we consider.

Apart from LSST, other future cosmological surveys will also be sensitive to the effects of




Figure 3.5: 1- σ (solid lines) and 2- σ (dashed lines) confidence intervals on 2-d subspace of the parameters $\sum m_\nu$ and w from LSST. The red curves are the results from 4 source bin and 4 lens bin analysis. The black curves represent the results when 6 source bins and 6 lens bins are used.

neutrino mass on galaxy clustering and weak lensing observables. It is, therefore, useful to check how the constraints on the cosmological parameters depend on survey specifications, such as the sky coverage, redshift depth, and number densities of galaxies. Apart from the LSST survey specifications, we consider two other surveys - one with specifications similar to EUCLID [126], and the other with survey parameters similar to the WFIRST survey [127]. For the EUCLID-like survey we assume $f_{\text{sky}} = 0.375$, an average source redshift $z = 0.7$, and number density $n_{\text{source}} = 20 \text{ arcmin}^{-2}$. The 1- σ constraint on $\sum m_\nu$ for this survey is 0.060 eV. For the WFIRST-like mission, we assume $f_{\text{sky}} = 0.0675$, an average source redshift $z = 1.4$, and number density $n_{\text{source}} = 40 \text{ arcmin}^{-2}$, and we find a 1- σ constraint of 0.067 eV on $\sum m_\nu$. Note that the constraints for WFIRST are not very different from the other surveys even though it has small f_{sky} due to the greater survey depth. Therefore, if WFIRST continues beyond its nominal three year mission, the sky coverage f_{sky} could become larger, and the constraints will improve accordingly. We find that our constraint forecasts are weaker compared to a similar forecast for the EUCLID mission in [145] in the case where the authors assume that the linear galaxy bias is known exactly for the lens galaxies. When this assumption is relaxed, the authors find a similar constraint on the neutrino mass as the ones presented here. For the gain by combining all three surveys, see also [173].

To account for systematic errors in our constraint forecasts at LSST coming from the shear bias of the source redshift bins, along with the photometric redshift uncertainties, we introduce the nuisance parameters m_i for each source redshift bin, as mentioned in § 3.3. As with the bias parameters, these are marginalized over to yield constraints on the cosmological parameters of interest. Note that photometric redshift uncertainty leads to biased estimates of the distances to galaxies, which are not exactly degenerate with shear calibration. But it is sufficiently accurate for the purposes of this study to fold the two biases into a single bias parameter per redshift bin.

Unlike with the galaxy bias parameters, which must be measured from the same dataset used for cosmology, we test the effect of priors on these shear bias parameters as they can be estimated with image simulations or high resolution imaging. We then study how the constraints on the sum of neutrino masses and the equation of state for dark energy vary with the imposed prior. We find that when flat priors are imposed on the m_i , as is the case for our fiducial results, the constraint on the neutrino mass from LSST weak lensing obtained above degrades by $\sim 17\%$, whereas the constraint on w degrades by only $\sim 5\%$, compared to the case where these parameters are completely ignored.

We illustrate how the constraints on the cosmological parameters change as we impose stronger priors in Fig. 3.6. For an imposed prior of 0.02, which may be achievable in LSST, we find that the neutrino mass constraint degraded by only about 3% compared to the case where the effect of these parameters is ignored. The degradation on the constraint on w is once again, even smaller - roughly 1%. Obtaining percent-level priors on shear bias would require extensive and realistic image simulations, as planned by LSST, and/or deep, high resolution imaging of a large enough subset of the source galaxies - e.g. with space based imaging by WFIRST. An alternative is to actually use CMB lensing to calibrate the shear - this approach was investigated by [174] and more recently in some detail for a joint analysis of LSST and CMB Stage 4 lensing by [169]. They find a calibration of shear bias in LSST at the percent level is indeed feasible. Note however that this still leaves redshift bias as a systematic, so for our fiducial results we use flat priors on the m_i parameters.

The behavior of the constraints as seen in Fig. 3.6 suggests that significant self-calibration is possible even with no external priors. This is due to several factors: the combination of auto- and




Figure 3.6: Degradation of the 1- σ constraints on the parameters as a function of the prior assumed on m_i , the shear calibration of the source bins at LSST. The constraints change very weakly with the value of the prior, especially for w .

cross-spectra across redshift bins gives $N(N+1)/2$ shear spectra for N redshift bins, while there are only N shear bias parameters. Additional self-calibration comes from galaxy-shear cross-spectra which have a different scaling with redshift, and some information from galaxy-galaxy autospectra alone, which are not affected by the shear bias and photometric redshift errors of the source redshift bins. Even though the galaxy-galaxy autospectra depend on the *a priori* unknown galaxy bias of the lenses, the comparatively small error bars that will be achievable in LSST, coming mainly from the large sky coverage f_{sky} , means that these autospectra are also sensitive to parameters which modify the shape of the various C_l . Massive neutrinos change both the amplitude and the shape of the 3 dimensional matter power spectrum, along with the different C_l . Changes in the dark energy equation of state changes the cosmological growth factor, thereby changing the amplitude of the various C_l^{gg} . In addition, the dark energy equation of state also affects the distance-redshift relation, which in turn, affects the positions of the BAO peaks in C_l^{gg} . Since this second effect is not degenerate with a change in the different bias parameters, the shapes of C_l^{gg} from different redshift bins provides strong constraints on w .

3.4.2 Comparison to CMB Stage 4 lensing

Next, we compare the constraints on the parameters obtained from weak lensing in galaxy surveys to the constraints derived from CMB Stage 4 lensing, with the inclusion of *Planck* priors. This comparison is illustrated in Fig. 3.7. The black solid line represents the results from LSST, while the red lines represent the results from CMB Stage 4 experiments, for different sky coverages. We find that the constraints on the sum of neutrino masses from the two experiments are very similar, when we consider the most optimistic scenario for CMB Stage 4 lensing ($f_{\text{sky}} = 0.75$), where the latter gives a constraint of $\sigma(\sum m_\nu) = 0.046 \text{ eV}$.

It is interesting to compare the sensitivity of a CMB lensing survey and a weak lensing photometric to the sum of neutrino masses under the assumption that they cover the same fraction of the sky, and that they are sample variance dominated out to a similar range in multipole l . As can be seen from Fig. 3.2 and Fig. 3.3, CMB lensing is sample variance dominated out to $l \sim 1000$, whereas with the number densities of lenses and sources assumed for LSST, the signal is sample variance dominated out to $l \sim 400$. For this comparison, therefore, we tune the number densities of lenses and sources in LSST to 4 times their fiducial value, and then compare the constraints on $\sum m_\nu$ from the photometric galaxy survey and CMB lensing using $l_{\text{max}} = 1000$ and $f_{\text{sky}} = 0.5$ in both cases. Under these assumptions, we find $\sigma(\sum m_\nu) = 0.018 \text{ eV}$ for the photometric survey and $\sigma(\sum m_\nu) = 0.055 \text{ eV}$ for CMB lensing. This shows that the tomographic information in a photometric survey allows us to sample more modes, and therefore have a higher Signal to Noise Ratio (SNR) compared to a CMB lensing survey where the source redshift and the lensing kernel is fixed by the CMB last scattering surface. Quantitative comparisons of the SNR for a photometric survey like LSST and CMB lensing has been performed in [169].

For CMB Stage 4 lensing, we find that our constraint on $\sum m_\nu$ is weaker by about a factor of 2 compared to those in [12]. There are two main reasons behind this difference. First, [12] assumes priors on other cosmological parameters coming from DESI BAO measurements, which will be much tighter than the *Planck* priors that we use in this work. Secondly, N_{eff} , which has large degeneracies with the neutrino mass is included in our analysis, while the constraints quoted in [12] assume that extremely accurate measurements of the CMB primaries from the Stage 4 experiments will constrain N_{eff} independently. Refs. [147, 135, 148] also find very similar constraints on the

neutrino mass when using DESI BAO priors in addition to CMB S4 lensing.

To check how our estimates improve with the inclusion of low redshift information coming from DESI, we use the parameter forecasts provided in [146] in our analysis. This extra information is especially helpful in tightening the constraint on the parameters Ω_m and w , coming from precise measurements of the BAO feature at low redshifts. We find that the $1\text{-}\sigma$ constraint on the sum of the neutrino masses from LSST clustering and lensing is 0.032 eV when information from DESI is included, compared to our fiducial result of 0.041 eV . For CMB Stage 4 lensing, the improvement is even more marked. With the extra constraints from DESI, Stage 4 CMB lensing can constrain the sum of neutrino masses to 0.029 eV at the $1\text{-}\sigma$ level. The larger improvement in the CMB S4 bound is understandable since the CMB lensing kernel is peaked near $z \sim 2$, and therefore the information from low redshift provided by DESI is mostly complementary to the information contained in CMB lensing. On the other hand, for LSST, some of the information about low redshifts is already included as the lens bins in our analysis extend down to $z = 0.4$. So adding in the DESI priors to LSST analysis do not improve the existing bounds on the neutrino mass by a lot.

Next, we add additional priors on the parameter N_{eff} based on the forecasts for CMB Stage 4 experiments. We assume that ΔN_{eff} will be measured using the CMB primaries at a $1\text{-}\sigma$ error level of 0.03. When this prior is included in our calculation for the LSST clustering and lensing, the $1\text{-}\sigma$ constraint on the sum of neutrino masses tightens further to 0.028 eV . Similarly, when this prior on N_{eff} is included in the CMB Stage 4 lensing estimates, the constraint on the neutrino mass becomes 0.023 eV . These numbers are then comparable to the constraints on the neutrino mass forecast in Ref. [12].

When compared to the current constraints on the neutrino mass coming from CMB lensing in the *Planck* experiment [13], we find that both LSST weak lensing and CMB Stage 4 lensing will improve the bounds on the neutrino mass by almost an order of magnitude. This is especially true once low redshift information from a DESI-like experiment is included for the CMB lensing analysis. For the *Planck* value, we have used the constraint obtained from the CMB lensing from *Planck* itself combined with *Planck* primaries only, with no additional external datasets. The $1\text{-}\sigma$ constraint on the neutrino mass from this dataset only is $\sim 0.29\text{ eV}$.




Figure 3.7: Marginalized probability distributions for the sum of neutrino masses, and dark energy equation of state w from CMB Stage 4 lensing (red curves) compared to LSST (black curve). The solid red curve corresponds to $f_{\text{sky}} = 0.75$, the dashed red curve corresponds to $f_{\text{sky}} = 0.5$ and the dot-dashed curve corresponds to $f_{\text{sky}} = 0.25$ for CMB Stage 4 experiments. While neutrino mass constraints from the two experiments are similar, w can be constrained more tightly by the LSST survey.

Since LSST weak lensing and CMB Stage 4 lensing are sensitive to very different systematics, the fact that they are expected to provide very similar constraints on the sum of neutrino masses is an important result. Since the errors from the two experiments do not correlate, the statistical significance of a detection from one of experiments will be greatly enhanced when combined with the data from the other.

While LSST weak lensing and CMB stage 4 lensing provide similar constraints on the sum of the neutrino mass, the constraint on w from LSST is much stronger than the constraint from CMB Stage 4, even when we consider the most optimistic case for the latter. The CMB lensing kernel peaks at around $z \sim 2$, when dark energy forms a negligible fraction of the energy budget. On the other hand, LSST weak lensing is sensitive to much lower redshifts when dark energy starts to dominate. Further, the equation of state for dark energy affects the growth rate of the power spectrum. Since we use spectra from multiple redshift bins in LSST, the growth rate, and in turn, w , are better constrained than in CMB lensing.

3.4.3 Combination of LSST and CMB S4

It is also possible to combine measurements from LSST and CMB Stage 4 lensing assuming that the survey windows of the two overlap each other. We use data from these overlapping surveys

to measure the cross-spectra $C_l^{g\kappa_{\text{CMB}}}$, along with the galaxy autospectra C_l^{gg} from LSST and the shear-shear autospectrum $C_l^{\kappa_{\text{CMB}}\kappa_{\text{CMB}}}$ from CMB lensing. Using CMB lensing κ_{CMB} measurements reduces some of the systematic uncertainties, since the multiplicative shear bias and source redshifts uncertainties present in the LSST lensing measurements no longer affect the results.

For this analysis, we use the same lens redshift bins for LSST, as well as the same experimental sensitivity for as mentioned in §3.3. We assume $f_{\text{sky}} = 0.5$ for both experiments with complete overlap of survey windows. We find that this combination provides 1- σ constraints of $\sigma(\sum m_\nu) = 0.031$ eV and $\sigma(w) = 0.016$. Both these constraints are slightly stronger than the ones obtained from LSST only and the most optimistic case for CMB Stage 4 lensing, as shown in Fig. 3.8.

One can also check how this joint constraint on the neutrino mass coming from LSST and CMB lensing improves when priors from the DESI experiment is included. Once again, we account for the extra information from DESI by adding stronger priors, especially on Ω_m and w , whose values are given in [146]. We find that including these stronger priors improves the 1- σ constraint on the sum of neutrino masses to 0.020 eV. This bound is competitive with the bounds presented in [12], even though in our analysis, we marginalize over one extra cosmological parameter in N_{eff} . On the other hand, this estimate may be a little optimistic given that we have assumed complete overlap of the survey volume between the LSST and CMB Stage 4 experiment.

While the constraints on the cosmological parameters do improve when shear measurements from CMB stage 4 are used, the improvement in the constraints is not dramatic unless extra information from a survey like DESI is included. On the other hand, adding CMB lensing to LSST helps with systematics of the LSST lensing measurements, especially in terms of constraining the multiplicative shear biases m_i of the different redshift bins. This has been studied in detail in Ref. [169]. Therefore, even for cases where the overlap of survey volumes of the experiments is not perfect, useful information can be obtained by looking at the cross-correlations of observables from the two.

3.5 Dark radiation

Thermal dark radiation is a natural and generic prediction of many theories beyond the standard model. For instance, pseudo-Nambu-Goldstone bosons (pNGBs) are naturally light and are realized




Figure 3.8: Marginalized probability distributions for the sum of neutrino masses $\sum m_\nu$, and dark energy equation of state w for joint LSST and CMB lensing analysis (green curve) compared to those coming from LSST only (black curve) and CMB lensing only (red curve). The constraints improve somewhat marginally compared to the LSST only results.

in many extensions of the SM, arising from the spontaneous breaking of a Peccei-Quinn symmetry (axions) [175, 176], lepton number symmetry (majorons) [177, 178], family symmetry (familons) [179], or dark number symmetry [180]. Given suitably large reheating temperatures and low symmetry breaking scales, these pGNBs can be thermally populated in the early universe. Beyond these minimal models, string compactifications often yield a proliferation of axion-like particles with a much expanded range of masses and interaction strengths, making pNGBs a natural and generic source of thermal bosonic dark radiation [181, 182, 183]. Such stringy constructions can also yield dark $U(1)$ gauge bosons [184, 185] with sub-eV masses, and possibly also fermionic dark radiation in the form of photini [186].

