Dark Begets Light: Exploring Physics Beyond the Standard Model with Cosmology

Citation

DePorzio, Nicholas Anthony. 2023. Dark Begets Light: Exploring Physics Beyond the Standard Model with Cosmology. Doctoral dissertation, Harvard University Graduate School of Arts and Sciences.

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Physics

have examined a dissertation entitled

Dark Begets Light: Exploring Physics Beyond the Standard Model with Cosmology

presented by Nicholas Anthony DePorzio

candidate for the degree of Doctor of Philosophy and hereby
certify that it is worthy of acceptance.

Signature ____
Typed name: Professor Cora Dvorkin, Chair

Signature ____
Typed name: Professor Lisa Randall, Co-Chair

Signature ____
Typed name: Professor Roxanne Guenette (Univ. of Manchester)

Signature ____
Typed name: Professor John Kovac

Date: May 3, 2023

Dark Begets Light: Exploring Physics Beyond the Standard Model with Cosmology

a dissertation presented
by
Nicholas Anthony DePorzio
to
The Department of Physics

in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of
Physics

Harvard University
Cambridge, Massachusetts
May 2023

Dissertation Advisor: Professor Cora Dvorkin
Dissertation Advisor: Professor Lisa Randall

Nicholas Anthony DePorzio

Dark Begets Light: Exploring Physics Beyond the Standard Model with Cosmology

Abstract

After nearly a century of inquiry, the particle nature of dark matter remains unknown. As a wide array of Dark Matter (DM) phenomenologies map onto similar variations in cosmological observables, cosmology places far-reaching constraints on the theoretical parameter space of Dark Matter models beyond the Standard Model (SM) of particle physics. This dissertation presents the results of several inquiries that demonstrate the use of cosmology to measure the parameter space for different classes of DM models assuming only gravitational interactions. This dissertation first considers an extension of ΛCDM involving additional particle degrees of freedom of non-zero mass in early thermal contact with the SM, which are called Light but Massive Relics (LiMRs). LiMRs introduce a characteristic scale into the cosmology through their free-streaming while relativistic. Effects on the distribution of galaxies and the Cosmic Microwave Background (CMB) are explored and used to impose constraints on the mass, abundance, and degrees of freedom of LiMRs. Implications for measuring the mass and hierarchy of massive neutrinos, a special case of LiMR, are considered. Representing another class of DM models, ultralight axion-like particles (ALPs), like LiMRs, introduce characteristic scales into the cosmology: one set by the oscillation time of the field; another characterized by the macroscopic wavelength of the field, the Jeans Scale. The nonlinear connection between matter perturbations and halo perturbations, the halo bias, is modelled in the presence LiMRs, neutrinos, and ultralight ALPs. Gravitational waves (GWs), a burgeoning source of cosmological information, motivate the concluding study of this dissertation which considers the degree to which mHz-Hz frequency GW detectors can infer the formation channels of black holes - a process which may be influenced by the introduction of new physics to the SM and ΛCDM.

Contents

Acknowledgments

I have voyaged into the dark, and found safe harbor again - there is no less than a community of individuals to thank for steering me away from the many perils that could have befallen my time as a student. First, my advisors Cora Dvorkin and Lisa Randall have provided no shortage of brilliant examples of how to tactfully helm a life of scientific research. They have challenged me profoundly only to repeatedly demonstrate limits of my own capability I had not yet known. Especially memorable of my interactions with Lisa are the layers of physical insight which would unfold in the period of reflection following our conversations, a hallmark of Lisa's seemingly boundless ability to identify connections between seemingly disparate problems and concepts in physics. What skill I have developed to efficiently inquire into new lines of research has been strongly influenced through exposure to Lisa's own ability to incisively question an argument to the point of its understanding, or demise. My time with Cora is epitomized by her boundless energy to pursue a question, through trial and setback, to its completion – a characteristic which has imbued and enlivened my own efforts. Yet in the rush of progress, Cora has encouraged the exercise of demonstrating the same result through varied channels has often saved me the waste of time which inevitably results from pursuing a false sense of understanding; this careful curation of insight and attention to detail has been ingrained into my scientific character as a result of my exposure to the same qualities in Cora.

I am deeply thankful of the other members of my committee, Roxanne Guénette and Johh Kovac, both for their helpful comments on my research development as well as for complementing my understanding of cosmology and particle physics with lessons from their own research fields. I am also appreciative of Roxanne for the time spent mentoring me in the beginning of my PhD - at a time when I was still wildly unsure of which topics I wanted to pursue, she offered me a foray into experimental particle physics, a welcoming community, a bridge to the rest of the department, and mentorship throughout the entirety of my time at Harvard. For their own unique contributions to my understanding of the science and the culture of the physics community (and how it can be built), I am also immensely grateful for my interactions with Melissa Franklin who I could always count on for brutally fantastic advisement drawn from a lifetime of experience on all aspects of life; to Vinny Manoharan for being a phenomenally lucid lecturer through the obscurity of statistical mechanics and analysis; and to Tracy Slatyer and Jessie Shelton for being deeply thoughtful collaborators, for expanding my research horizons, and for teaching me an incredible amount of physics. The privilege of having grown within the lively scientific environment of the High Energy Theory group also owes much thanks to Matt Reece, Matt Schwartz, Matt Strassler, and Howard Georgi for the examples you set and for the many stories, insights, and laughs that composed some of the best parts of this environment. Likewise,there are a host of other postdocs who played no lesser part in this accomplishment and who I have been immensely fortunate to have interacted with and learned from: Francis-Yan Cyr-Racine, Hayden Lee, David Pinner, Prateek Agrawal, Sid Mishra-Sharma, Georgios Valogiannis, Bryan Ostdiek, Pouya Asadi, Motoo Suzuki, John Stout, Sam Homiller, and Severin Lüst.

My collaborators Julian Muñoz, Linda Xu, Zhong-Zhi Xianyu and Rashmish Mishra have made my moments of turmoil bearable and my triumphs celebratory. Julian and Linda have been with me since my first project, and bore with me through the obscene number of questions necessary to make my first contributions as researcher. I have learned so much from each of you. To Rashmish especially, it is difficult to convey the gravity to which you contributed to my time at Harvard – it is profoundly rare, from my experience, to encounter a person you can place your trust in, advise you both honestly and frankly, and yet never make you feel less than perfectly capable of achieving what you want with well-placed effort. You have been both a friend and a bottomless wealth of physics knowledge in my many moments of uncertainty.