Perhaps the leading candidate for thermal fermionic dark radiation is a sterile neutrino. One way to thermalize sterile neutrinos is to add new interactions in the neutrino sector that are motivated by unification and by measurements of neutrino mixing [187, 188, 189, 190]. Similar to NNaturalness, thermal relic populations of both sterile neutrinos and dark photons can also arise in *mirror sectors*, where the matter content of the SM is replicated, wholly or in part, in a hidden sector. This long-standing idea, reviewed in [191], has been motivated by parity restoration as well as asymmetric dark matter and solutions to the hierarchy problem [102, 192].

More generally, hidden sectors constitute a generic possibility for physics beyond the standard model. Cosmologically, thermal dark sectors are very well motivated as a source of DM. Such

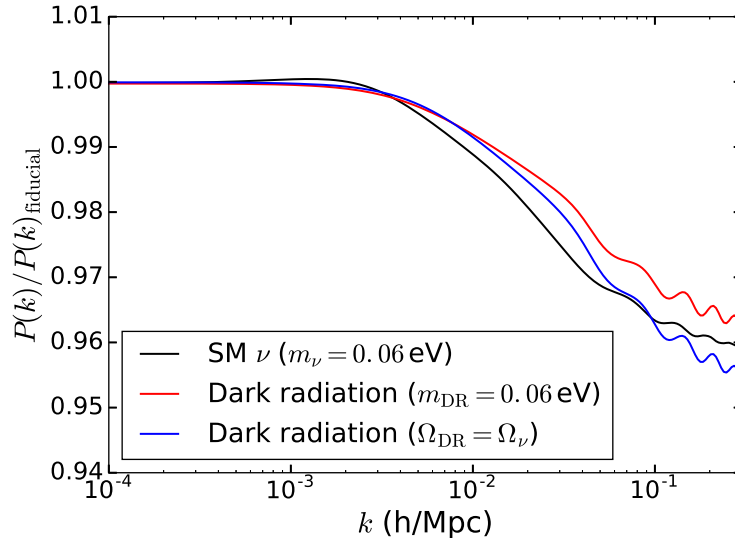


Figure 3.9: Damping in the power spectrum caused by a Standard Model neutrino (black curve) with mass $m_\nu = 0.06 \text{ eV}$ compared to the damping due a fermionic dark radiation (red curve) particle producing $\Delta N_{\text{eff}} = 0.15$ at CMB and having the same mass, $m_{\text{DR}} = 0.06 \text{ eV}$. The shapes and the amplitudes of the damping are different for the two particles, implying that dark radiation and Standard Model neutrino masses can be measured simultaneously in surveys like LSST.

hidden sector models may be thermal, e.g. [193, 194], or nonthermal, e.g. [195]. In either case, the entropy of the hidden sector must generically be either deposited into the SM thermal plasma or carried by dark radiation. Dark radiation is thus a generic component of dark sector model-building. The ratio of the dark radiation temperature to the temperature of the SM will generally depend on the physics of reheating [196] as well as on the degrees of freedom in both sectors and the strength of the leading coupling(s) between them, and for our purposes can be treated as a free parameter. Models where a dark radiation species is directly involved in the freezeout of DM are frequently motivated by structure formation, and in such models the radiation does not always free-stream. However, in the general case, the lightest stable state(s) in the dark sector may be well-described by a thermal free-streaming state, as in the SM; e.g. [197].

An extra light degree of freedom should in principle show up in measurements of N_{eff} from the primary CMB, and the size of the signal is proportional to $T_{\text{DR,CMB}}^4$, where $T_{\text{DR,CMB}}$ is the temperature of the dark radiation species at the epoch of CMB last scattering. *Planck* has already constrained $\Delta N_{\text{eff}} < 0.33$, and future CMB Stage 4 experiments should be able to improve the constraints to $\Delta N_{\text{eff}} \lesssim 0.03$ [12]. However, it should also be noted that if the particles are relativistic

at CMB last scattering, ΔN_{eff} is insensitive to the mass of the dark radiation particle. Apart from a signature on N_{eff} , these light degrees of freedom can also damp the matter power spectrum at late times. The amount of damping is proportional to the energy density of the species today, Ω_{DR} . Assuming that these particles are non-relativistic at late times, the energy density is proportional to the mass of these dark radiation particles m_{DR} , and to the late time number density n_{DR} , i.e $\Omega_{\text{DR}} \propto m_{\text{DR}} n_{\text{DR}}$.

Since the damping of the matter power spectrum on small scales depends on Ω_{DR} , surveys like LSST will be sensitive to this parameter. We illustrate this with the following example, in which we suppose that $\Delta N_{\text{eff}} = 0.15$ is observed in future CMB experiments. We then repeat our Fisher matrix analysis, adding the mass of the dark radiation particle as an extra free parameter. The relevant power spectra and transfer functions were produced using the publicly available Boltzmann code CLASS [198, 199], which allows for easy implementation of bosons as an extra light species. We find that for fermionic dark radiation, the constraint forecast is $\sigma(m_{\text{DR}}) = 0.162 \text{ eV}$, and for bosonic dark radiation, $\sigma(m_{\text{DR}}) = 0.137 \text{ eV}$. Note that the mass bound is somewhat different for fermions than for bosons, simply due to the difference in density arising from Fermi-Dirac or Bose-Einstein statistics [119]. We further note that these results have been obtained after marginalizing over all other cosmological parameters including the neutrino mass. This means that LSST is potentially sensitive to multiple light species with different temperatures and masses, as the damping signatures will have different amplitudes and shapes, as illustrated in Fig. 3.9. The uncertainty on the inferred mass scales inversely with the magnitude of ΔN_{eff} observed in the CMB.

3.6 NNaturalness

NNaturalness [102] is a new approach to solving the Higgs hierarchy problem. NNaturalness posits that there are N copies of the Standard Model with differing Higgs masses. The Higgs mass squared parameters are distributed between $-\Lambda^2$ and Λ^2 , where Λ is the scale which cuts off quadratic divergences. If N is large enough, one of the copies will naturally have a Higgs mass parametrically smaller than the cutoff. The sector with the smallest negative Higgs mass squared is identified with the Standard Model.

Given current constraints on ΔN_{eff} , the question of why the Higgs mass is small has been

transformed into the question of why the sector with the lightest Higgs dominates the energy density of the universe. Crucially the lightest sector need to be dominantly reheated without making it otherwise special in any way. Not abiding by this rule would reintroduce the hierarchy problem.

To solve the problem, a new scalar field, the reheaton, was introduced. If the reheaton couples with equal strength to the various copies of the Standard Model through the interaction $\phi H_i^\dagger H_i$ and is light (with a mass comparable to the lightest Higgs), then it predominantly reheats the sector with the lightest negative Higgs mass. A light reheaton can only decay through off-shell Higgs bosons, which favors the lighter Higgs masses.

The cosmological signals of this scenario come from the fact that the reheaton inevitably populates some of the other copies of the Standard Model. Since the sector with a light Higgs boson can not be singled out in any way, the new sectors are very similar to us. For the sake of calculability we assume that they are identical copies, except for slightly heavier Higgs bosons (with and without a vev). Relaxing this assumption does not appreciably change the phenomenology. The new copies are reheated to temperatures slightly smaller than our own and like our sector have many light particles (including neutrinos) and a massless photon. The presence of these new light particles, with slightly lower temperatures than our own, makes this model an ideal candidate to be constrained and/or discovered by the techniques described before.

The cosmology of this realization of the NNaturalness paradigm is determined by two variables. The first is the mass of the reheaton, m_ϕ , and the second is the distribution of Higgs masses. The coupling of the reheaton to the Higgs cancels out once the temperature of our sector has been fixed. In principle, the differing Higgs masses can be drawn from any distribution. However any feature in the distribution would imply some assumptions on the dynamics related to the hierarchy problem at the scale Λ . So we take a uniform distribution with Higgs masses varying as

$$(m_H^2)_i = -\frac{\Lambda_H^2}{N}(2i + r), \quad -\frac{N}{2} \leq i \leq \frac{N}{2}. \quad (3.20)$$

r is a real and positive parameter that accounts for the possible probabilistic nature of the Higgs mass distribution or alternatively can be seen as a proxy for fine-tuning. If $r = 1$ the Higgs masses are equally spaced around zero. For $r < 1$ the lightest sector is closer to $m_h = 0$. Thus the




Figure 3.10: Ratio of the power spectrum at $z = 0$ for different NNaturalness scenarios (labeled by m_ϕ) compared to the Standard Model normal hierarchy scenario. The presence of extra free streaming species damps the power spectrum on the free streaming scale of the lightest of the extra species. On linear scales ($k > 0.1$ h/Mpc), the damping is below a percent level.

cosmology of this simple model of NNaturalness is determined by m_ϕ and r .

In order to test the visibility of this model, we consider a set of different reheaton masses. For each reheaton mass, we take the largest value of r such that ΔN_{eff} is small enough to satisfy current constraints. This leads to different combinations of masses and temperatures for the extra neutrino species. We implement these extra neutrinos using the CLASS code [198, 199]. The CLASS code allows for extra non-CDM species with different masses, temperatures and number densities, and therefore it is straightforward to check the ratio of the 3-dimensional linear matter power spectrum for different extra neutrino masses and temperatures. The results are plotted in Fig. 3.10. The fiducial cosmology was taken to be Λ CDM plus one massive neutrino species with $m_\nu = 0.06$ eV. For each reheaton mass, we see the generic feature of extra damping on small scales, where the amplitude of the damping is proportional to the energy density in the extra copies of the neutrinos. The scale at which these damping effects start showing up is given by the free streaming scale of the extra neutrinos. Since the extra copies have a higher mass and a lower temperature than the Standard Model neutrino, in each case we find that the free streaming scale of the extra neutrinos is smaller than that of the Standard Model neutrino.

We then proceed to use the formalism from § 3.2.3 to check the statistical significance of

the detection for these cosmologies with different reheaton masses. For the reheaton mass which produces the largest damping in the matter power spectrum, the calculated $\Delta\chi^2$ difference with the fiducial cosmology is 0.151.

This small $\Delta\chi^2$ suggests that these models will be hard to detect at LSST using information from linear scales only, along with the priors from *Planck*. To obtain more optimistic estimates, we modify some of our assumptions and survey parameters. First we set l_{max} for our various lens and source bins to 1000. This obviously violates our assumptions of linear theory, but if nonlinear scales can be modeled correctly, information from these small scales could, in principle, enhance the detectability of the NNaturalness models. We clarify that in this analysis, we continue to use linear power spectra. This is likely a somewhat conservative estimate of the signal at small scales since nonlinear terms raise the amplitude of the power spectrum on these scales. It is worth noting that pushing to $l_{\text{max}} \sim 1000$ at low redshifts also means that baryonic physics needs to be taken into account, which we can safely neglect at larger scales. This introduces a new systematic, and needs to be modeled accurately, either through simulations, or semi-analytics models to extract information about cosmological parameters from small scales, e.g. [200, 201]. The next assumption we make is to double the source and lens galaxy counts in the different bins. For the lens bins, the redMaGiC galaxies are only a fraction of all galaxies in the survey at that redshift, therefore doubling the number of galaxies for which the redshifts are known accurately seems an achievable target. Trying to double the number of sources seems a more optimistic assumption. Finally, we also add tighter priors on Ω_m and w from BAO measurements in DESI. We repeat our analysis to calculate $\Delta\chi^2$ for the NNaturalness models with respect to the fiducial cosmology, and we find that even under these assumptions, the largest $\Delta\chi^2$ for the NNaturalness realizations considered here is 0.53.

Since the combinations of m_ϕ and r that were considered here produced the highest ΔN_{eff} consistent with current bounds, measurements of N_{eff} from Stage 4 CMB experiments will easily detect the contribution from such models. But the detection will not be able to distinguish them from other scenarios e.g. [189, 180]. We have explored the detection of its signature in in large-scale structure. We find that even the lensing surveys of the next decade will be hard pressed to make a detection. The signal is below a percent-level suppression of the power spectrum, which

is challenging to achieve statistically and, even if we pushed survey parameter to achieve that, it would be daunting to control systematics at the desired level. One might hope that a completely different observable that goes to smaller scales or otherwise samples many more modes – like the Lyman- α forest or 21cm probes of structure - while achieving sub-percent control of systematics, will be sensitive to these cosmologies.

3.7 Summary and Discussion

We have obtained forecasts for the constraints on the neutrino mass from measurements of galaxy clustering and weak lensing in the LSST experiment, using information from only linear scales. The constraints on cosmological parameters, including the sum of the neutrino mass, improve with the number of redshift bins used in the analysis. We find that using 6 lens redshift bins and 6 source redshift bins, the sum of neutrino masses can be constrained to $\sigma(\sum m_\nu) = 0.041$ eV and the dark energy equation of state can be constrained to $\sigma(w) = 0.020$. While the constraint on the neutrino mass is very similar to that coming from Stage 4 CMB lensing experiments, LSST provides much stronger constraints on w than CMB lensing. These constraints were obtained by including priors on cosmological parameters coming from the *Planck* experiment. Further, the number of relativistic degrees of freedom N_{eff} was included as a parameter in our analysis. In principle, N_{eff} could be constrained independently from measurements of the CMB primaries in the CMB Stage 4 experiments. However, to provide a conservative estimate of the constraints, we initially include it as a parameter that is marginalized over when obtaining constraints on other cosmological parameters of interest. Our constraints are in agreement with similar forecasts in literature e.g. [133, 134].

When we include the expected tighter priors on other cosmological parameters from the upcoming survey DESI, the forecasted 1- σ constraint on the sum of neutrino masses from LSST goes down to 0.032 eV. A similar analysis for CMB Stage 4 lensing yields a constraint of 0.029 eV. These improvements come mainly from the tighter constraints on Ω_m and w , arising from accurate BAO measurements at low redshifts.

When we further included strong priors on N_{eff} , assuming that it will be measured independently in upcoming surveys, the forecasted constraint on the sum of neutrino masses at LSST tightens to

0.028 eV. For CMB Stage 4 lensing, this additional prior on N_{eff} yields a constraint of $\sigma(\sum m_\nu) = 0.023$ eV. These constraints on the sum of neutrino masses from CMB lensing are consistent to those presented in [135] when marginalizing over the same set of cosmological parameters.

We have also obtained constraints on $\sum m_\nu$ and w using a combination of galaxy clustering measurements from LSST and CMB stage 4 lensing. Once again, if low redshift information from DESI is included in the analysis, the constraints on the sum of neutrino masses becomes tighter. In fact, this combination can constrain $\sigma(\sum m_\nu) = 0.020$ eV, which is comparable to the bounds obtained from CMB lensing only in [12, 135, 147, 148], even though we marginalize over an extra parameter in N_{eff} for this comparison. Note that this combination of LSST, CMB Stage 4 lensing and DESI would be able to detect the minimal mass normal hierarchy of neutrinos at $\sim 3\sigma$. However, this is an relatively optimistic estimate. Firstly, this estimate assumed perfect overlap between the survey volumes of LSST and CMB Stage 4 experiment, and that the cross-correlations of data from the two surveys does not throw up unforeseen issues. This also assumes that all the major sources of systematics have been accounted for and behave as expected, and that all the surveys are able meet their target statistical error levels. On the other hand, it might be possible to extract even stronger constraints using this technique if smaller scales at LSST can be modeled more accurately.

One of the main sources of systematics in the LSST measurements will be shear bias and uncertainties in the photometric redshift of the source galaxies. We use redMaGiC galaxies, for which the photometric redshifts are measured to a very high accuracy, for lens galaxies, and therefore neglect lens redshift errors as a source of systematic errors. However, we account for redshift uncertainties of source galaxies and uncertainties in the shear bias using the parameterization in § 3.3. Our constraints were obtained without placing any priors on the nuisance parameters m_i . We have also investigated how the constraints improve when priors are placed on these parameters. The improvement is quite weak in $\sum m_\nu$, and even weaker for w . Placing a realistic prior of 0.02 on each of the m_i produces an improvement of $\sim 15\%$ on the neutrino mass constraint, and $\sim 5\%$ improvement on the error on w .

We note that we have made some simplifying assumptions in this study. We have assumed a flat cosmology in our analysis, and so Ω_k is fixed to 0 for all our calculations. We have also

assumed that the equation of state for dark energy is time invariant, setting $w_a = 0$. The effect of marginalizing over Ω_k and w_a as extra parameters on the neutrino mass constraint has been studied in e.g. [135] for Stage 4 CMB lensing. It should also be pointed out we do not consider the effects of extended theories of gravity on the power spectrum, which can be degenerate with the effect of massive neutrinos [202]. In this analysis, we have obtained the constraints on the sum of neutrino masses with the minimal mass normal hierarchy as the fiducial model. Various authors have looked into distinguishing the normal hierarchy of neutrino masses from the inverted hierarchy using cosmological observables, [145, 143, 203, 204, 205] for example. Moving on to systematics, we have neglected the effect of intrinsic alignments of galaxies and other effects which as discussed above are expected to be subdominant to the uncertainties considered in our analysis.

Beyond Standard Model neutrinos, thermal dark radiation can also lead to observable signatures on both N_{eff} at CMB, as well as on the late time matter power spectrum. Dark radiation candidates which were relativistic at CMB, but non-relativistic today damp the power spectrum on scales below their free streaming scale with an amplitude that is proportional to the mass of the particle. Here, we have studied the bounds that can be placed on the mass of dark radiation particles if they produce a detection of $\Delta N_{\text{eff}} = 0.15$ in the CMB primaries. The constraint on the mass depends on whether this new thermal species follows a Bose-Einstein distribution or a Fermi-Dirac distribution. For the former, we find a constraint on the mass $\sigma(m_{\text{DR}}) = 0.137 \text{ eV}$ at a $1\text{-}\sigma$ level. For the latter, the mass constraint is $\sigma(m_{\text{DR}}) = 0.162 \text{ eV}$ at the $1\text{-}\sigma$ level. We find that the mass constraints scale inversely as ΔN_{eff} . Since the neutrino mass was one of the parameters that was marginalized over to obtain this estimate, it suggests that LSST will be potentially sensitive to multiple light species with different temperatures and masses.

A specific model that posits the presence of extra light degrees of freedom is the NNaturalness model outlined in [102]. This model predicts the presence of multiple copies of neutrinos with different temperatures and number densities, all of which can contribute to N_{eff} . We have analyzed how effective LSST will be at detecting these extra neutrino species. We find that even when the ΔN_{eff} produced by these extra neutrinos is near the current bounds, LSST measurements will only be able to detect the effects of these at a statistical significance of 0.151σ . Even a more optimistic estimate yields a maximum $\Delta\chi^2 = 0.53$ between the fiducial cosmology and a cosmology including

extra neutrinos from the NNaturalness models. Since CMB Stage 4 will improve current constraints on N_{eff} by almost an order of magnitude, the contribution to N_{eff} from the NNaturalness model will be detectable. However, even for an unambiguous detection of N_{eff} from measurements of the CMB primaries, it will be difficult to distinguish NNaturalness from other exotic models which also predict extra light degrees of freedom during the CMB epoch e.g. [189, 180].

Through most of our analysis, we have only used modes for which linear perturbation theory is still valid. This allowed us to use a linear matter power spectrum instead of the full nonlinear matter power spectrum, as well to use the simplifying assumption that the covariance of various l modes are independent of each other. Moreover, at these linear scales, we could make the added simplifying assumption that galaxy bias is scale-independent. However, smaller scales which are mildly nonlinear, potentially have a lot of useful information due to the fact that there are many such independent patches on the sky. Extracting this information could lead to stronger constraints on the neutrino mass, and statistically significant detection of extra light degrees of freedom. To do this, we first need to accurately model the matter power spectrum in the presence of massive neutrinos and other light degrees of freedom. Secondly, we need to take into account off-diagonal covariances between different l modes in the nonlinear regime. And finally, to use the galaxy-galaxy autocorrelation and the galaxy-galaxy lensing spectra on small scales, we need an accurate model of galaxy bias on these scales, where it can be strongly scale dependent. Using this extra information will allow for even stronger constraints on the neutrino mass, as well a more statistically significant detections of other light degrees of freedom in the future.

Our study of dark radiation models makes clear that surveys planned for the next decade have interesting new discovery space. Indeed, for models such as NNaturalness, even more ambitious surveys will be required. One possibility is 3-dimensional surveys that can access more modes of mass fluctuations than any of the surveys discussed in this study. Surveys discussed in the Cosmic Visions whitepaper [206], typically for the 2030's, would achieve fractional errors on the power spectrum of well below a percent. Another option for an ambitious future survey would be an imaging survey, possibly from space, that has the sky coverage of LSST but with shear and photo- z information out to higher redshifts. We perform a naive mode-counting exercise to compare the number of linear modes that will be available in these futuristic surveys compared to LSST. To

do this, we assume that the redshift coverage in such a survey would extend out to $z \sim 3$, and that we can have 12 lens bins and 12 source bins extending from $z = 0.4$ to $z = 2$, and from $z = 1$ and $z = 2.8$ respectively. For each redshift bin, we determine l_{\max} following the procedure in §3.3. Then, we calculate the total number of modes in each bin as $N = f_{\text{sky}} l_{\max}^2$, and sum over all the tomographic bins to get the total number of modes in the survey [207]. Comparing a survey with these specifications to LSST, we find that these future surveys will increase the total number of modes in a survey by a factor of ~ 10 , and so the signal-to-noise ratio is expected to improve by more than a factor of 3. Therefore, such surveys would enable high significance detections of standard model neutrino properties and also test a variety of models for new particle species.

Chapter 4

Structure formation in Anisotropic Self-Interacting Dark Matter cosmologies

We study the impact of anisotropic self-interactions of Dark Matter on the structure and formation of dark matter halos through full cosmological simulations. We study the impact of these self-interactions on the position of the splashback radius of dark matter halos, as well as on the statistics of subhalos inside host halos.

4.1 Introduction

In the standard Λ -CDM paradigm, the Cold Dark Matter (CDM) component is generally considered to be collisionless and non-interacting. However, in recent years, there has been considerable interest in scenarios where self-interactions in the dark sector are important [21, 104, 105, 106, 107, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117, 118], for example. Self-interacting Dark Matter has been invoked as a solution to some of the small scale problems of the Λ -CDM cosmology. The simplest case of elastic self-interactions with an isotropic differential cross-section has been implemented in full cosmological simulations [23, 24, 27, 25, 26]. The slightly more complex case of elastic self-interactions with an anisotropic differential cross-section, where forward scattering is favored over large-angle scattering, has been studied in a number of papers [34, 208], but full cosmological simulations with this form of self-interactions is yet to be carried out.

The reason these anisotropic self-interaction models is interesting to study is because they can possibly explain two recent observational results, which are potentially in conflict with the predictions of the collisionless Cold Dark Matter scenario. The first observation relates to the structure of the outer profiles of dark matter halos. It has been pointed out in [30, 31] that there should be a sharp feature in the stacked density profile of halos which corresponds to the boundary of the multi-streaming region of halo. This feature occurs at the “splashback” radius, and the

position of the feature depends on various properties of the halo such as its accretion rate and the redshift at which the halo forms. The splashback feature has recently been detected in observations [28, 29]. However, both studies find that the location of the observed splashback feature is offset from that expected from simulations by $\sim 20\%$, and that the observed splashback radius is always smaller than in the simulations. In (cite More), the authors consider self-interactions of Dark Matter as one of the possible resolutions for this discrepancy.

The second observation has to do with the offsets of the light profiles and mass profiles of galaxies in large clusters. Assuming that the galaxies reside at the center of subhalos inside the host halo, the center of the light profile and the the center of the mass profile are expected to remain locked to each other, as long as both the stars and the dark matter are collisionless. However, if dark mater has extra self-interactions, the stars and the dark matter will start developing offsets, proportional to the size of the extra interactions. Recent studies have detected such offsets in multiple systems [32, 33].

It was shown in [34], that for anisotropic self-interactions of dark matter, where the two particle interactions are elastic and energy conserving, the main effect can be described by an effective drag in the direction of the relative velocity on the subhalos as they travel through the host halo. This drag force, then, can explain both the shortening of the splashback radius, as well as the offset between stars and dark matter. This effective drag is in sharp contrast to what is seen in models of isotropic scattering for dark matter, where the main effect is subhalo evaporation. This is because for isotropic scattering, large angle scatterings, corresponding to large momentum transfers are more likely. These large momentum transfers kick the subhalo particles out of the subhalo, leading to mass loss. Therefore, anisotropic SIDM can exhibit very different signatures on the structures of dark matter halos compared to isotropic self-interactions, and it is therefore interesting to study these effects through full cosmological simulations.

This chapter is organized as follows. In §4.2, we set down the formalism to describe anisotropic self-interactions. In §4.3 we set up a toy problem with a known analytic solution, and describe the two methods to simulate this toy model. In §4.4, we describe how we apply implement the self-interactions in the cosmological code L-Gadget, and provide the results of full cosmological simulations run with both our methods. We draw the main conclusions of our work in §4.5, and




Figure 4.1: Elastic collision of two dark matter particles in the center of mass frame. For anisotropic SIDM, the scattering angle θ is drawn from a distribution where small values of θ are favored over large values of θ .

discuss future avenues extensions and refinements.

4.2 Formalism for anisotropic SIDM

In this section, we lay down the basic formalism for describing anisotropic SIDM following the the conventions in [34]. Similar to that paper, we will consider a differential cross section such that

$$\frac{d\sigma}{d\Omega} = \mathcal{K} \frac{1 + \cos^2 \theta}{1 - \cos^2 \theta}, \quad (4.1)$$

where \mathcal{K} represents the normalization constant. We point out that for this choice, the total cross section,

$$\sigma = \int \frac{d\sigma}{d\Omega} d\Omega, \quad (4.2)$$

actually diverges, which means that for the short range interactions considered here, two particles at the same point will always interact with each other. However, as we will show later, physical quantities like the momentum transfer in the direction parallel to the relative velocity of the colliding particles, as well as the energy transfer in the perpendicular direction are finite and well-behaved, even for this $d\sigma/d\Omega$.

We begin by considering the elastic collision of two dark matter particles in the center of mass frame. The particles have initial velocities \vec{v} and $-\vec{v}$. The final velocities are \vec{v}' and $-\vec{v}'$, and the scattering angle is given by θ , where θ is drawn from a distribution described by Eq. 4.1. In terms

of θ , we can write down the changes in various quantities before and after the collision:

$$\delta(v^2) = 0 \quad (4.3)$$

$$\delta(v_{\parallel}) = -v(1 - \cos \theta) \quad (4.4)$$

$$\delta(v_{\parallel}^2) = -v^2 \sin^2 \theta \quad (4.5)$$

$$\delta(v_{\perp}^2) = v^2 \sin^2 \theta \quad (4.6)$$

$$\delta(\vec{v}_{\perp}) = v \sin \theta \hat{\theta} \quad (4.7)$$

where v_{\parallel} is the measured in direction parallel to the relative velocity direction, and v_{\perp} is measured in the direction perpendicular to the relative velocity. Since θ is small for most collisions, the change in both v_{\parallel} and v_{\perp} will be small compared to the initial velocity of the particles.

We now consider what happens when two cold, localized streams of dark matter pass through each other. We point out that since the interactions considered here are elastic in nature, a single cold, isolated stream will not be affected by these interactions. We focus on the changes in the parallel and perpendicular velocities of one of the streams, and assume that the other stream has a local density given by ρ . For a particle in the first stream, the number of collisions it suffers in a time interval dt is given by

$$dC = \frac{\rho v_{\text{rel}} dt}{m} \frac{d\sigma}{d\Omega} d\Omega, \quad (4.8)$$

where m is the mass of the dark matter particle, and v_{rel} is the relative velocity of the two particles. For two identical particles, $|v_{\text{rel}}| = 2|v|$.

If we want to compute the average change in the parallel velocity of the particles in the first stream, using Eq. 4.4 and Eq. 4.8 we can write it as

$$d\langle v_{\parallel} \rangle = -\frac{\rho v_{\text{rel}} v}{m} dt \int \frac{d\sigma}{d\Omega} (1 - \cos \theta) d\Omega. \quad (4.9)$$

Since the integral has no dependence on ϕ , it can be rewritten as

$$d\langle v_{\parallel} \rangle = -\frac{\rho v_{\text{rel}} v}{m} dt (4\pi\mathcal{K}) \int_0^{\pi/2} \frac{1 + \cos^2 \theta}{1 - \cos^2 \theta} (1 - \cos \theta) \sin \theta d\theta, \quad (4.10)$$

where the trivial integral over ϕ has been carried out, and the θ integral has been split into two parts from 0 to $\pi/2$, and $\pi/2$ to π . Assuming that the particles are indistinguishable, the two integrals are symmetric, and introduces a factor of 2 in the above equation.