I have also been lucky to have an incredible group of peers throughout the physics community – lifesavers amidst a multitude of stumbles overboard! Thank you to my fellow group members Çağan Şengül, Georges Obied, Ana Díaz Rivero, Chandrika Chandrashekar, Arthur Tsang, Anmol Raina, Prish Chakraborty, Aizhan Akhmetzhanova, Shu-Fan Chen, Gemma Zhang, Jackie Lodman and to some of the other high energy physics enthusiasts of the Boston physics area Katie Fraser, Aditya Parikh, Hengameh Bagherian, Wenzer Qin, Yitian Sun, Sarah Geller, and Josh Foster for the wonderful time spent discussing physics and especially for the time spent sharing meals and ventures together. To my incoming cohort of graduate students in general, and to Ian Davenport, Barak Gabai, Eric Knall, Sepehr Ebaudi, Cari Cesarotti, and Nicolo Foppiani in particular, thank you for the various misadventures we shared and for the post-graduation plans we have yet to look forward to! As valuable as the skills I have developed in my PhD are these friendships that I have forged. So it seems that some of the best friendships are forged in the hottest of fires - thank you also to Qianshu Lu for being a true friend and magnificent human throughout these years, sharing in many of the same turmoils and helping me find a path guilded with laughter and encouragement through the chaotic mists of our time at Harvard.

In ways that are difficult to see but impossible to deny, many others in my life have played critical roles in my arrival at this point. Within the department Lisa Cacciabaudo, Jacob Barandes, Hannah Belcher, Dionne Clark, and Frank Torres have built the tools I needed to be successful as a student, advocated for me when I could not myself and provided an environment within which we can work efficiently - beyond all of this, they have lent their time to listen when needing someone to talk to. Chief amongst these counselors in times of need has been Carol Davis - thank you for the many moments of laughter we shared, and for your compassion in times of hardship. To my college mentors Darien Wood, Lindley Winslow, Emanuela Barberis, Jonna Iacono, Toyoko Orimoto, Swastik Kar, Erin Cram, Bill Fowler, and Carey Rappaport, thank you for helping me discover many things I want to do in the world, and for getting me started on one. For supporting me through my early missteps, encouraging my early interest in research, and providing me resources to pursue that curiosity, thank you to my high school mentors Kate Pruett, Steve Coover, and Paul McElligott. Ensuring I remember, and partake in, the many joys of life outside of the world of physics I have to thank the many friends I have gained through life, and especially those who have seen me through this phase of my life: Alec, Luca, and Haley Luyendyk; Bogdan and Marley Korishev; Tyler Braun, Jake Vergara, Sarah MacClellan, Ian McLarney, Eric Yu, Dan Polnerow, Matt Curran, and Rachel Rubin.

A constant throughout all my endeavors in life, my family has played a monumental role in my achievements and aspirations. To Kerri and Brian Coveney, thank you for welcoming me into your family, for your patience at my lowest points, for providing a second place I can call home, and for inviting me to share in the bright moments of your lives. To my mother Nicole, my father Brian, my sister Gabrielle, and my brother Vincent, I have never felt anything but the highest support from each of you. You all have given me perspective when lost in my work and am beyond thankful for the joyful excursions you refused to leave me out on these past few years. Each of the communities of people I have given thanks for evokes particular aspects of who I am, exposing specific flaws and revealing particular strengths. There are only two souls that have seen the entirety of my being, every aspect of who I am in the rawest of forms. In my time at Harvard and beyond, Katlyn Coveney, my partner, has helped me to bear every burden I have faced; to find clarity in my path when I could not see through the fog of my own doubt; to recognize and celebrate the moments which make life beautiful before they slip away; to realize what is important from what is fleeting; and to summon the strength to tirelessly pursue those ideals. Through a sea of tribulation, she has guided me here. And the second soul? Our dog, Wrigley, and he has fulfilled these roles to no lesser degree.

DEDICATED TO KATLYN.

List of Figures

List of Tables

"Truly, Water is become now fairer than my heart imagined, neither had my secret thought conceived the snowflake, nor in all my music was contained the falling of the rain."

J.R.R. Tolkien

The following authors contributed to Chapter 1: Weishuang Linda Xu, Julian Muñoz and Cora Dvorkin.

The following authors contributed to Chapter 2: Weishuang Linda Xu, Julian Muñoz and Cora Dvorkin.

The following authors contributed to Chapter 3: Weishuang Linda Xu, Julian Muñoz and Cora Dvorkin.

The following authors contributed to Chapter 4: Zhong-Zhi Xianyu and Lisa Randall.

0

Introduction

For beyond a century, the thread of Dark Matter (DM) has been joined deeper into the weave of knowledge comprising our understanding of physics, astrophysics, and cosmology; yet, its mechanisms remain elusive to the best observational efforts and theoretical guidance of these disciplines. We have yet to confirm the non-gravitational interactions, if any, that DM has with the Standard Model of Particle Physics (SM). More generally, the particle nature of DM remains unknown. As with any difficult problem, it is best to seek insight from a variety of sources and approaches. And so we have: the pursuit of DM extends across dozens of orders of magnitude of energy; collider physics; planetary, solar, galactic, and extragalactic astronomy; the astrophysics of black holes, supernovae, and clusters; and the entire cosmological history of the Universe. Indeed, the enthrall of DM is evidenced by the scale of its presence in the modern resource prioritization of these communities [155,56]. A natural response to the persistence of the DM problem is to allow our theoretical postulations of its properties to become increasingly imaginative. Caution, however, must be taken. Unbounded theoretical creativity without diligent reduction to observational consequence is a recipe for failing to appreciate where we have already labored experimentally to exclude implausible physics. Ignorance puts at risk the time, financial, and human resources of the scientific community - potentially slowing our progress for decades more to come. Worse yet, it threatens a misdirection of our physical intuition which has sweeping consequences across all domains of physics. It is therefore imperative that the collage of increasingly kaleidoscopic DM theories be continuously corralled by incisive mappings onto the interconnected landscape of experimental constraint. The use of cosmological datasets to constrain, and expand knowledge of, broad swathes of phenomenological descriptions of the particle nature of DM shall be the overarching theme of this dissertation.

Cosmological data provide a powerful tool in the search for physics beyond the Standard Model. An interesting target are light relics, new degrees of freedom which decoupled from the SM while relativistic. Nearly massless relics contribute to the radiation energy budget, and are commonly parametrized as variations in the effective number $N_{\rm eff}$ of neutrino species. Additionally, relics with masses greater than 10^{-4} eV become non-relativistic before today, and thus behave as matter instead of radiation. This leaves an imprint in the clustering of the large-scale structure of the universe, as light relics have important streaming motions, mirroring the case of massive neutrinos. In one line of inquiry, we forecast how well current and upcoming cosmological surveys can probe light massive relics (LiMRs). We consider minimal extensions to the SM by both fermionic and bosonic relic degrees of freedom. The broad theoretical coverage attained by this study is a result of assuming only gravitational interactions between LiMRs and the rest of the SM - any interacting theory must at least satisfy the bounds set here. By combining current and upcoming cosmic-microwave-background and large-scale-structure surveys, we forecast the significance at which each LiMR, with different masses and temperatures, can be detected. We find that a very large coverage of parameter space will be attainable by upcoming experiments, opening the possibility of exploring uncharted territory for new physics beyond the SM.