We now define the momentum transfer cross-section in the following manner:

$$\sigma_T = (4\pi\mathcal{K}) \int_0^{\pi/2} \frac{1 + \cos^2 \theta}{1 - \cos^2 \theta} (1 - \cos \theta) \sin \theta \, d\theta. \quad (4.11)$$

While the total cross section diverged for the differential cross section in Eq. 4.1, the momentum transfer cross section can be easily evaluated, and is finite. This is because even as the total number of interactions get large, most interactions produce a small change in the velocity because of the peaked distribution in θ , and therefore the integrated effect is finite.

With the definition of the momentum transfer cross section, we can rewrite Eq. 4.10 as

$$d \langle v_{\parallel} \rangle = -\frac{\rho v_{\text{rel}} v \sigma_T}{m} dt \quad (4.12)$$

Since the two streams of dark matter that we considered are symmetric, one can derive a similar expression for the change in the parallel velocity of the second stream. Anisotropic self-interactions, therefore, produce an effective drag on each stream in the direction of the relative velocity of the two streams, reducing the velocity along that direction. Further, since this drag is expected to be primarily responsible for the observable effects that we are looking for, we will calibrate the strength of the anisotropic self-interactions using this momentum transfer cross section σ_T , rather than the total cross section σ , as is usually the practice when considering isotropic self-interactions.

Next we consider the amount of energy transferred in the perpendicular direction. For a single collision, this was given by

$$\delta(v_{\perp}^2) = v^2 \sin^2 \theta. \quad (4.13)$$

If we once again integrate over all collisions, we get the following expression for each stream:

$$d \langle v_{\perp}^2 \rangle = \frac{\rho v_{\text{rel}} v^2}{m} dt \int \frac{d\sigma}{d\Omega} \sin^2 \theta d\Omega. \quad (4.14)$$

As earlier, we can split up the $d\Omega$ integral into the θ and ϕ components, and define the viscosity

cross section as

$$\sigma_V = (4\pi\mathcal{K}) \int_0^{\pi/2} \frac{1 + \cos^2\theta}{1 - \cos^2\theta} \sin^2\theta \sin\theta d\theta. \quad (4.15)$$

Once again, this is a finite quantity in spite of the total cross section being divergent. Therefore physical quantities such as the momentum transfer and energy transfer as result of such interactions are well defined quantities. the definition of σ_V , Eq. 4.13 becomes

$$d\langle v_{\perp}^2 \rangle = \frac{\rho v_{\text{rel}} v^2 \sigma_V}{m} dt \quad (4.16)$$

Note that while $d\langle v_{\perp}^2 \rangle$ is non-zero, $d\langle v_{\parallel}^2 \rangle = 0$, since the direction in which the particles will scatter will be random. Also note that since $\delta(v^2) = 0$ in the center of mass frame, there is also an equal effect in the energy in the direction parallel to the relative velocity. The overall effect of the self-interactions, therefore, is to reduce the average velocity in the parallel direction, and introduce a velocity dispersion in both the parallel and perpendicular directions.

4.3 Toy Problem

In this section, we describe a toy problem to calibrate our simulations methods. We then describe in detail our two methods for simulating the self interactions, and compare the results of the two approaches when applied to our toy problem.

4.3.1 Setup

In our toy problem, we consider two localized sets of particles - one set made up of N_{halo} particles which are initially at rest, and another set of N_{subhalo} particles which are moving toward the first set of particles with velocity v . This is to imitate a situation where an actual subhalo in the full cosmological simulations falls into a host halo. For concreteness, we consider cases where $N_{\text{halo}} = 10N_{\text{subhalo}}$. All particles are taken to have the same mass, and so, the density of the “halo” is 10 times the density of the “subhalo”. We assume that the two sets of particles interact with each other over some time interval Δt . We ignore gravity, so that we can use the expressions derived in §4.2 to calculate the analytical expectations for this toy model. Note, characterize the strength of the self-interactions using the momentum transfer cross section σ_T , as opposed to the value of the

total cross section, which is divergent in this case. We implement the self-interactions using two methods and then compare the results from these simulations with the analytic expectations for the average parallel and perpendicular velocities of the subhalo particles after all interactions have taken place.

4.3.2 Deterministic Drag Method

In the first method, for each particle, we find all neighbors within some search radius at any given timestep. For each neighbor, we find the direction of the relative velocity between the original particle and the neighbor. Then, in the center of mass frame of the two particles, we use Eq. 4.12 to compute the change in the velocity of the first particle in the direction of the relative velocity, so that

$$\Delta v_{\parallel} = -\frac{\rho v_{\text{rel}} v \sigma_T}{m} \Delta t. \quad (4.17)$$

ρ represents the density of the neighbor at the position of the first particle. If all simulation particles have the same mass, ρ can be calculated using the mass and the search volume over which we search for neighbors. We compute the energy lost by the particle in the parallel direction using Eq. 4.17, and kick the particle in the plane perpendicular to the parallel velocity, where the magnitude of the kick is adjusted so that there is no overall change in the energy of the particle in the center of mass frame. The exact direction of the kick in the perpendicular plane is determined by drawing a random angle ϕ between 0 and 2π .

In order to ensure overall momentum conservation, we need to ensure that the neighboring particle is also kicked in exactly the opposite direction in the perpendicular plane. To ensure this, we assign a unique particle ID to each particle. If the ID of the first particle is larger than the ID of the neighboring particle, we change the velocity of both particles in the perpendicular plane, using the randomly generated ϕ for the first particle, and $\phi' = \phi + \pi$ for the neighboring particle. If the ID of the first particle is smaller than the ID of the neighbor, then we only change the velocity of the first particle in the parallel direction. We loop over all neighbors using this prescription, and then loop over all particles to make sure we account for all self-interactions.

When the two sets of particles have no overlap with each other, our method ensures that the self-interactions produce no effect on the individual bunches. This is because Eq. 4.17 depends

explicitly on v_{rel} , which is 0 for all particles belonging to the same bunch. Only when the two bunches overlap, and our neighbor search algorithm finds particles from the other bunch, $v_{\text{rel}} \neq 0$, and the self-interactions start affecting the evolution of the particles.

4.3.3 Random Angle Method

In this method, we treat the “halo” and “subhalo” particles as we would treat microscopic interactions - that is, we assume that for each two-body interaction, we can draw a scattering angle θ from the underlying distribution of the differential cross section, and calculate the post-collision velocities using this angle θ . Most studies of isotropic self-interactions implement this method in their simulations. However, since the total interaction cross section diverges for the specific type of anisotropic self-interactions that we consider here, care needs to be taken to ensure that the implementation is self-consistent, while not being exorbitantly expensive in terms of the computational time.

From Eq. 4.1 we see that there is a divergence at $\cos \theta = 1$. In our implementation, therefore, we choose a cutoff ϵ so that for the differential cross section, we use

$$\frac{d\sigma}{d\Omega} = \begin{cases} \mathcal{K} \frac{1+\cos^2 \theta}{1-\cos^2 \theta}, & \text{if } \cos \theta < 1 - \epsilon \\ \mathcal{K} \frac{1+(1-\epsilon)^2}{1-(1-\epsilon)^2}, & \text{otherwise} \end{cases}$$

This gets rid of the divergence, and we can then easily compute the cumulative distribution function (CDF) that is required to generate the angle θ in the simulations. We choose $\epsilon = 0.005$ throughout. We illustrate the analytic PDF and the cut-off PDF in Fig. 4.2 for this choice of ϵ . We also plot the CDF using the cut-off distribution. For a given ϵ , we can compute the physically relevant quantities such as the momentum transfer cross section, and the viscosity cross section. We find that both σ_T and σ_V computed with the cut-off distribution are within 0.5% of the analytic value. Therefore, by truncating the distribution function, we are almost exclusively ignoring those interactions which contribute very little to the parallel and perpendicular changes in the velocities of the interacting particles. Further, we are not biasing the distribution of the angles drawn from $d\sigma/d\Omega$ in a way which affects physical quantities in any significant manner.

To implement this method in the simulations, we proceed in the following way - once again, for

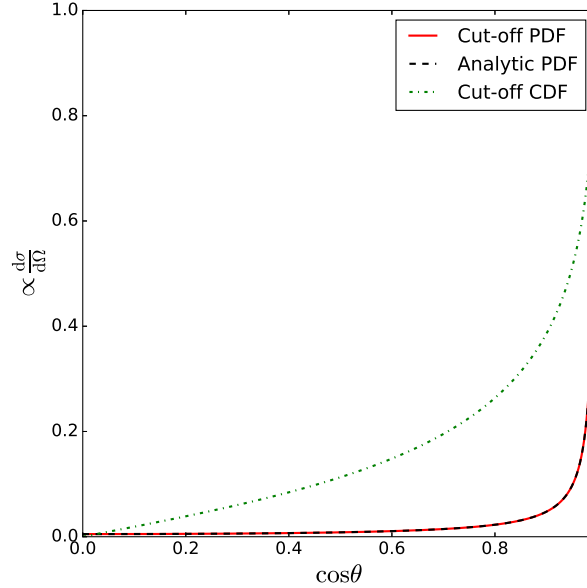


Figure 4.2: Comparison of the analytic PDF (black dashed line) and the cut-off PDF (red solid line) for the choice of $\epsilon = 0.005$. The normalization is adjusted to fit within the plot range. We also plot the CDF (blue dot-dashed line) for cut-off PDF. We find that the momentum transfer cross section σ_T computed with the cut-off PDF is within 0.5% of the analytic value for this choice of ϵ .

a given particle, we search for all neighbors within a given search volume. For each neighbor, we calculate the probability for the interaction of the two particles, given by

$$P_{\text{int}} = \frac{\rho v_{\text{rel}} \Delta t}{2m} \sigma, \quad (4.18)$$

where σ is the total interaction cross section computed using the cut-off distribution. The extra factor of 2 in the denominator of the right hand side accounts for the fact that if particle j is the neighbor of particle i , then we also find particle i when looping over the neighbors of particle j . The factor of 2 ensures that we do not double count the interactions for each pair. Once again ρ can be computed given the search volume for neighbors. We then generate a random number between 0 and 1, and if the value of the random number is smaller than P_{int} for the given pair of particles, we assume that the two particles interact with each other.

Once we have determined that two particles interact with each other, we draw a uniform random number between 0 and 1. We then use the CDF of the distribution function to convert the random number into a random angle θ . We change the velocity in the direction parallel to the relative

velocity using

$$\Delta(v_{\parallel}) = -v(1 - \cos \theta). \quad (4.19)$$

In the plane perpendicular to the relative velocity, we need to pick another uniform random number between 0 and 2π to define the angle ϕ , and therefore a specific direction in the plane. We change the magnitude of the perpendicular velocity using

$$\Delta(v_{\perp}) = v \sin \theta. \quad (4.20)$$

We perform these operations for all neighbors of a given particle, and then loop over all particles in the simulation. Once again, for two isolated cold blobs of dark matter, this method ensures that self-interactions do not affect the dynamics. This is because the probability of interaction itself depends on v_{rel} , and for all particles moving together, $v_{\text{rel}} = 0$, implying $P_{\text{int}} = 0$ for each pair. In the presence of particles from the other blob, however, the self-interactions start taking effect.

Notice that in this random angle method, the momentum transfer cross section σ_T or the viscosity cross section σ_V do not appear explicitly in any of the steps that are required to compute the change in the velocities of the particles. What appears instead is the total cross section evaluated using the cut-off PDF distribution. However, as discussed in §4.3.1, we will use σ_T to calibrate the strength of the cross section, and so we will match the results from this method with that from the other method by making sure that σ_T computed using the cutoff distribution matches for the two methods.

4.3.4 Results

We implement both methods described above to the toy problem laid out in §4.3.1. We look at three different cases, that is, $N_{\text{subhalo}} = 50$, $N_{\text{subhalo}} = 100$, and $N_{\text{subhalo}} = 200$. We start with the same initial conditions for all simulations, but we change the particle mass in three cases above so that the total mass of the “halo” and the “subhalo” is the same across the three cases. This is equivalent to changing the mass resolution of a full cosmological simulation. We also run each of the three cases 10 times so that we can compare the distributions of the final results. In particular we




Figure 4.3: Results of our toy model for anisotropic SIDM interactions. We plot the average parallel velocity of all subhalo particles in the left panel, and the average of the square of the perpendicular velocity in the right panel. We look at three different cases, $N_{\text{subhalo}} = 50$, $N_{\text{subhalo}} = 100$, and $N_{\text{subhalo}} = 200$, and 10 realizations of each case. The 10 realizations are slightly displaced along the x-axis to help with the visualization. The analytic expectations are plotted using the blue lines. The results from the deterministic drag method are represented using black points, and the random angle method using the green dots.

will be interested in the average parallel velocity, and the average of the square of the perpendicular velocity of all subhalo particles after they have finished interacting with the halo particles.

We present our results in Fig. 4.3. We compare the results from the two methods, represented by the black dots for the deterministic drag method, and by green dots for the random angle method to the analytic expectations plotted using the blue lines. We find that for the deterministic drag method, we always get the parallel velocity to match the analytic expectation, irrespective of the number of particles present in our simulations. This is expected, since all subhalo particles experience the same drag due to the presence of the halo particles, and therefore there is no dispersion about the mean result. On the other hand, for the perpendicular direction, the black dots do not coincide with each other. This is because there are random kicks in the perpendicular direction. As expected, when we increase the number of particles in our simulation, and therefore the number of interactions faced by each particle, the results converge, as can be seen by comparing the results on the right panel of Fig. 4.3 for higher values of N_{subhalo} . It is, however, important to note that the results do not converge to the analytic value. This happens because in our

implementation, we compute the energy lost due to the drag in the parallel direction in the center of mass frame, and convert it completely into a kick in the perpendicular direction only, thereby injecting more energy than expected in the perpendicular direction. As it turns out, there is no way in which we can use the expression for the deterministic drag in Eq. 4.17, and ensure the correct transfer of energy in the perpendicular direction, while at the same time making sure of the global energy conservation in a particle based simulation. Even though this method suffers from this drawback, it has the advantage that it gets the parallel drag correct, and since this is the effect that we will be interested in for the cosmological simulations, we implement it in next section for full cosmological runs.

Next, we look at the results from the random angle method. We can see from the left panel of Fig. 4.3 that even for the parallel velocity of the subhalo particles, and for a fixed N_{subhalo} , there is a scatter about the analytic result. This is expected, since the number of interactions, as well as the value of θ which determines the change in velocity in a given interaction is different in each different realization. Even though the scatter is large compared to the deterministic drag method, the mean of all 10 realizations is close to the analytic expectation. We find similar results in terms of the scatter about the mean for the perpendicular direction, as shown in the right panel. However, unlike the deterministic drag method, in this case, the mean of the 10 realizations is very close to the analytic expectation. This is basically a check that we are performing the Monte-Carlo integration over the distribution function correctly - that is - we are effectively calculating the expected value for the momentum transfer cross section σ_T and the viscosity cross section σ_V .

However, it is interesting to note that for this method, increasing the number of particles, while holding the total mass the same, has no effect on the convergence of the results to the analytic value. This is true for both the parallel and perpendicular directions. This result can be understood by examining Eq. 4.18. If the mass of an individual particle goes down by a factor of 2, for example, then the density deposited by that particle at the location of another particle also goes down by a factor of 2. This means a lower probability of interaction. However, since the number of particles has to be increased by the same factor of 2 to ensure that the total mass remains constant, the overall results do not change. Similarly, taking smaller time steps will not change the overall results too much, because even though the probability of interaction becomes smaller when individual time

steps decrease, adding up the probabilities from individual time steps leads to the same result.

This has the following implication when implemented in full cosmological simulations - we do not expect the random angle method to converge faster to analytic expectation as the mass resolution of the cosmological simulation is increased, or if the individual time steps are reduced. This is contrast to the gravitational interactions in these simulations, which converge with both the mass resolution and the time resolution. However, when averaged over many realizations, this method does yield the correct result in both the parallel and perpendicular directions.

As we have illustrated in this section, both the deterministic drag method, as well as the the random angle method, have their own advantages and drawbacks, which will be present when applied to cosmological simulations. Therefore, we it is interesting to implement both methods, and compare their results in terms of the observables that we are interested in - how the splashback radius changes in the presence of these self-interactions, and how stars and dark matter subhalos behave in terms of offsets between the two as they fall into a large host halo.

4.4 Cosmological simulations with anisotropic SIDM

In this section, we describe the implementation of anisotropic self-interactions in full cosmological simulations by modifying the L-Gadget code. L-Gadget is a modified version of the GADGET-2 code, optimized for cosmological N-body simulations. We then present the results from these simulations, and discuss the implications of these results.

4.4.1 Implementation in L-Gadget

L-Gadget is a modified version of the GADGET-2 cosmological code [79]. This is an N-body PPM code, where long range gravitational forces are calculated by solving the Poisson equation on a grid, while the short range gravitational force is calculated using a tree-walk method. Since the tree-walk that L-Gadget employs for the gravitational interaction is already set up to find the nearest neighbors of a given particle, we incorporate the steps for implementing the methods discussed in the previous section into the same tree walk. With this, there no significant change in the efficiency of the code even when self-interactions are included.

To implement the self-interactions, we define a fixed search radius, and assume that only par-

ticles within this search radius can possibly interact with each other. We set this search radius to be 1.5 times the gravitational softening length h in L-Gadget. For most simulations, h is between $1/60$ to $1/30$ of the mean inter-particle separation. We check that for both implementations, our final results do not depend strongly on the exact value of our search radius.

We could also have defined a density-dependent search radius, or use the N nearest neighbors, in keeping with what is usually done in SPH simulations, but we do not do so for the following reason. Consider what happens to a single subhalo lying near the subhalo center, as the entire subhalo travels through the host halo. If we use a density dependent search radius, then near the center, the search radius would be small, and most particles within this search radius would belong to subhalo itself, rather than from the host halo. The same thing obviously happens if we choose the N nearest neighbors of the subhalo particle. For these particles, the relative velocities will be small compared to the true relative velocity of the subhalo relative to the host halo, and since the effect of the self interactions scale with the relative velocity, we will underestimate the overall effect. On the other hand, using a fixed search radius, and ensuring that this size is large compared to the size of subhalo centers, we get around this problem.

We also assume a particular profile for each particle $\propto (1 - r/r_{\text{search}})$ to calculate the density of that particle as seen at the position of a second particle. Here r is the radial separation between the two particles, r_{search} is the search radius, and $r < r_{\text{search}}$ for the two particles to interact. Once again, we check that our results are not sensitive to the profile we choose - the results are similar even if we choose a top-hat profile for the particles .

4.4.2 Results

We run cosmological simulations for a box with side $150h^{-1}\text{Mpc}$ using 512^3 particles. We compare the case of $\sigma_T = 1\text{ cm}^2/\text{g}$ (for both the deterministic drag method, as well as the random angle method) to $\sigma_T = 0$, or the non-interacting case. For each simulation, we run the Rockstar Halo Finder [89] to identify halos at $z = 0$.

First, we look at the density profile and position of the splashback radius defined by Dark Matter particles for clusters in the mass range of $1 - 4 \times 10^{14}h^{-1}M_{\odot}$. For this, we use particles from all the clusters identified in the simulation box within this mass range to compute the stacked




Figure 4.4: Stacked density profiles for halos of mass $1 - 4 \times 10^{14} h^{-1} M_{\odot}$ for different implementations of self-interactions. The random angle method changes the density profile more than the deterministic drag method.

density profiles, and then compute the slope of the density profile. The splashback radius is defined as the point where the slope of the density profile is at a minimum.

We plot the stacked density profiles in Fig. 4.4. Comparing to $\sigma_T = 0$, we find that the random angle method produces a larger effect on the density profile than the deterministic drag method, for the same interaction cross section. However, the effect is not very large, except in the central region of the halo. We also show the slopes of the stacked density profiles in Fig. 4.5. Once again, it is clear that the random angle method shows a larger effect of the slope. However, it is interesting to note that the position of the splashback, defined by the point where the slope goes to its minimum value, hardly changes in the presence of the self interactions. Since the interactions are elastic in nature, a single dark matter particle being dragged during its orbit through the halo, implies that another particle has gained energy in its orbit, through energy conservation. It is therefore, possible that the effects balance each other to keep the splashback radius defined by dark matter particles to be the same.

However, in observations, one measures the density profile and splashback radius using visible galaxies, which live inside subhalos of the host halo. Therefore, to obtain a fairer representation of what might be seen in actual observations, we identify the subhalos of the host halos in the mass

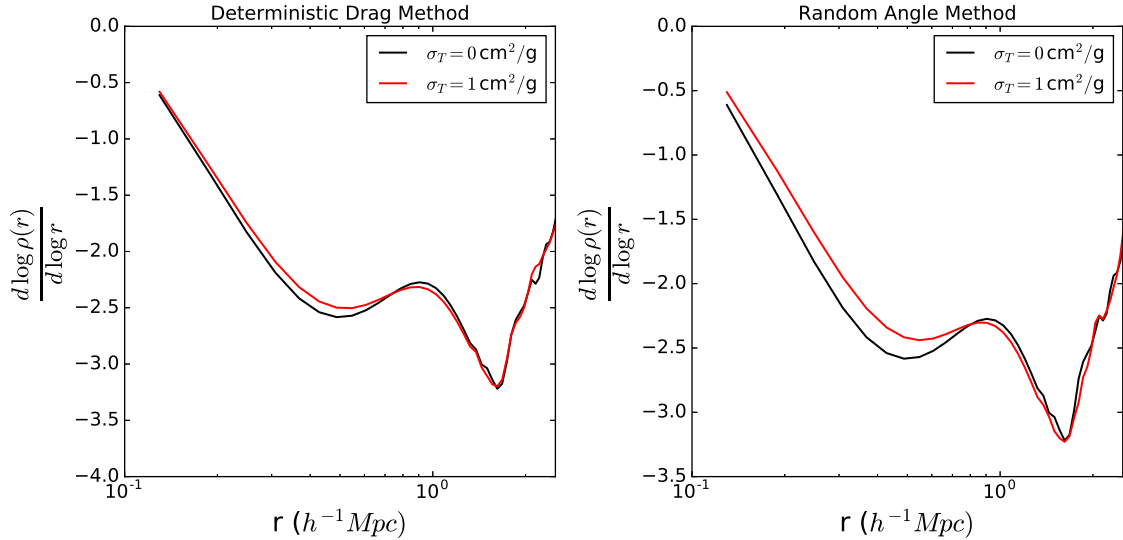


Figure 4.5: Slope calculated from the stacked density profiles of halos with mass $1 - 4 \times 10^{14} h^{-1} M_{\odot}$ for different implementations of self-interactions. It is clear that the random angle method changes the slope of the density profile more than the deterministic drag method, but there is no effect on the splashback radius of dark matter particles.

bin we use. These are also identified from the Rockstar Halo Catalog. Given the force resolution of our simulations, we only use subhalos with mass greater than $10^{12} h^{-1} M_{\odot}$. Once we have identified the subhalos, we can compute the stacked number density profiles of these subhalos as a function of distance to the host halo center.

These profiles are plotted in Fig. 4.6. Compared to $\sigma_T = 0$, the subhalo number profile is very different in the deterministic drag case. The number densities are lower throughout the whole profile, especially near the center. For the random angle method, there are differences near the center, but very small differences on the outskirts. Since the overall profile does not change much as shown in Fig. 4.4, this means that there is much greater subhalo evaporation in the deterministic drag method. We had seen in the toy model experiments in §4.3.4 that the deterministic drag method led to too much energy being transferred in the direction perpendicular to the direction of relative velocity. This could explain why we find a higher rate of subhalo evaporation for this method, compared to the random angle method, which gave the correct result on average for the energy transfer in the perpendicular direction.

It would have been very interesting to calculate the splashback radius using the subhalo number profiles. This is a much noisier measurement since it involves taking derivatives of noisy data, and




Figure 4.6: Stacked subhalo density profiles for halos of mass $1 - 4 \times 10^{14} h^{-1} M_{\odot}$. Only subhalos with mass greater than $10^{12} h^{-1} M_{\odot}$ were used to obtain this profile. In this case, we find that the deterministic drag method shows a larger effect for the same cross section. This implies a higher subhalo evaporation rate for this method compared to the random angle method.

then determining where the derivative has a minimum. Given the overall small number of clusters in our simulation volume, and the resolution of our simulations, the measurement is too noisy to make a meaningful detection of the shift of the splashback radius. We plan to check for the shift using simulations with larger physical volumes, as well as larger numbers of particles - which will allow us to resolve even smaller subhalos.

4.5 Summary and Discussion

We have studied the effects of elastic but anisotropic self-interactions of Dark Matter on the profiles of dark matter halos, in particular, looking at the stacked density profile and splashback radius for Dark Matter particles, and stacked number profiles for subhalos. These interactions are frequent, but with low momentum transfer per collision, which led to the prediction of an effective drag on subhalos as they fall into the host halo [34].

Here, we have developed the formalism to describe these self-interactions. We find that for this form of self-interactions, the physical quantities are the momentum transfer cross section σ_T , in the direction parallel to the relative velocity of two particles, and the viscosity cross section σ_V ,

which is a measure of the energy transferred in the direction perpendicular to the relative velocity. Even though we have considered the case where the total cross section σ diverges, these quantities remain finite.

We have developed two different methods to implement this form of self interactions, and test it on a toy model. We find that one of the methods, deterministic drag method in §4.3.2, always gives the correct answer for the drag in the parallel direction, and also converges quickly with the number of interactions. However, in the perpendicular direction, it transfers more energy than expected in the perpendicular direction, effectively yielding an incorrect value for σ_V for the toy model. The other method, the random angle method outlined in §4.3.3, we find that the scatter about the correct result is much larger. However, on average, this method yields the correct value for both the parallel drag and the transfer of energy in the perpendicular direction.

We implemented both these methods in the cosmological code L-Gadget2, and ran simulations with $\sigma_T = 1 \text{ cm}^2/\text{g}$. Using Dark Matter particles, we looked at the stacked density profile and its slope for halos in the mass range $1 - 4 \times 10^{14} h^{-1} M_\odot$. We found that both methods show some deviation from the non-interacting case, especially toward the center of the host halo. The effect is larger when the interactions are implemented using the random angle method. However, we find that for both methods, we seem to find no movement in the splashback radius of Dark Matter particles. This would seem odd, given the expected drag force between particles. However, for single dark matter particles, every time one particle is dragged, or loses energy in the rest frame of the halo, another particle is kicked, or gains energy in the same frame. It is, therefore possible that the overall splashback radius for particles does not move under the influence of this form of self-interactions.

Since observations use galaxies to define the splashback radius of clusters, we also looked at the number density profiles for the halos by stacking all their subhalos above a certain mass threshold. In this case, we found that the deterministic drag method produces a large change in the subhalo number profiles, implying a much larger subhalo evaporation rate than in random angle case. We saw that the former method transfers too much energy in the perpendicular direction in the toy model, and this is likely the reason behind the enhanced evaporation rate. It would have been interesting to see if the splashback radius of the subhalos shifts in the presence of the self-

interactions, but our current statistics of subhalos does not allow for a significant detection of the splashback radius.

As an extension to this work, we plan to run simulations on larger cosmological volumes, and larger numbers of particles, get better statistics for the subhalo number counts, and therefore be able to actually measure the splashback radius. We also plan to run simulations for a range of values of σ_T , and see if there are any qualitative differences in the results for stronger self-interactions. Finally, we plan to implement this form of self interactions in simulations which include baryonic physics, to calculate the effect of self-interactions on the displacements of the light profiles of stars and the mass profile from the dark matter subhalo. Given our current results, it seems that the random angle method yields more reasonable results than the deterministic drag method, and this is the method that should be implemented in future simulations. However, further tests are needed to make sure that this is still true for higher resolution simulations, and is not an artifact of our choice of simulation and self-interaction parameters.