The massive neutrino is a particle with studied and observed SM interactions which conforms to the definition of a LiMR. As in the case of the generic LiMR, a promising avenue to measure the total, and potentially individual, mass of neutrinos consists of leveraging cosmological datasets, such as the cosmic microwave background and surveys of the large-scale structure of the universe. In order to obtain unbiased estimates of the neutrino mass, however, many effects ought to be included. Here we forecast, via a Markov Chain Monte Carlo likelihood analysis, whether measurements by two galaxy surveys: DESI and *Euclid*, when added to the CMB-S4 experiment, are sensitive to two effects that can alter neutrino-mass measurements. The first is the slight difference in the suppression of matter fluctuations that each neutrino-mass hierarchy generates, at fixed total mass. The second is the growth-induced scale-dependent bias (GISDB) of haloes produced by massive neutrinos. We find that near-future surveys can distinguish hierarchies with the same total mass only at the 1σ level; thus, while these are poised to deliver a measurement of the sum of neutrino masses, they cannot significantly discern the mass of each individual neutrino in the foreseeable future. We further find that neglecting the GISDB induces up to a 1σ overestimation of the total neutrino mass, and we show how to absorb this effect via a redshift-dependent parametrization of the scale-independent bias.

Where our ability to cosmologically constrain LiMRs is driven by observational sensitivity to the energy scale introduced by their free streaming, we can generalize our study further by considering other classes of models which introduce characteristic energy scales into the cosmology. Ultra-light axions with masses $10^{-33} < m_\varphi/\mathrm{eV} < 10^{-22}$ are allowed to constitute only a small fraction of the observed dark matter abundance. Nevertheless, they may yet produce a visible impact on the cosmology due to their macroscopic quantum scale. Next generation galaxy survey data are poised to challenge this possibility, but in order to do so, all aspects of structure formation in this quasi-linear regime must be accounted for consistently and precisely. This includes modeling not only the effect of these axions on the background cosmology and matter fluctuations, but also on the halo bias that governs the tracers we observe, namely galaxies. In this work we discuss the effect of ultra-light axions on cosmological observables, and present a prescription for computing the growth-induced scale-dependent bias in their presence. We find that these axions introduce a step in the halo bias at their characteristic Jeans scale, representing – even at percent-level abundances – a sizable increase in the total scale-dependence of the bias, compared to the ΛCDM fiducial. We implement this prescription as a function of axion mass and relic abundance, in a public package which we dub `RelAxiFast`, an extension of the extant `RelicFast`.

The first observation of the gravitational wave signature of a binary black hole (BBH) merger has provided another avenue of inquiry into unknown physics. Since this observation, there have been further detections of BBH mergers of anomalous mass - casting doubt as to the astrophysical formation channels of black holes. One candidate for both DM and black hole formation is a theory for primordial black holes (PBH) formed early in the Universe. We consider how these different formation astrophysical and cosmological formation channels can lead to differences in the residual orbital eccentricity of BBH systems in the regimes of gravitation wave experiments. Contrasted with the study of individual black holes and binary systems, we consider how resolving characteristics of black hole populations offers insight into the mechanisms which may have formed the black holes in our Universe. This work considers how one population characteristic, the distribution in the orbital eccentricity of black hole binary pairs, might be studied within the experimental landscape of mHz-Hz frequency gravitational detectors. We expand on prior works which considered these effects at fixed eccentricity. As an example, we present how the eccentricity distributions corresponding to different formation channels produce dramatic shifts in the number of observed binaries in a mHz range detector like LISA. We also demonstrate how adding the capability to observe highly eccentric orbits offers a stark improvement in formation channel distinguishability due to the breaking of dataset degeneracies, offering motivation for the development of eccentric gravitational waveform templates.

In all the studies presented in this dissertation, we seek to infer limits on new physics through only the gravitational effects of those processes. In this sense, these approaches cover an extremely broad set of phenomenological parameter space, as any model with additional Standard Model interactions will need to at least satisfy the constraints derived here. Our ability to use gravitational interactions to derive competitive bounds with other observational channels is a recent phenomenon that is only enabled by the observational and sensitivity scales achieved by modern detectors. The promise of a fleet of upcoming galaxy, gravitational wave, and electromagnetic surveys in the approaching decade is a compelling reason to further refine the methods presented in this dissertation so that we can quickly narrow the field of DM candidates using information from the only channel through which DM is known to interact - gravity.

1

Finding eV-scale Light Relics with Cosmological Observables

1.1 INTRODUCTION

The nature of the dark sector is one of the major puzzles of fundamental physics, integral to the understanding of our universe across almost every epoch. Searches for the composition of the dark sector and, more broadly, of physics beyond the Standard Model (SM), take place at different energy scales, and use data ranging from particle colliders to astrophysical and cosmological surveys. The interactions of the dark sector with the SM are central to many of these searches. Yet, the small energies and interaction cross-sections expected in many models often result in low experimental sensitivity to new physics. In contrast, by exploring the entropic effects of new dark-sector physics, cosmological data is in an exciting position to make robust discoveries.

Numerous extensions of the SM happen to posit the existence of light, feebly interacting particles, including axions and axion-like particles [167,201,192,32], dark photons [13,46,31,84], and light fermions [61,97,86]. One broad category are *light relics*, stable particles which were in thermal contact with the SM in the early universe and decoupled while relativistic. Consequently, their cosmic abundance was frozen and survived until $z = 0$. The quintessential example within the SM are neutrinos, but they need not be the only light relics to populate our universe. Different proposed new light relics include a fourth, sterile neutrino, whose existence is suggested by different anomalous experimental results [92,148,19] (see Ref. [73] for a recent review); as well as the gravitino, the supersymmetric partner of the graviton [45].

New relics that are sufficiently light will manifest as dark radiation, and can be searched for through their effect on the cosmic microwave background (CMB) anisotropies [37,112,40], typically parametrized by the effective number of neutrino species, $N_{\rm eff}$ (which is 3.045 in the standard cosmological model [139,72,22]). Massive relics can, on the other hand, become non-relativistic at some point in cosmic history, and behave as other components of matter in the Universe thereafter. However, their decoupling while relativistic gives these relics significant streaming motion, which sets a scale below which they cannot cluster, thus altering the large-scale structure (LSS) of our universe. This has allowed cosmology to set the leading constraints on neutrino mass, at $\Sigma m_\nu < 0.26$ eV (95% C.L.), assuming standard cosmology [17]. In this work we will search for new Light—but Massive—Relics (LiMRs) using cosmological observables.

Cosmological data from near-future surveys are expected to provide exquisite measurements of the distribution of matter in our universe. LiMRs that have become non-relativistic before $z = 0$ (with masses $m_X \gtrsim 10^{-3}$ eV), will impact that distribution by behaving as hot dark matter[50,36,44,81,42]. In addition to the relic mass, two relevant parameters determine the relic abundance. The first is their number g_X of degrees of freedom. The second is their temperature $T_X^{(0)}$ today. Due to comoving-entropy conservation, any relic that was in equilibrium with the SM in the early universe ought to have $T_X^{(0)} \geq 0.91$ K. This minimum temperature gives rise to different values of $\Delta N_{\rm eff}$ for each type of relic[54]: 0.027 for scalars ($g_X = 1$), 0.047 for Weyl fermions ($g_X = 2$), 0.054 for massless gauge bosons ($g_X = 2$), and 0.095 for Dirac fermions ($g_X = 4$). In addition, relics with masses in the eV-scale will become non-relativistic before $z = 0$, leaving an imprint in the form of suppressed matter fluctuations. Here we forecast how well eV-scale LiMRs can be observed by joint CMB and LSS surveys.

This paper is structured as follows. In Section 1.2 we briefly review light relics and their effects on cosmological observables. In Section 1.3 we detail the datasets we consider, which we employ in Section 1.4 to forecast constraints on LiMRs within the mass range 10^{-2} eV- 10^{1} eV. We conclude in Section 1.5.

1.2 Light relics and their effect on cosmological observables

We begin with an overview of the physics of light relics and their effects on cosmological observables. A LiMR X is characterized by its present-day temperature $T_X^{(0)}$ and mass m_X, as well as its statistics, bosonic or fermionic, and its number g_X of degrees of freedom. The present-day temperature of a light relic (massive or not) is set by the time at which it decouples from the SM thermal bath, which is found as

$$T_X^{(0)} = \left(\frac{g_{*S}^{(0)}}{g_{*S}^{(\rm dec)}} \right)^{1/3} T_\gamma^{(0)}, \tag{1.1}$$

where $g_{*S}^{(0/\mathrm{dec})}$ denotes the entropy degrees of freedom in the universe today/when the relic decoupled, and $T_\gamma^{(0)} = 2.725$ K is the present-day temperature of the photon bath. In this way, the conservation of comoving entropy provides a minimal light relic temperature assuming the SM with no additional degrees of freedom (other than the relic),

$$T_X^{(0)} \gtrsim \left(\frac{3.91}{106.75}\right)^{1/3} T_\gamma^{(0)} \approx 0.91 \text{ K}, \tag{1.2}$$

where just after the electroweak phase transition we have $g_{*s}^{(\mathrm{dec})} = 106.75$ encompasses all the known degrees of freedom of the Standard Model, and the present-day value of $g_{*s}^{(0)} = 3.91$ includes photons and decoupled, cooler neutrinos. As an example, the SM (active) neutrinos have $T_\nu^{(0)} = 1.95$ K, as they decoupled just prior to electron-positron annihilation where $g_{*s}^{(\mathrm{dec},\nu)} = 10.75$. Note that the baryonic and cold-dark matter (CDM) contributions are negligible, given their exponentially suppressed abundance.

In contrast, light relics decoupled while relativistic, and so are cosmologically abundant, with number densities comparable to that of photons or neutrinos. For instance, a Weyl fermion decoupling as early as possible (with minimal present-day temperature 0.91 K) will have a number density today of 11 cm^{-3}, and a vector boson that decouples just before $\mathrm{e}^+\mathrm{e}^-$ annihilation (with a temperature today of 1.95 K, as neutrinos) will have a present-day number density of 150 cm^{-3}. Thus, the contribution of light relics to the cosmic energy budget can be significant.

It is often enlightening to describe the cosmological effects of other relics in relation to those of neutrinos, given their common origin as light relics. As advanced in the introduction, relics in the early universe (while $T_X \gg m_X$) behave as radiation, and their cosmological impact while relativistic can be encapsulated in the number of effective neutrinos, N_{eff}, defined with respect to their contribution to the radiation energy density,

$$\begin{aligned}\rho_{\rm rad}(z) &= \frac{\pi^2}{30}\left(\sum_{\rm bosons} g_b T_b^4(z) + \frac{7}{8}\sum_{\rm fermions} g_f T_f^4(z)\right) \\ &\equiv \frac{\pi^2}{30}\left(2T_\gamma^4(z) + \frac{7}{4}N_{\rm eff}T_\nu^4(z)\right), \end{aligned} \tag{1.3}$$

where $T_{\gamma/\nu}(z)$ is the temperature of photons and neutrinos at redshift z, g_b/g_f are the degrees of freedom, and T_b/T_f are the temperatures of each boson/fermion, respectively.

Introducing an entropically significant light relic will generate a contribution to Eq. (1.3) of $(\pi^2/30)g_X T_X^4$ for bosonic species, or 7/8 times that for fermionic species. We can then describe any departure from the predicted value of $N_{\rm eff}^{\Lambda{\rm CDM}} = 3.045$ in the standard ΛCDM model by the quantity $\Delta N_{\rm eff}$, given by

$$\Delta N_{\rm eff} = c_1^\gamma \left(\frac{g_X}{g_\nu}\right)\left(\frac{T_X^{(0)}}{T_\nu^{(0)}}\right)^4, \tag{1.4}$$

in terms of the neutrino parameters $g_\nu = 2$ and $T_\nu^{(0)} = 1.95$ K. The factor $c_1 = 8/7$ accounts for the difference between the Bose-Einstein ($\gamma = 1$) and Fermi-Dirac ($\gamma = 0$) distributions.

This discussion is encapsulated in Fig. 1.1, showing the relation between the present-day relic temperature to the time of relic decoupling, and its corresponding contribution to $N_{\rm eff}$. Note that the present-day temperature of a relic for fixed decoupling epoch does not depend on particle species, but its contribution to radiation energy does.

Current limits on $\Delta N_{\rm eff}$ arise primarily from observables at two epochs. The first is recombination. Measurements of radiation at recombination are sensitive to relics lighter than $\sim$ 0.1 eV. The *Planck* 2018 analysis reports a measurement of $N_{\rm eff} = 2.99^{+0.34}_{-0.33}$ (TT+TE+EE+lowE+lensing+BAO) at 95% C.L.[17]. The proposed CMB-Stage 4 (CMB-S4) experiment is expected to refine this measurement to the $\sigma(N_{\rm eff}) = 0.03$ level[2]. The second is the Helium abundance, from where we can infer the number of relativistic species present during big bang nucleosynthesis (BBN). The 68%