References

- [1] F. Zwicky, *Die Rotverschiebung von extragalaktischen Nebeln*, *Helvetica Physica Acta* **6** (1933) 110–127.
- [2] SUPERNOVA COSMOLOGY PROJECT collaboration, S. Perlmutter et al., *Discovery of a supernova explosion at half the age of the Universe and its cosmological implications*, *Nature* **391** (1998) 51–54, [[astro-ph/9712212](#)].
- [3] SUPER-KAMIOKANDE COLLABORATION collaboration, Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino et al., *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (Aug, 1998) 1562–1567.
- [4] SNO COLLABORATION collaboration, Q. R. Ahmad, R. C. Allen, T. C. Andersen, J. D. Anglin, G. Bühler, J. C. Barton et al., *Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the sudbury neutrino observatory*, *Phys. Rev. Lett.* **87** (Jul, 2001) 071301.
- [5] K2K COLLABORATION collaboration, M. H. Ahn, E. Aliu, S. Andringa, S. Aoki, Y. Aoyama, J. Argyriades et al., *Measurement of neutrino oscillation by the k2k experiment*, *Phys. Rev. D* **74** (Oct, 2006) 072003.
- [6] S. F. King, *Neutrino mass models*, *Rept. Prog. Phys.* **67** (2004) 107–158, [[hep-ph/0310204](#)].
- [7] S. F. King, *Models of Neutrino Mass, Mixing and CP Violation*, *J. Phys.* **G42** (2015) 123001, [[1510.02091](#)].
- [8] <https://www.katrin.kit.edu/>.
- [9] <http://snoplus.phy.queensu.ca/Home.html/>.
- [10] S. Dodelson, *Modern cosmology*. Academic Press, 2003.
- [11] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel et al., *LSST Science Book, Version 2.0, ArXiv e-prints* (Dec., 2009) , [[0912.0201](#)].
- [12] CMB-S4 collaboration, K. N. Abazajian et al., *CMB-S4 Science Book, First Edition*, [1610.02743](#).
- [13] PLANCK collaboration, P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13, [[1502.01589](#)].

- [14] A. J. Cuesta et al., *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in the correlation function of LOWZ and CMASS galaxies in Data Release 12*, *Mon. Not. Roy. Astron. Soc.* **457** (2016) 1770–1785, [[1509.06371](#)].
- [15] B. Moore, S. Ghigna, F. Governato, G. Lake, T. R. Quinn, J. Stadel et al., *Dark matter substructure within galactic halos*, *Astrophys. J.* **524** (1999) L19–L22, [[astro-ph/9907411](#)].
- [16] A. A. Klypin, A. V. Kravtsov, O. Valenzuela and F. Prada, *Where are the missing Galactic satellites?*, *Astrophys. J.* **522** (1999) 82–92, [[astro-ph/9901240](#)].
- [17] J. Diemand, M. Kuhlen, P. Madau, M. Zemp, B. Moore, D. Potter et al., *Clumps and streams in the local dark matter distribution*, *Nature* **454** (2008) 735–738, [[0805.1244](#)].
- [18] V. Springel, J. Wang, M. Vogelsberger, A. Ludlow, A. Jenkins, A. Helmi et al., *The Aquarius Project: the subhalos of galactic halos*, *Mon. Not. Roy. Astron. Soc.* **391** (2008) 1685–1711, [[0809.0898](#)].
- [19] W. J. G. de Blok, F. Walter, E. Brinks, C. Trachternach, S.-H. Oh and R. C. Kennicutt, Jr., *High-Resolution Rotation Curves and Galaxy Mass Models from THINGS*, *Astron. J.* **136** (2008) 2648–2719, [[0810.2100](#)].
- [20] W. J. G. de Blok, *The Core-Cusp Problem*, *Advances in Astronomy* **2010** (2010) 5, [[0910.3538](#)].
- [21] D. N. Spergel and P. J. Steinhardt, *Observational evidence for selfinteracting cold dark matter*, *Phys. Rev. Lett.* **84** (2000) 3760–3763, [[astro-ph/9909386](#)].
- [22] B. D. Wandelt, R. Dave, G. R. Farrar, P. C. McGuire, D. N. Spergel and P. J. Steinhardt, *Selfinteracting dark matter*, in *Sources and detection of dark matter and dark energy in the universe. Proceedings, 4th International Symposium, DM 2000, Marina del Rey, USA, February 23-25, 2000*, pp. 263–274, 2000. [astro-ph/0006344](#).
- [23] M. Rocha, A. H. G. Peter, J. S. Bullock, M. Kaplinghat, S. Garrison-Kimmel, J. Onorbe et al., *Cosmological Simulations with Self-Interacting Dark Matter I: Constant Density Cores and Substructure*, *Mon. Not. Roy. Astron. Soc.* **430** (2013) 81–104, [[1208.3025](#)].
- [24] A. H. G. Peter, M. Rocha, J. S. Bullock and M. Kaplinghat, *Cosmological Simulations with Self-Interacting Dark Matter II: Halo Shapes vs. Observations*, *Mon. Not. Roy. Astron. Soc.* **430** (2013) 105, [[1208.3026](#)].
- [25] J. Zavala, M. Vogelsberger and M. G. Walker, *Constraining Self-Interacting Dark Matter with the Milky Way’s dwarf spheroidals*, *Monthly Notices of the Royal Astronomical Society: Letters* **431** (2013) L20–L24, [[1211.6426](#)].
- [26] M. Vogelsberger, J. Zavala and A. Loeb, *Subhaloes in Self-Interacting Galactic Dark Matter Haloes*, *Mon. Not. Roy. Astron. Soc.* **423** (2012) 3740, [[1201.5892](#)].
- [27] S. Y. Kim, A. H. G. Peter and D. Wittman, *In the Wake of Dark Giants: New Signatures of Dark Matter Self Interactions in Equal Mass Mergers of Galaxy Clusters*, **1608.08630**.

- [28] S. More et al., *Detection of the Splashback Radius and Halo Assembly bias of Massive Galaxy Clusters*, *Astrophys. J.* **825** (2016) 39, [[1601.06063](#)].
- [29] E. Baxter, C. Chang, B. Jain, S. Adhikari, N. Dalal, A. Kravtsov et al., *The Halo Boundary of Galaxy Clusters in the SDSS*, [1702.01722](#).
- [30] B. Diemer and A. V. Kravtsov, *Dependence of the outer density profiles of halos on their mass accretion rate*, *Astrophys. J.* **789** (2014) 1, [[1401.1216](#)].
- [31] S. Adhikari, N. Dalal and R. T. Chamberlain, *Splashback in accreting dark matter halos*, *JCAP* **1411** (2014) 019, [[1409.4482](#)].
- [32] D. Harvey, R. Massey, T. Kitching, A. Taylor and E. Tittley, *The non-gravitational interactions of dark matter in colliding galaxy clusters*, *Science* **347** (2015) 1462–1465, [[1503.07675](#)].
- [33] R. Massey et al., *The behaviour of dark matter associated with four bright cluster galaxies in the 10 kpc core of Abell 3827*, *Mon. Not. Roy. Astron. Soc.* **449** (2015) 3393–3406, [[1504.03388](#)].
- [34] F. Kahlhoefer, K. Schmidt-Hoberg, M. T. Frandsen and S. Sarkar, *Colliding clusters and dark matter self-interactions*, *Mon. Not. Roy. Astron. Soc.* **437** (2014) 2865–2881, [[1308.3419](#)].
- [35] P. J. E. Peebles, *The large-scale structure of the universe*. 1980.
- [36] PLANCK collaboration, P. A. R. Ade et al., *Planck 2013 results. XVI. Cosmological parameters*, *Astron. Astrophys.* **571** (2014) A16, [[1303.5076](#)].
- [37] H. Mo, F. C. van den Bosch and S. White, *Galaxy Formation and Evolution*. May, 2010.
- [38] K. Heitmann, P. M. Ricker, M. S. Warren and S. Habib, *Robustness of cosmological simulations. 1. Large scale structure*, *Astrophys. J. Suppl.* **160** (2005) 28–58, [[astro-ph/0411795](#)].
- [39] K. Heitmann, Z. Lukić, P. Fasel, S. Habib, M. S. Warren, M. White et al., *The cosmic code comparison project*, *Computational Science and Discovery* **1** (Oct., 2008) 015003, [[0706.1270](#)].
- [40] A. Jenkins, C. S. Frenk, S. D. M. White, J. M. Colberg, S. Cole, A. E. Evrard et al., *The Mass function of dark matter halos*, *Mon. Not. Roy. Astron. Soc.* **321** (2001) 372, [[astro-ph/0005260](#)].
- [41] J. L. Tinker, A. V. Kravtsov, A. Klypin, K. Abazajian, M. S. Warren, G. Yepes et al., *Toward a halo mass function for precision cosmology: The Limits of universality*, *Astrophys. J.* **688** (2008) 709–728, [[0803.2706](#)].
- [42] M. S. Warren, K. Abazajian, D. E. Holz and L. Teodoro, *Precision determination of the mass function of dark matter halos*, *Astrophys. J.* **646** (2006) 881–885, [[astro-ph/0506395](#)].
- [43] J. F. Navarro, C. S. Frenk and S. D. M. White, *A Universal density profile from hierarchical clustering*, *Astrophys. J.* **490** (1997) 493–508, [[astro-ph/9611107](#)].

- [44] J. M. Colberg, R. K. Sheth, A. Diaferio, L. Gao and N. Yoshida, *Voids in a lambda-CDM universe*, *Mon. Not. Roy. Astron. Soc.* **360** (2005) 216–226, [[astro-ph/0409162](#)].
- [45] M. Zemp, *The Structure of Cold Dark Matter Halos: Recent Insights from High Resolution Simulations*, *Modern Physics Letters A* **24** (2009) 2291–2305, [[0909.4298](#)].
- [46] J. Diemand and B. Moore, *The Structure and Evolution of Cold Dark Matter Halos*, *Advanced Science Letters* **4** (Feb., 2011) 297–310, [[0906.4340](#)].
- [47] J. Stadel, D. Potter, B. Moore, J. Diemand, P. Madau, M. Zemp et al., *Quantifying the heart of darkness with GHALO - a multi-billion particle simulation of our galactic halo*, *Mon. Not. Roy. Astron. Soc.* **398** (2009) L21–L25, [[0808.2981](#)].
- [48] V. Springel, S. D. White, A. Jenkins, C. S. Frenk, N. Yoshida, L. Gao et al., *Simulations of the formation, evolution and clustering of galaxies and quasars*, *nature* **435** (2005) 629–636.
- [49] L. Gao, V. Springel and S. D. M. White, *The Age dependence of halo clustering*, *Mon. Not. Roy. Astron. Soc.* **363** (2005) L66–L70, [[astro-ph/0506510](#)].
- [50] J. L. Tinker, B. E. Robertson, A. V. Kravtsov, A. Klypin, M. S. Warren, G. Yepes et al., *The large-scale bias of dark matter halos: numerical calibration and model tests*, *The Astrophysical Journal* **724** (2010) 878.
- [51] K. Eguchi, S. Enomoto, K. Furuno, J. Goldman, H. Hanada, H. Ikeda et al., *First results from kamland: evidence for reactor antineutrino disappearance*, *Physical Review Letters* **90** (2003) 021802.
- [52] F. An, J. Bai, A. Balantekin, H. Band, D. Beavis, W. Beriguete et al., *Observation of electron-antineutrino disappearance at daya bay*, *Physical Review Letters* **108** (2012) 171803.
- [53] J. Lesgourgues and S. Pastor, *Neutrino cosmology and Planck*, *New Journal of Physics* **16** (June, 2014) 065002, [[1404.1740](#)].
- [54] A. G. Doroshkevich, Y. B. Zeldovich, R. A. Sunyaev and M. Y. Khlopov, *Astrophysical Implications of the Neutrino Rest Mass - Part Two - the Density Perturbation Spectrum and Small-Scale Fluctuations in the MICROWAVE_ Background*, *Soviet Astronomy Letters* **6** (Apr., 1980) 252–259.
- [55] A. G. Doroshkevich, M. I. Khlopov, R. A. Sunyaev, A. S. Szalay and I. B. Zeldovich, *Cosmological impact of the neutrino rest mass*, *Annals of the New York Academy of Sciences* **375** (Dec., 1981) 32–42.
- [56] J. Brandbyge and S. Hannestad, *Grid Based Linear Neutrino Perturbations in Cosmological N-body Simulations*, *JCAP* **0905** (2009) 002, [[0812.3149](#)].
- [57] M. Archidiacono, T. Brinckmann, J. Lesgourgues and V. Poulin, *Physical effects involved in the measurements of neutrino masses with future cosmological data*, [1610.09852](#).
- [58] Y. Ali-Haïmoud and S. Bird, *An efficient implementation of massive neutrinos in non-linear structure formation simulations*, *Mon. Not. Roy. Astron. Soc.* **428** (2012) 3375–3389, [[1209.0461](#)].

- [59] F. Villaescusa-Navarro, F. Marulli, M. Viel, E. Branchini, E. Castorina, E. Sefusatti et al., *Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies*, *JCAP* **1403** (2014) 011, [[1311.0866](#)].
- [60] J. Brandbyge and S. Hannestad, *Resolving Cosmic Neutrino Structure: A Hybrid Neutrino N-body Scheme*, *JCAP* **1001** (2010) 021, [[0908.1969](#)].
- [61] F. Villaescusa-Navarro, M. Vogelsberger, M. Viel and A. Loeb, *Neutrino Signatures on the High Transmission Regions of the Lyman-alpha Forest*, *ArXiv e-prints* (June, 2011) , [[1106.2543](#)].
- [62] F. Villaescusa-Navarro, S. Bird, C. Pena-Garay and M. Viel, *Non-linear evolution of the cosmic neutrino background*, *JCAP* **1303** (2013) 019, [[1212.4855](#)].
- [63] J. Brandbyge, S. Hannestad, T. Haugbølle and Y. Y. Y. Wong, *Neutrinos in Non-linear Structure Formation - The Effect on Halo Properties*, *JCAP* **1009** (2010) 014, [[1004.4105](#)].
- [64] E. Castorina, E. Sefusatti, R. K. Sheth, F. Villaescusa-Navarro and M. Viel, *Cosmology with massive neutrinos II: on the universality of the halo mass function and bias*, *JCAP* **1402** (2014) 049, [[1311.1212](#)].
- [65] M. Costanzi, F. Villaescusa-Navarro, M. Viel, J.-Q. Xia, S. Borgani, E. Castorina et al., *Cosmology with massive neutrinos III: the halo mass function and an application to galaxy clusters*, *JCAP* **1312** (2013) 012, [[1311.1514](#)].
- [66] D. Inman, J. D. Emberson, U.-L. Pen, A. Farchi, H.-R. Yu and J. Harnois-Déraps, *Precision reconstruction of the cold dark matter-neutrino relative velocity from N-body simulations*, *Phys. Rev.* **D92** (2015) 023502, [[1503.07480](#)].
- [67] E. Castorina, C. Carbone, J. Bel, E. Sefusatti and K. Dolag, *DEMNUi: The clustering of large-scale structures in the presence of massive neutrinos*, *JCAP* **1507** (2015) 043, [[1505.07148](#)].
- [68] C. Carbone, M. Petkova and K. Dolag, *DEMNUi: ISW, Rees-Sciama, and weak-lensing in the presence of massive neutrinos*, *ArXiv e-prints* (May, 2016) , [[1605.02024](#)].
- [69] P. Bode, J. P. Ostriker and N. Turok, *Halo formation in warm dark matter models*, *Astrophys. J.* **556** (2001) 93–107, [[astro-ph/0010389](#)].
- [70] F. Villaescusa-Navarro and N. Dalal, *Cores and Cusps in Warm Dark Matter Halos*, *JCAP* **1103** (2011) 024, [[1010.3008](#)].
- [71] V. Avila-Reese, P. Colin, O. Valenzuela, E. D’Onghia and C. Firmani, *Formation and structure of halos in a warm dark matter cosmology*, *Astrophys. J.* **559** (2001) 516–530, [[astro-ph/0010525](#)].
- [72] J. Wang and S. D. M. White, *Discreteness effects in simulations of Hot/Warm dark matter*, *Mon. Not. Roy. Astron. Soc.* **380** (2007) 93–103, [[astro-ph/0702575](#)].
- [73] R. E. Angulo, O. Hahn and T. Abel, *The Warm DM halo mass function below the cut-off scale*, *Mon. Not. Roy. Astron. Soc.* **434** (2013) 3337, [[1304.2406](#)].