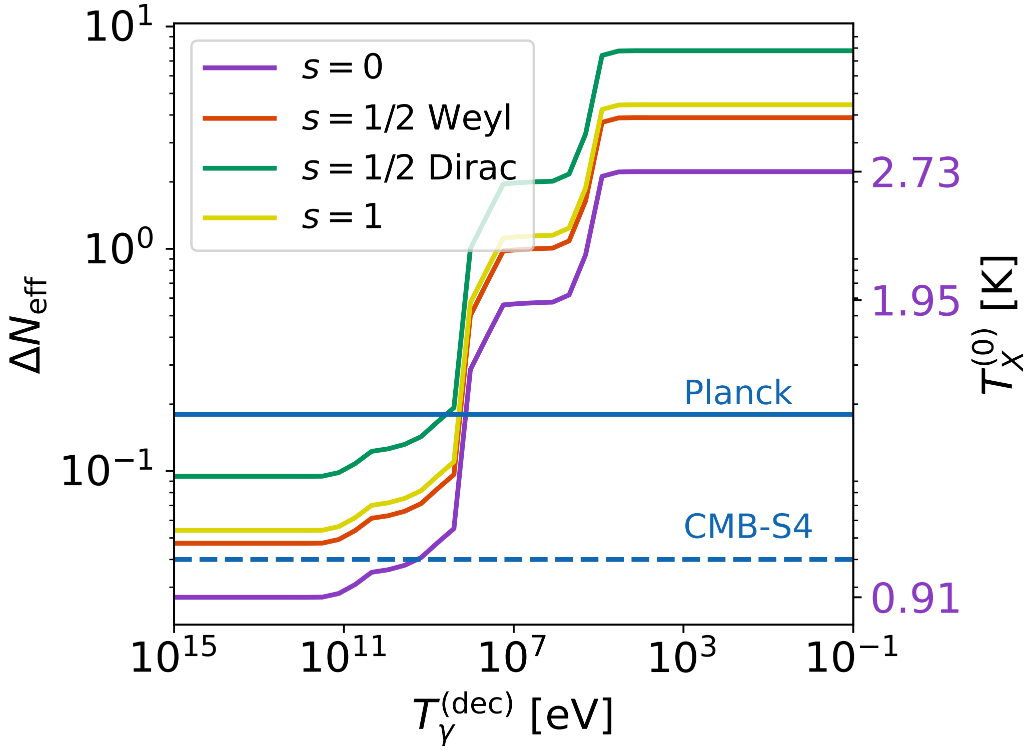

Figure 1.1: Cosmic evolution of $\Delta N_{\rm eff}$ due to a light relic that decoupled when the universe had a temperature $T_\gamma^{\rm (dec)}$. We assume four different types of relics with spin s, as described in the text, and show the 68% C.L. constraints achieved by *Planck* as a horizontal solid line, and the forecast by CMB-S4 in dashed lines. The right vertical axis shows what the temperature of the relic would be at $z = 0$, following the violet (lowest) curve plotted for $s = 0$. Note that these constraints only apply to relics with $m_X \approx 0.1$ eV or lighter.

C.L. measurement during that era is $N_{\rm eff} = 2.85 \pm 0.28$[69], which is valid for all relics lighter than $m_X \lesssim 10^6$ eV. Note that this does not affect dark matter (DM) produced via the freeze-in mechanism, as it can contribute negligibly to $N_{\rm eff}$[60,83].

In this work we consider detection prospects for four types of LiMRs: scalars, vectors, and both Dirac and Weyl fermions. We study relics with eV-scale masses, 10^{-2} eV $\leq m_X \leq 10^1$ eV, such that they all behave as matter at $z = 0$, with the highest mass candidates constituting up to $\sim$10% of DM abundance. Finally, we also consider a range of temperatures, bounded by $T_X^{(0)} \geq 0.91$ K from below. Our maximum temperature is informed by the constraint $\Delta N_{\rm eff} \leq 0.36$ from *Planck*, corresponding to a single additional species of Weyl fermion at $T_X^{(0)} \leq 1.5$ K. This bound could be further improved by combining with BBN measurements of e.g. D/H ratios[87], Lyman-α forest flux power spectrum data[184,164], as well as Baryon Acoustic Oscillations (BAO) and galaxy power spectrum measurements[23,207,41].

Effect on the LSS of the Universe

LiMRs can become non-relativistic at some point in cosmic history, and comprise a fraction of DM at $z = 0$. Unlike CDM, which is expected to compose the majority of the matter sector, LiMRs have significant thermal motions, even if non-relativistic. Thus, these relics will stream away from structures below their free-streaming scale, which during matter domination is given by[25,48]

$$k_{\rm fs} = \frac{0.08}{\sqrt{1+z}} \left(\frac{m_X}{0.1 \rm eV}\right) \left(\frac{T_X^{(0)}}{T_\nu^{(0)}}\right)^{-1} h \, \rm Mpc^{-1}. \tag{1.5}$$

Throughout this section we assume a Weyl fermionic relic, and we will relax this assumption later. This presents another way of searching for LiMRs: through their effect on the matter fluctuations. LiMRs produce a suppression in the matter power spectrum at scales smaller than $k_{\rm fs}$, which we discuss below. The size of this suppression depends on the present abundance of the LiMR, which (if non-relativistic) is given by

$$\Omega_X h^2 = \frac{m_X}{93.14 \, \rm eV} \frac{g_X}{g_\nu} \left(\frac{T_X^{(0)}}{T_\nu^{(0)}}\right)^3 . \tag{1.6}$$

From Eq. (1.6) we see that there is a maximum allowed particle mass, found by saturating the observed DM abundance $\Omega_{\rm cdm} h^2 = 0.12$[17]. For a relic temperature $T_X^{(0)} {\approx} 1.5$ K, this is $m_X {\approx} 10$ eV. Additionally, in this work we are interested in the relics that become non-relativistic before today.

Thus, the mass range we will study encompasses

$$10^{-2}\,\text{eV} \leq m_X \leq 10^{1}\,\text{eV}. \tag{1.7}$$

LiMRs produce a suppression in matter fluctuations, similar to neutrinos, due to two reasons. The first is simply that the light relic does not cluster at small scales, and its fluctuation δ_X at small-scale roughly follows $\delta_X = (k/k_{\rm fs})^{-2}\,\delta_m$ with respect to the matter overdensity δ_m. The second is that the absence of relic fluctuations at small scales slows down the growth of CDM (and baryon) overdensities. Together, these two factors produce a suppression of roughly $(1-14f_X)$ in the matter power spectrum[129], where f_X is the fraction of matter that is composed of the LiMR X. This suppression is less pronounced for relics that stay relativistic for longer, which yields the well-known result of $(1-8f_\nu)$ for neutrinos comprising a fraction f_ν of matter, as neutrinos only become non-relativistic during matter domination. These numbers are for illustration purposes only, and in all cases we find the full effect of LiMRs on the cosmological observables using the publicly available software `CLASS`[49]. Nevertheless, they provide intuition about the physical effect of such a relic. While the mechanism that produces the suppression is the same as for neutrino masses, the free-streaming scale $k_{\rm fs}$ for a LiMR is not fully determined by its mass (or abundance), as their temperature today is unknown. Relics that are still relativistic at $z=0$ (with $m_{\rm X} \lesssim 10^{-3}$ eV) will have never collapsed into structures and thus their observable effects can be fully included into $\Delta N_{\rm eff}$. In practice, this is the case for LiMRs with masses below ~ 0.1 eV, as we will show, so we will use our results for a 10^{-2} eV relic for lighter masses.

To study LiMRs, the relevant observables are the fluctuations of baryons and cold dark matter, as only those will gravitationally bind to form the visible structures we observe as galaxies, the relics being too light to cluster (see, however, Ref.[136]). The power spectrum of baryonic plus cold dark-matter fluctuations is modeled by

$$P_{\rm cb}(k) = P_\zeta(k)\Big(f_b T_b(k) + f_c T_c(k)\Big)^2, \tag{1.8}$$

where P_ζ is the primordial power spectrum, the transfer functions T_b and T_c are found using `CLASS`[49], and the fractional abundances are defined by

$$f_{b/c} \equiv \frac{\omega_{b/c}}{\omega_b + \omega_c}, \tag{1.9}$$

where ω_b and ω_c are the baryon and CDM abundances.

We show the suppression in $P_{\rm cb}$ in Fig. 1.2 (upper panel) for a fermion with $m_X = 0.02$ eV and $T_X = 0.91$ K, for degrees of freedom $g_X = 2$, 3 and 4. In all cases the high-k power is more suppressed, as expected. Increasing the abundance of the LiMR, by augmenting g_X, produces a more marked suppression, while keeping the shape fixed. Moreover, increasing the relic abundance produces wiggles at the BAO scale, as the LiMR both contributes as radiation at recombination and free streams – like neutrinos – changing the BAO phase[37].

The suppression of matter fluctuations produces a change in the biasing of galaxies, which has been calculated for both neutrinos and other relics[133,152,63], and accounted for in neutrino-mass forecasts in our companion paper[205]. This produces a growth in the galaxy power spectrum that partially compensates the relic-induced suppression. Here we account for this growth induced scale-dependent bias (GISDB) by multiplying the Lagrangian bias by a k-dependent factor

$$g(k) = R_L^{\Lambda{\rm CDM}}(k) R_L^X(k) R_L^\nu(k), \tag{1.10}$$

where the functions R_L^i account for different effects, following Ref.[152]. First, $R_L^{\Lambda{\rm CDM}}$ accounts for the step-like change in the growth rate of fluctuations before and after matter-radiation equality, parametrized as

$$R_L^{\Lambda\mathrm{CDM}}(k) = 1 + \Delta_{\Lambda\mathrm{CDM}} \tanh\left(\frac{\alpha k}{k_{\mathrm{eq}}}\right), \tag{1.11}$$

where $\Delta_{\Lambda\mathrm{CDM}} = 4.8 \times 10^{-3}$ and $\alpha = 4$ determine the amplitude and location of the step, given the scale k_{eq} of matter-radiation equality. The two other factors account for the effect of a LiMR on the matter power spectrum, also taken to be a step-like function

$$R_L^i(k) = 1 + \Delta_i \tanh\left(1 + \frac{\ln q_i(k)}{\Delta q}\right), \tag{1.12}$$

with an amplitude $\Delta_i = 0.6 f_i$ determined by the fraction f_i of matter composed of the relic i (X or ν), width $\Delta q = 1.6$, and where we have defined $q_i(k) \equiv 5k/k_{\mathrm{fs,i}}$, given the free-streaming scale $k_{\mathrm{fs},i}$ of each LiMR.

Effect on the CMB

The CMB is sensitive to the presence of LiMRs in the universe, through their mean energy density[168,78] and their perturbations[38,112]. Their additional energy density changes the expansion rate of the universe, which in turn affects the CMB damping tail. Since matter-radiation equality is very well measured through the location of the first acoustic peak, this causes the power spectrum to be suppressed on short-wavelength modes. In addition to this effect, their perturbations cause a change in the amplitude and a shift in the location of the CMB acoustic peaks (for a review of the phase shift in the acoustic peaks in the CMB, see Ref.[40]).

We show an example of the effect of a LiMR on the CMB in Fig. 1.2 – again for a fermion with $m_X = 0.02$ eV and $T_X = 0.91$ K, for degrees of freedom $g_X = 2$, 3 and 4. The amplitude and phase shift of the BAO is clearly seen to increase with g_{X}.

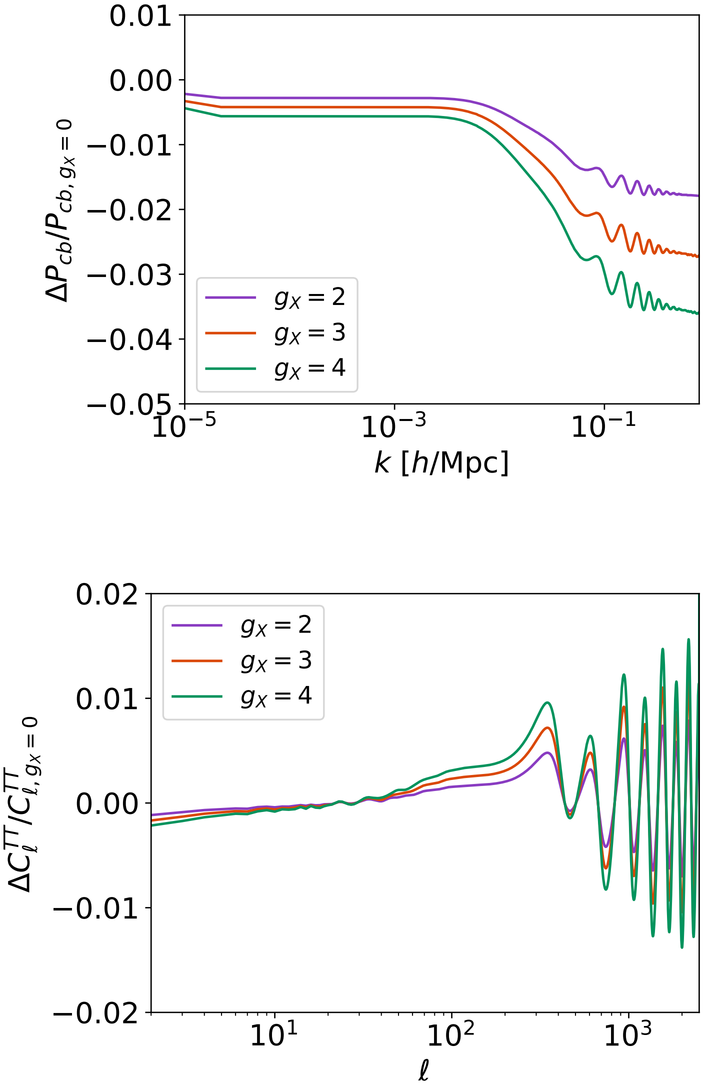

Figure 1.2: Effect of introducing a fermion with degrees of freedom g_X, temperature $T_X = 0.91$ K and mass $m_X = 0.02$ eV on the CDM+baryon power spectrum (upper panel) and the CMB temperature power spectrum (lower panel). Here all cosmological parameters are fixed when introducing the LiMR so the fraction of the matter or radiation energy occupied by the LiMR before and after its non-relativistic transition will increase with its abundance. Since the LiMR energy density is not counted in the CDM plus baryon power power spectrum, an increase in LiMR abundance will manifest as an overall suppression to P_{cb}. We note that an effective fractional number of degrees of freedom may be achieved as a result of out-of-equilibrium processes.

Types of Relic

Throughout this work we will study four major types of LiMRs, two fermionic and two bosonic, which we now describe.

In the fermionic category, the first type we study are the neutrino-like Weyl fermions, with non-zero mass, spin $s = 1/2$, and two degrees of freedom ($g_X = 2$). In addition to sterile neutrinos, an intriguing example is the gravitino, the supersymmetric partner of the graviton. While the gravitino has $s = 3/2$, only the longitudinal modes couple to the Standard Model and hence behaves equivalently to an $s = 1/2$ particle with $g_X = 2$. The gravitino is predicted in models of supersymmetric gravity to have a mass in the eV range [161,144], within the range relevant to our study. The second type we tackle are the related Dirac fermions, such as the axino [66], which simply have twice as many degrees of freedom ($g_X = 4$).