- [74] Y. Hezaveh, N. Dalal, G. Holder, M. Kuhlen, D. Marrone, N. Murray et al., *Dark Matter Substructure Detection Using Spatially Resolved Spectroscopy of Lensed Dusty Galaxies*, *Astrophys. J.* **767** (Apr., 2013) 9, [[1210.4562](#)].
- [75] Y. D. Hezaveh et al., *Detection of lensing substructure using ALMA observations of the dusty galaxy SDP.81*, *Astrophys. J.* **823** (2016) 37, [[1601.01388](#)].
- [76] A. Hobbs, J. Read, O. Agertz, F. Iannuzzi and C. Power, *Novel Adaptive softening for collisionless N-body simulations: eliminating spurious haloes*, *Mon. Not. Roy. Astron. Soc.* **458** (2016) 468–479, [[1503.02689](#)].
- [77] J. J. Monaghan, *Smoothed particle hydrodynamics*, *Annual review of astronomy and astrophysics* **30** (1992) 543–574.
- [78] R. W. Hockney and J. W. Eastwood, *Computer simulation using particles*. CRC Press, 1988.
- [79] V. Springel, *The Cosmological simulation code GADGET-2*, *Mon. Not. Roy. Astron. Soc.* **364** (2005) 1105–1134, [[astro-ph/0505010](#)].
- [80] E. F. Toro, *Riemann solvers and numerical methods for fluid dynamics: a practical introduction*. Springer Science & Business Media, 2009.
- [81] A. Harten, *High resolution schemes for hyperbolic conservation laws*, *Journal of computational physics* **49** (1983) 357–393.
- [82] P. Roe, *Characteristic-based schemes for the euler equations*, *Annual review of fluid mechanics* **18** (1986) 337–365.
- [83] J. Binney and S. Tremaine, *Galactic dynamics*. Princeton university press, 2011.
- [84] B. Van Leer, *Towards the ultimate conservative difference scheme. iv. a new approach to numerical convection*, *Journal of computational physics* **23** (1977) 276–299.
- [85] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, *ArXiv e-prints* (Apr., 2011) , [[1104.2932](#)].
- [86] D. Blas, J. Lesgourgues and T. Tram, *The cosmic linear anisotropy solving system (class). part ii: Approximation schemes*, *Journal of Cosmology and Astroparticle Physics* **2011** (2011) 034.
- [87] J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics*, *JCAP* **1109** (2011) 032, [[1104.2935](#)].
- [88] N. E. Chisari and M. Zaldarriaga, *Connection between Newtonian simulations and general relativity*, *Phys. Rev.* **D83** (2011) 123505, [[1101.3555](#)].
- [89] P. Behroozi, R. Wechsler and H.-Y. Wu, “Rockstar: Phase-space halo finder.” Astrophysics Source Code Library, Oct., 2012.
- [90] M. LoVerde, *Halo bias in mixed dark matter cosmologies*, *Phys. Rev.* **D90** (2014) 083530, [[1405.4855](#)].

- [91] M. LoVerde, *Neutrino mass without cosmic variance*, *Phys. Rev.* **D93** (2016) 103526, [[1602.08108](#)].
- [92] E. Massara, F. Villaescusa-Navarro, M. Viel and P. M. Sutter, *Voids in massive neutrino cosmologies*, *JCAP* **1511** (2015) 018, [[1506.03088](#)].
- [93] R. J. Scherrer and D. H. Weinberg, *Constraints on the effects of locally-biased galaxy formation*, *Astrophys. J.* **504** (1998) 607–611, [[astro-ph/9712192](#)].
- [94] N. Dalal, O. Dore, D. Huterer and A. Shirokov, *The imprints of primordial non-gaussianities on large-scale structure: scale dependent bias and abundance of virialized objects*, *Phys. Rev.* **D77** (2008) 123514, [[0710.4560](#)].
- [95] B. Jain and J. Khoury, *Cosmological tests of gravity*, *Annals of Physics* **325** (July, 2010) 1479–1516, [[1004.3294](#)].
- [96] J. A. Fillmore and P. Goldreich, *Self-similar spherical voids in an expanding universe*, *The Astrophysical Journal* **281** (1984) 9–12.
- [97] E. Jennings, Y. Li and W. Hu, *The abundance of voids and the excursion set formalism*, [1304.6087](#).
- [98] C. Sánchez, J. Clampitt, A. Kovacs, B. Jain, J. García-Bellido, S. Nadathur et al., *Cosmic Voids and Void Lensing in the Dark Energy Survey Science Verification Data*, *ArXiv e-prints* (May, 2016) , [[1605.03982](#)].
- [99] J. Clampitt, B. Jain and C. Sánchez, *Clustering and Bias Measurements of SDSS Voids*, *Mon. Not. Roy. Astron. Soc.* **456** (2016) 4425–4431, [[1507.08031](#)].
- [100] A. Upadhye, J. Kwan, A. Pope, K. Heitmann, S. Habib, H. Finkel et al., *Redshift-space distortions in massive neutrino and evolving dark energy cosmologies*, *Phys. Rev. D* **93** (Mar, 2016) 063515.
- [101] F. Marulli, C. Carbone, M. Viel, L. Moscardini and A. Cimatti, *Effects of massive neutrinos on the large-scale structure of the universe*, *Monthly Notices of the Royal Astronomical Society* **418** (2011) 346.
- [102] N. Arkani-Hamed, T. Cohen, R. T. D’Agnolo, A. Hook, H. D. Kim and D. Pinner, *Nnaturalness*, [1607.06821](#).
- [103] D. J. Eisenstein and W. Hu, *Power spectra for cold dark matter and its variants*, *Astrophys. J.* **511** (1997) 5, [[astro-ph/9710252](#)].
- [104] F.-Y. Cyr-Racine and K. Sigurdson, *Cosmology of atomic dark matter*, *Phys. Rev.* **D87** (2013) 103515, [[1209.5752](#)].
- [105] J. Fan, A. Katz, L. Randall and M. Reece, *Double-Disk Dark Matter*, *Phys. Dark Univ.* **2** (2013) 139–156, [[1303.1521](#)].
- [106] L. Ackerman, M. R. Buckley, S. M. Carroll and M. Kamionkowski, *Dark Matter and Dark Radiation*, *Phys. Rev.* **D79** (2009) 023519, [[0810.5126](#)].

- [107] J. L. Feng, M. Kaplinghat, H. Tu and H.-B. Yu, *Hidden Charged Dark Matter*, *JCAP* **0907** (2009) 004, [[0905.3039](#)].
- [108] S. Tulin, H.-B. Yu and K. M. Zurek, *Beyond Collisionless Dark Matter: Particle Physics Dynamics for Dark Matter Halo Structure*, *Phys. Rev.* **D87** (2013) 115007, [[1302.3898](#)].
- [109] E. D. Carlson, M. E. Machacek and L. J. Hall, *Self-interacting dark matter*, *Astrophys. J.* **398** (1992) 43–52.
- [110] Y. Hochberg, E. Kuflik, T. Volansky and J. G. Wacker, *Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles*, *Phys. Rev. Lett.* **113** (2014) 171301, [[1402.5143](#)].
- [111] J. Kopp, J. Liu, T. R. Slatyer, X.-P. Wang and W. Xue, *Impeded Dark Matter*, *JHEP* **12** (2016) 033, [[1609.02147](#)].
- [112] J. A. Dror, E. Kuflik and W. H. Ng, *Co-Decaying Dark Matter*, *Phys. Rev. Lett.* **117** (2016) 211801, [[1607.03110](#)].
- [113] L. Forestell, D. E. Morrissey and K. Sigurdson, *Non-Abelian Dark Forces and the Relic Densities of Dark Glueballs*, [1605.08048](#).
- [114] D. Pappadopulo, J. T. Ruderman and G. Trevisan, *Dark matter freeze-out in a nonrelativistic sector*, *Phys. Rev.* **D94** (2016) 035005, [[1602.04219](#)].
- [115] E. Kuflik, M. Perelstein, N. R.-L. Lorier and Y.-D. Tsai, *Elastically Decoupling Dark Matter*, *Phys. Rev. Lett.* **116** (2016) 221302, [[1512.04545](#)].
- [116] K. Agashe, Y. Cui, L. Necib and J. Thaler, *(In)direct Detection of Boosted Dark Matter*, *JCAP* **1410** (2014) 062, [[1405.7370](#)].
- [117] A. Loeb and N. Weiner, *Cores in Dwarf Galaxies from Dark Matter with a Yukawa Potential*, *Phys. Rev. Lett.* **106** (2011) 171302, [[1011.6374](#)].
- [118] Z. Chacko, Y. Cui, S. Hong and T. Okui, *Hidden dark matter sector, dark radiation, and the CMB*, *Phys. Rev.* **D92** (2015) 055033, [[1505.04192](#)].
- [119] S. Hannestad, A. Ringwald, H. Tu and Y. Y. Y. Wong, *Is it possible to tell the difference between fermionic and bosonic hot dark matter?*, *JCAP* **0509** (2005) 014, [[astro-ph/0507544](#)].
- [120] <https://pole.uchicago.edu/>.
- [121] <http://act.princeton.edu/>.
- [122] <https://simonsobservatory.org/>.
- [123] <http://www.darkenergysurvey.org/>.
- [124] <http://subarutelescope.org/Projects/HSC/HSCProject.html>.
- [125] <http://kids.strw.leidenuniv.nl/>.

- [126] R. Laureijs, J. Amiaux, S. Arduini, J. . Auguères, J. Brinchmann, R. Cole et al., *Euclid Definition Study Report, ArXiv e-prints* (Oct., 2011) , [[1110.3193](#)].
- [127] D. Spergel, N. Gehrels, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi et al., *Wide-Field InfraRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA Final Report, ArXiv e-prints* (May, 2013) , [[1305.5422](#)].
- [128] T. Baldauf, M. Mirbabayi, M. Simonović and M. Zaldarriaga, *LSS constraints with controlled theoretical uncertainties*, [1602.00674](#).
- [129] W. Hu, D. J. Eisenstein and M. Tegmark, *Weighing neutrinos with galaxy surveys*, *Phys. Rev. Lett.* **80** (1998) 5255–5258, [[astro-ph/9712057](#)].
- [130] K. N. Abazajian and S. Dodelson, *Neutrino mass and dark energy from weak lensing*, *Phys. Rev. Lett.* **91** (2003) 041301, [[astro-ph/0212216](#)].
- [131] J. Lesgourgues, L. Perotto, S. Pastor and M. Piat, *Probing neutrino masses with cmb lensing extraction*, *Phys. Rev.* **D73** (2006) 045021, [[astro-ph/0511735](#)].
- [132] M. Kaplinghat, L. Knox and Y.-S. Song, *Determining neutrino mass from the CMB alone*, *Phys. Rev. Lett.* **91** (2003) 241301, [[astro-ph/0303344](#)].
- [133] S. Hannestad, H. Tu and Y. Y. Y. Wong, *Measuring neutrino masses and dark energy with weak lensing tomography*, *JCAP* **0606** (2006) 025, [[astro-ph/0603019](#)].
- [134] T. D. Kitching, A. F. Heavens, L. Verde, P. Serra and A. Melchiorri, *Finding Evidence for Massive Neutrinos using 3D Weak Lensing*, *Phys. Rev.* **D77** (2008) 103008, [[0801.4565](#)].
- [135] R. Allison, P. Caucal, E. Calabrese, J. Dunkley and T. Louis, *Towards a cosmological neutrino mass detection*, *Phys. Rev.* **D92** (2015) 123535, [[1509.07471](#)].
- [136] K. Ichiki, M. Takada and T. Takahashi, *Constraints on Neutrino Masses from Weak Lensing*, *Phys. Rev.* **D79** (2009) 023520, [[0810.4921](#)].
- [137] S. Saito, M. Takada and A. Taruya, *Nonlinear power spectrum in the presence of massive neutrinos: Perturbation theory approach, galaxy bias, and parameter forecasts*, *Phys. Rev. D* **80** (Oct, 2009) 083528.
- [138] S. Saito, M. Takada and A. Taruya, *Neutrino mass constraint from the sloan digital sky survey power spectrum of luminous red galaxies and perturbation theory*, *Phys. Rev. D* **83** (Feb, 2011) 043529.
- [139] N. Canac, G. Aslanyan, K. N. Abazajian, R. Easther and L. C. Price, *Testing for New Physics: Neutrinos and the Primordial Power Spectrum*, *JCAP* **1609** (2016) 022, [[1606.03057](#)].
- [140] R. Hložek, D. J. E. Marsh, D. Grin, R. Allison, J. Dunkley and E. Calabrese, *Future CMB tests of dark matter: ultra-light axions and massive neutrinos*, [1607.08208](#).
- [141] B. A. Reid, L. Verde, R. Jimenez and O. Mena, *Robust Neutrino Constraints by Combining Low Redshift Observations with the CMB*, *JCAP* **1001** (2010) 003, [[0910.0008](#)].