In the bosonic category we study two types of particles as well: firstly scalars, with only one degree of freedom ($g_X = 1$). A realization of this model could be a Goldstone boson, which can have naturally small masses. The second type are spin-1 vectors. We assume that they have a Stueckelberg mass, as it is technically natural [180] and avoids complications from Higgs mechanisms. While this relic will be non-relativistic today, its longitudinal mode was decoupled in the early universe (while it was relativistic), and thus only two of the three degrees of freedom were populated. Therefore, this relic has $g_X = 2$.

Instead of modifying the distribution function for each type of relic, we will take advantage of the fact that any relic, whether bosonic or fermionic, can be recast onto an equivalent Weyl relic (i.e., a neutrino with $g_W = 2$), with some temperature T_W^{eq} and mass m_W^{eq} [152,50]. Justification for this procedure is based on the results of other works which considered the significance of the distribution shapes for different species [152]. Assuming a relic of temperature T_X, with g_X degrees of freedom, the equivalent Weyl relic has

$$T_W^{\mathrm{eq}} = T_X \, (g_X/g_W)^{1/4} \, c_1^{\gamma/4} \tag{1.13}$$

$$m_W^{\mathrm{eq}} = m_X \, (g_X/g_W)^{1/4} \, c_1^{\gamma/4} c_2^{\gamma}, \tag{1.14}$$

where we correct for the different distributions of these particles by setting $\gamma = 1$ for bosons (and $\gamma = 0$ for our base case of fermions as before), with constants $c_1 = 8/7$ (as in Eq. 1.4) and $c_2 = 7/6$. Note that our normalization is slightly different from that found in Ref.[152], as there fermionic degrees of freedom contributed by 3/2.

1.3 METHODS

We now present our forecasting methods. In this first exploratory work we will follow a Fisher-matrix approach, in order to efficiently explore the 2D parameter space ($T_X^{(0)}$, m_X) of possible LiMRs. We encourage the reader to visit Appendix 1.6 for a comparison against MCMC results. We will also cover different combinations of datasets. For the CMB, we will study the current *Planck* satellite[17] as well as the upcoming ground-based CMB-S4[1]. On the galaxy-survey side we will consider the current BOSS[70], the ongoing DESI[16], and the upcoming *Euclid*[28] surveys.

1.3.1 PARAMETERS

We are interested in forecasting how well different LiMRs with varied temperatures and masses can be detected. Therefore, a simple Fisher forecast of the relic mass and temperature, assuming a particular fiducial relic, is insufficient. Instead, we will find how well LiMRs of varying mass m_X and temperature $T_X^{(0)}$ can be observed by different experiments. The parameter we will forecast is g_X, the number of degrees of freedom of the LiMR.[*] Then, $g_X/\sigma(g_X)$ is a good proxy for the significance at which a LiMR of a particular m_X and $T_X^{(0)}$ can be detected.

In order to properly search for a LiMR we have to marginalize over the six ΛCDM parameters. These include the baryon and cold dark-matter abundances, ω_b and $\omega_{\rm cdm}$ (with fiducial values of $\omega_b = 0.02226$ and $\omega_{\rm cdm} = 0.1127$), the (reduced) Hubble constant $h = 0.701$, and the optical depth $\tau_{\rm reio} = 0.0598$ to reionization. The last two parameters are the amplitude A_s, and tilt n_s, of primordial fluctuations, with fiducial values of $A_s = 2.2321 \times 10^{-9}$ and $n_s = 0.967$. In addition, we marginalize over the effect of neutrino masses. We assume for our fiducial model the existence of three degenerate massive neutrinos, with $\sum m_\nu = 0.06$ eV, and we will report constraints both with and without marginalization over neutrino masses. Unless explicitly stated, no prior will be assumed for these parameters in the Fisher forecasts used to provide parameter constraints. For a discussion about the effect of the neutrino hierarchy see Refs. [30,205].

1.3.2 CMB EXPERIMENTS

We will model both *Planck* and CMB-S4 as having a single effective observing frequency, to avoid marginalizing over foregrounds. For *Planck* we will use CMB temperature (T) and E-mode polarization data, covering the range $\ell = [2 - 2500]$. We take noises of $\Delta_T = 43\mu$K-arcmin and $\Delta_E = 81\mu$K-arcmin, with a $\theta_{\rm FWHM} = 5$ arcmin angular resolution. This well approximates the (more complex) *Planck* data likelihood.

For CMB-S4 we take $\Delta_T = 1\mu$K-arcmin, and $\Delta_E = \sqrt{2}\Delta_T$, with an angular resolution of $\theta_{\rm FWHM} = 3$ arcmin. Additionally, we include lensing data, where we perform iterative delensing of B-modes to lower the noise, as in Refs. [106,160]. All modes cover the range $\ell = [30 - 5000]$, except for the TT autocorrelation, where we do not go beyond $\ell = 3000$ to avoid foreground contamina-

---

[*] We note that, while g_X appears to be a fixed quantity for a given relic, e.g. $g_X = 1$ for a scalar, changing g_X simply means altering the amount of relic particles (as both $\Delta N_{\rm eff} \propto g_X$ and $\Omega_X \propto g_X$) while keeping their thermal properties identical. That makes g_X a useful variable to forecast.

tion[2]. We add a Gaussian prior on the optical depth of reionization of $\sigma(\tau_{\rm reio}) = 0.01$, instead of the $\ell < 30$ modes in this case. This follows the prescription in the CMB-S4 Science Book[2], as well as our companion paper[205] and is the sensitivity reported from the *Planck* 2018 results. As such, it serves as a conservative estimate for futuristic surveys, such as CMB-S4.[17]

The CMB data will perform two main roles. First, it will very precisely measure the standard cosmological parameters, breaking many degeneracies in the LSS data. Second, the CMB is sensitive to the effects of a LiMR both during recombination and in the matter fluctuations at lower redshifts, through the weak lensing information.

1.3.3 GALAXY SURVEYS

For the LSS data we will consider three surveys, all of them spectroscopic. We leave for future work studying the promise of photometric surveys, such as the Vera Rubin Observatory[117], and weak-lensing surveys, such as the Dark Energy Survey[9].

We take the luminous red galaxy (LRG) sample of the Sloan Digital Sky Survey Baryon Oscillation Spectroscopic Survey (BOSS)[70], which will serve as an indication of the power of current data. To showcase the promise of upcoming surveys we study the emission-line galaxy (ELG) sample of the Dark Energy Spectroscopic Instrument (DESI)[16], and the more futuristic Hα-emitters of *Euclid*[28]. We restrict our analysis to a single tracer, the most populous for each survey, though more optimistic results are expected for multi-tracer approaches[51]. The noise per redshift bin for each sample is reported in Table 1.1. We assume sky coverages of 10,000 deg^2 for BOSS; 14,000 deg^2 for DESI; and 15,000 deg^2 for *Euclid*.