- [142] C. Carbone, L. Verde, Y. Wang and A. Cimatti, *Neutrino constraints from future nearly all-sky spectroscopic galaxy surveys*, *JCAP* **3** (Mar., 2011) 030, [[1012.2868](#)].
- [143] A. J. Cuesta, V. Niro and L. Verde, *Neutrino mass limits: robust information from the power spectrum of galaxy surveys*, *Phys. Dark Univ.* **13** (2016) 77–86, [[1511.05983](#)].
- [144] R. Jimenez, T. Kitching, C. Pena-Garay and L. Verde, *Can we measure the neutrino mass hierarchy in the sky?*, *JCAP* **1005** (2010) 035, [[1003.5918](#)].
- [145] J. Hamann, S. Hannestad and Y. Y. Y. Wong, *Measuring neutrino masses with a future galaxy survey*, *JCAP* **1211** (2012) 052, [[1209.1043](#)].
- [146] DESI collaboration, A. Aghamousa et al., *The DESI Experiment Part I: Science, Targeting, and Survey Design*, [1611.00036](#).
- [147] W. L. K. Wu, J. Errard, C. Dvorkin, C. L. Kuo, A. T. Lee, P. McDonald et al., *A Guide to Designing Future Ground-based Cosmic Microwave Background Experiments*, *Astrophys. J.* **788** (2014) 138, [[1402.4108](#)].
- [148] TOPICAL CONVENERS: K.N. ABAZAJIAN, J.E. CARLSTROM, A.T. LEE collaboration, K. N. Abazajian et al., *Neutrino Physics from the Cosmic Microwave Background and Large Scale Structure*, *Astropart. Phys.* **63** (2015) 66–80, [[1309.5383](#)].
- [149] G.-B. Zhao, S. Saito, W. J. Percival, A. J. Ross, F. Montesano, M. Viel et al., *The clustering of galaxies in the sdss-iii baryon oscillation spectroscopic survey: weighing the neutrino mass using the galaxy power spectrum of the cmass sample*, *Monthly Notices of the Royal Astronomical Society* **436** (2013) 2038–2053.
- [150] C. Carbone, C. Fedeli, L. Moscardini and A. Cimatti, *Measuring the neutrino mass from future wide galaxy cluster catalogues*, *Journal of Cosmology and Astroparticle Physics* **2012** (2012) 023.
- [151] S. Gratton, A. Lewis and G. Efstathiou, *Prospects for Constraining Neutrino Mass Using Planck and Lyman-Alpha Forest Data*, *Phys. Rev.* **D77** (2008) 083507, [[0705.3100](#)].
- [152] S. Wang, Z. Haiman, W. Hu, J. Khoury and M. May, *Weighing neutrinos with galaxy cluster surveys*, *Phys. Rev. Lett.* **95** (2005) 011302, [[astro-ph/0505390](#)].
- [153] Y. Takeuchi and K. Kadota, *Probing Neutrinos from Planck and Forthcoming Galaxy Redshift Surveys*, *JCAP* **1401** (2014) 046, [[1310.0037](#)].
- [154] A. Font-Ribera, P. McDonald, N. Mostek, B. A. Reid, H.-J. Seo and A. Slosar, *DESI and other dark energy experiments in the era of neutrino mass measurements*, *JCAP* **1405** (2014) 023, [[1308.4164](#)].
- [155] A. C. Hall and A. Challinor, *Probing the neutrino mass hierarchy with cosmic microwave background weak lensing*, *Monthly Notices of the Royal Astronomical Society* **425** (2012) 1170–1184.
- [156] R. de Putter, E. V. Linder and A. Mishra, *Inflationary Freedom and Cosmological Neutrino Constraints*, *Phys. Rev.* **D89** (2014) 103502, [[1401.7022](#)].

- [157] R. de Putter, O. Zahn and E. V. Linder, *Cmb lensing constraints on neutrinos and dark energy*, *Phys. Rev. D* **79** (Mar, 2009) 065033.
- [158] O. Lahav, A. Kiakotou, F. B. Abdalla and C. Blake, *Forecasting neutrino masses from galaxy clustering in the Dark Energy Survey combined with the Planck measurements*, *Monthly Notices of the Royal Astronomical Society* **405** (2010) 168–176, [0910.4714].
- [159] E.-M. Mueller, F. de Bernardis, R. Bean and M. D. Niemack, *Constraints on massive neutrinos from the pairwise kinematic Sunyaev-Zel’dovich effect*, *Phys. Rev. D* **92** (2015) 063501, [1412.0592].
- [160] W. Hu and B. Jain, *Joint galaxy - lensing observables and the dark energy*, *Phys. Rev. D* **70** (2004) 043009, [astro-ph/0312395].
- [161] A. Lewis and A. Challinor, “CAMB.” <http://camb.info/>.
- [162] W. Hu and T. Okamoto, *Mass reconstruction with cmb polarization*, *Astrophys. J.* **574** (2002) 566–574, [astro-ph/0111606].
- [163] D. Hanson, “QUICKLENS.” <https://github.com/dhanson/quicklens>.
- [164] E. Rozo, E. S. Rykoff, A. Abate, C. Bonnett, M. Crocce, C. Davis et al., *redMaGiC: selecting luminous red galaxies from the DES Science Verification data*, *Mon. Not. Roy. Astron. Soc.* **461** (Sept., 2016) 1431–1450, [1507.05460].
- [165] DES collaboration, J. Kwan et al., *Cosmology from large scale galaxy clustering and galaxy-galaxy lensing with Dark Energy Survey Science Verification data*, *Mon. Not. Roy. Astron. Soc.* (2016) , [1604.07871].
- [166] C. M. Hirata and U. c. v. Seljak, *Intrinsic alignment-lensing interference as a contaminant of cosmic shear*, *Phys. Rev. D* **70** (Sep, 2004) 063526.
- [167] M. A. Troxel and M. Ishak, *The Intrinsic Alignment of Galaxies and its Impact on Weak Gravitational Lensing in an Era of Precision Cosmology*, *Phys. Rept.* **558** (2014) 1–59, [1407.6990].
- [168] C. Chang, S. Kahn, J. Jernigan, J. Peterson, Y. AlSayyad, Z. Ahmad et al., *Spurious shear in weak lensing with the large synoptic survey telescope*, *Monthly Notices of the Royal Astronomical Society* **428** (2013) 2695–2713.
- [169] E. Schaan, E. Krause, T. Eifler, O. Doré, H. Miyatake, J. Rhodes et al., *Looking through the same lens: shear calibration for LSST, Euclid and WFIRST with stage 4 CMB lensing*, *ArXiv e-prints* (July, 2016) , [1607.01761].
- [170] E. Krause, T. Eifler and J. Blazek, *The impact of intrinsic alignment on current and future cosmic shear surveys*, *Mon. Not. Roy. Astron. Soc.* **456** (2016) 207–222, [1506.08730].
- [171] Joachimi, B. and Bridle, S. L., *Simultaneous measurement of cosmology and intrinsic alignments using joint cosmic shear and galaxy number density correlations*, *Astron. Astrophys.* **523** (2010) .
- [172] A. Hall and A. Taylor, *Intrinsic alignments in the cross-correlation of cosmic shear and CMB weak lensing*, *Mon. Not. Roy. Astron. Soc.* **443** (2014) L119, [1401.6018].

- [173] B. Jain et al., *The Whole is Greater than the Sum of the Parts: Optimizing the Joint Science Return from LSST, Euclid and WFIRST*, [1501.07897](#).
- [174] S. Das and D. N. Spergel, *Measuring distance ratios with CMB-galaxy lensing cross-correlations*, *Phys. Rev.* (Feb., 2009) 043509, [[0810.3931](#)].
- [175] S. Weinberg, *A New Light Boson?*, *Phys. Rev. Lett.* **40** (1978) 223–226.
- [176] F. Wilczek, *Problem of Strong p and t Invariance in the Presence of Instantons*, *Phys. Rev. Lett.* **40** (1978) 279–282.
- [177] Y. Chikashige, R. N. Mohapatra and R. D. Peccei, *Are There Real Goldstone Bosons Associated with Broken Lepton Number?*, *Phys. Lett.* **B98** (1981) 265–268.
- [178] G. B. Gelmini and M. Roncadelli, *Left-Handed Neutrino Mass Scale and Spontaneously Broken Lepton Number*, *Phys. Lett.* **B99** (1981) 411–415.
- [179] F. Wilczek, *Axions and Family Symmetry Breaking*, *Phys. Rev. Lett.* **49** (1982) 1549–1552.
- [180] S. Weinberg, *Goldstone Bosons as Fractional Cosmic Neutrinos*, *Phys. Rev. Lett.* **110** (2013) 241301, [[1305.1971](#)].
- [181] P. Svrcek and E. Witten, *Axions In String Theory*, *JHEP* **06** (2006) 051, [[hep-th/0605206](#)].
- [182] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper and J. March-Russell, *String Axiverse*, *Phys. Rev.* **D81** (2010) 123530, [[0905.4720](#)].
- [183] J. Jaeckel and A. Ringwald, *The Low-Energy Frontier of Particle Physics*, *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 405–437, [[1002.0329](#)].
- [184] M. Goodsell, J. Jaeckel, J. Redondo and A. Ringwald, *Naturally Light Hidden Photons in LARGE Volume String Compactifications*, *JHEP* **11** (2009) 027, [[0909.0515](#)].
- [185] M. Cicoli, M. Goodsell, J. Jaeckel and A. Ringwald, *Testing String Vacua in the Lab: From a Hidden CMB to Dark Forces in Flux Compactifications*, *JHEP* **07** (2011) 114, [[1103.3705](#)].
- [186] A. Arvanitaki, N. Craig, S. Dimopoulos, S. Dubovsky and J. March-Russell, *String Photini at the LHC*, *Phys. Rev.* **D81** (2010) 075018, [[0909.5440](#)].
- [187] N. Engelhardt, A. E. Nelson and J. R. Walsh, *Apparent CPT Violation in Neutrino Oscillation Experiments*, *Phys. Rev.* **D81** (2010) 113001, [[1002.4452](#)].
- [188] L. A. Anchordoqui and H. Goldberg, *Neutrino cosmology after WMAP 7-Year data and LHC first Z' bounds*, *Phys. Rev. Lett.* **108** (2012) 081805, [[1111.7264](#)].
- [189] B. Dasgupta and J. Kopp, *Cosmologically Safe eV -Scale Sterile Neutrinos and Improved Dark Matter Structure*, *Phys. Rev. Lett.* **112** (2014) 031803, [[1310.6337](#)].
- [190] P. Ko and Y. Tang, *$\nu\Lambda$ MDM: A model for sterile neutrino and dark matter reconciles cosmological and neutrino oscillation data after BICEP2*, *Phys. Lett.* **B739** (2014) 62–67, [[1404.0236](#)].

- [191] R. Foot, *Mirror dark matter: Cosmology, galaxy structure and direct detection*, *Int. J. Mod. Phys. A* **29** (2014) 1430013, [[1401.3965](#)].
- [192] Z. Chacko, H.-S. Goh and R. Harnik, *The Twin Higgs: Natural electroweak breaking from mirror symmetry*, *Phys. Rev. Lett.* **96** (2006) 231802, [[hep-ph/0506256](#)].
- [193] M. Pospelov, A. Ritz and M. B. Voloshin, *Secluded WIMP Dark Matter*, *Phys. Lett.* **B662** (2008) 53–61, [[0711.4866](#)].
- [194] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer and N. Weiner, *A Theory of Dark Matter*, *Phys. Rev.* **D79** (2009) 015014, [[0810.0713](#)].
- [195] K. M. Zurek, *Asymmetric Dark Matter: Theories, Signatures, and Constraints*, *Phys. Rept.* **537** (2014) 91–121, [[1308.0338](#)].
- [196] P. Adshead, Y. Cui and J. Shelton, *Chilly Dark Sectors and Asymmetric Reheating*, *JHEP* **06** (2016) 016, [[1604.02458](#)].
- [197] J. L. Feng and Y. Shadmi, *WIMPlless Dark Matter from Non-Abelian Hidden Sectors with Anomaly-Mediated Supersymmetry Breaking*, *Phys. Rev.* **D83** (2011) 095011, [[1102.0282](#)].
- [198] D. Blas, J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes*, *JCAP* **7** (July, 2011) 034, [[1104.2933](#)].
- [199] J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics*, *JCAP* **1109** (2011) 032, [[1104.2935](#)].
- [200] T. Eifler, E. Krause, S. Dodelson, A. Zentner, A. Hearin and N. Gnedin, *Accounting for baryonic effects in cosmic shear tomography: Determining a minimal set of nuisance parameters using PCA*, *Mon. Not. Roy. Astron. Soc.* **454** (2015) 2451–2471, [[1405.7423](#)].
- [201] A. R. Zentner, E. Semboloni, S. Dodelson, T. Eifler, E. Krause and A. P. Hearin, *Accounting for baryons in cosmological constraints from cosmic shear*, *Phys. Rev. D* **87** (Feb, 2013) 043509.
- [202] N. Bellomo, E. Bellini, B. Hu, R. Jimenez, C. Pena-Garay and L. Verde, *Hiding neutrino mass in modified gravity cosmologies*, [1612.02598](#).
- [203] E. Giusarma, M. Gerbino, O. Mena, S. Vagnozzi, S. Ho and K. Freese, *Improvement of cosmological neutrino mass bounds*, *Phys. Rev.* **D94** (2016) 083522, [[1605.04320](#)].
- [204] A. Zablacki, *Constraining neutrinos and dark energy with galaxy clustering in the dark energy survey*, *Phys. Rev. D* **94** (Aug, 2016) 043525.
- [205] N. Palanque-Delabrouille et al., *Constraint on neutrino masses from SDSS-III/BOSS Ly α forest and other cosmological probes*, *JCAP* **1502** (2015) 045, [[1410.7244](#)].
- [206] S. Dodelson, K. Heitmann, C. Hirata, K. Honscheid, A. Roodman, U. Seljak et al., *Cosmic Visions Dark Energy: Science*, [1604.07626](#).
- [207] R. de Putter, O. Doré and M. Takada, *The Synergy between Weak Lensing and Galaxy Redshift Surveys*, [1308.6070](#).

- [208] A. Robertson, R. Massey and V. Eke, *Cosmic particle colliders: simulations of self-interacting dark matter with anisotropic scattering*, *Mon. Not. Roy. Astron. Soc.* **467** (2017) 4719–4730, [[1612.03906](#)].