As each of these surveys contain distinct tracers, the bias description of each will be somewhat different as well. Here we follow a simple approach, and parametrize the linear Eulerian bias as

$$b_1(k,z) = \left[1 + b_L(k,z) + \alpha_{k2} k^2\right] , \tag{1.15}$$

where the α_{k2} term (with a fiducial value of 1 Mpc^2) accounts for non-linearities in the bias[149]. We emphasize that we do not include the clustering of light relics in this description. We also note that while cold dark matter and baryons may demonstrate different clustering behaviors at small scales, we do not consider such scales in this work and so do not include corrections to the bias that would differentiate the baryon and cold dark matter clustering fields. An additional scale-dependence comes from the aforementioned GISDB effect, which enters in the Lagrangian bias,

$$b_L(k, z) = [b_0(z) - 1]\, g(k), \tag{1.16}$$

where $g(k)$ is as defined in Eq. (1.10). The redshift evolution of the bias is encapsulated in the term $b_0(z)$, which is chosen such that the scale-independent (i.e., $k \to 0$) behavior of the Eulerian bias matches with suggestions made elsewhere in the literature[152]. For the ELGs in DESI we match to

$$b_0(z) = \frac{\beta_0}{D(z)}, \tag{1.17}$$

where $D(z)$ is the growth factor and $\beta_0 = 1$[16]; whereas for the tracers in BOSS and *Euclid* we take

$$b_0(z) = \beta_0 (1 + z)^{0.5\beta_1}, \tag{1.18}$$

with fiducials $\beta_0 = 1.7$ and $\beta_1 = 1$ as in Ref.[191]. We marginalize over the nuisance parameters β_0, α_{k2}, as well as β_1 for BOSS and *Euclid*. We note that a full analysis of the data might require marginalization over the amplitude of the bias at each redshift bin independently, which would however lead to a loss in constraining power.

z	0.05	0.15	0.25	0.35	0.45	0.55	0.65	0.75	0.85	0.95	1.05	1.15	1.25	1.35	1.45	1.55	1.65	1.75	1.85	1.95
$\frac{\mathrm{d}N_{LRG}}{\mathrm{d}z\,\mathrm{ddeg}^2}$ [BOSS]	8	50	125	222	332	447	208	30	0	0	0	0	0	0	0	0	0	0	0	0
$\frac{\mathrm{d}N_{ELG}}{\mathrm{d}z\,\mathrm{ddeg}^2}$ [DESI]	0	0	0	0	0	0	309	2269	1923	2094	1441	1353	1337	523	466	329	126	0	0	0
$\frac{dN_{H\alpha}}{\mathrm{d}z\,\mathrm{ddeg}^2}$ [*Euclid*]	0	0	0	0	0	0	2434	4364	4728	4825	4728	4507	4269	3720	3104	2308	1514	1474	893	497

Table 1.1: Forecasted number of target galaxies measurable by each survey: LRGs for BOSS, ELGs for DESI, and Hα emitters for *Euclid* per redshift per deg^2 at each redshift bin z, taken from Refs. [16,28,89].

1.3.4 FISHER MATRIX

We will obtain forecasted constraints using the Fisher-matrix formalism [122,206,194]. For the CMB we follow the approach of Refs. [154,95]. For the galaxy observables we detail below how we construct our Fisher matrix.

As described in Section 1.2, LiMRs suppress the clustering of matter in our universe, and as a consequence, that of biased tracers of matter, such as galaxies. We take into account several effects to convert from matter to galaxy fluctuations. First, there are redshift-space distortions (RSD), induced by the gravitational infall into, and peculiar velocities of galaxies [120,55]. We write the galaxy power spectrum as

$$P_g(k,\mu) = \mathcal{R}(k,\mu)\mathcal{F}(k,\mu)P_{\mathrm{cb}}(k), \tag{1.19}$$

in terms of the power spectrum $P_{\rm cb}(k)$ of CDM + baryon fluctuations, where the two pre-factors $\mathcal{R}$ and $\mathcal{F}$ account for the RSD and the finger-of-god (FoG) effect, both of which make P_g anisotropic, as they depend on $\mu = \hat{k} \cdot \hat{n}$, the line-of-sight angle.

We model the linear RSD term simply as

$$\mathcal{R}(k,\mu) = \left[b_1(k) + f\mu^2\right]^2 , \tag{1.20}$$

where b_1 is the linear Eulerian bias, as described above, and $f \equiv d \ln D / d \ln a$ is the logarithmic derivative of the growth factor D, which can be well approximated by[131]

$$f(z) = \left(\frac{\Omega_{\rm cb}(1+z)^3}{\Omega_{\rm cb}\,(1+z)^3 + \Omega_\Lambda} \right)^\gamma , \tag{1.21}$$

with $\gamma = 0.55$. The non-linear FoG effect is included in the term

$$\mathcal{F}(k,\mu) = \exp\left[-k^2\mu^2\sigma_v^2/H^2\right] , \tag{1.22}$$

with $\sigma_v = (1+z)\sqrt{c^2\,\sigma_z^2 + \sigma_{\rm FoG}^2/2}$, where $\sigma_{\rm FoG} = \sigma_{\rm FoG}^{(0)}\sqrt{1+z}$, with $\sigma_{\rm FoG}^{(0)} \equiv 250\,{\rm km\,s^{-1}}$ [208] as the intrinsic velocity dispersion of galaxies, and we take a spectroscopic redshift error $\sigma_z \equiv 0.001c$ [16], which corresponds to the DESI precision requirement at $z = 1$.

In addition, we include the Alcock-Paczynski (AP) effect[24,33,128], which accounts for changes in the observed k and μ and the comoving volumes from assuming different cosmologies. For that, we write the observed galaxy power spectrum as[150]

$$\tilde{P}_g(k',\mu') = P_g(k,\mu) \left(\frac{H_{\rm true}}{H_{\rm fid}}\right) \left(\frac{D_{A,\rm fid}}{D_{A,\rm true}}\right)^2 , \tag{1.23}$$

where the subscript “fid” refers to fiducial, and the “true” wavenumber k' and angle μ' are given by

$$k' = k\left[(1-\mu^2)\frac{D^2_{A,\mathrm{fid}}(z)}{D^2_{A,\mathrm{true}}(z)} + \mu^2 \frac{H^2_{\mathrm{true}}(z)}{H^2_{\mathrm{fid}}(z)}\right]^{1/2} \tag{1.24}$$

$$\mu' = \mu \frac{k}{k'} \frac{H_{\mathrm{true}}(z)}{H_{\mathrm{fid}}(z)}. \tag{1.25}$$

Properly accounting for the AP effect, thus, implies evaluating the entire galaxy power spectrum at different wavenumbers for each cosmological-parameter change. That can be computationally consuming, so instead we will perform a simpler step that is accurate to first order in derivatives (as any further is not captured by Fisher). Therefore, we can write

$$\frac{\partial \tilde{P}_g(k', \mu')}{\partial \theta_i} = \frac{\partial P_g(k,\mu)}{\partial \theta_i} + \mathcal{C}_i(k), \tag{1.26}$$

for each parameter θ_i, where

$$\mathcal{C}_i(k) = \frac{\partial P_g}{\partial k}\frac{dk}{d\theta_i} + \frac{\partial P_g}{\partial \mu}\frac{d\mu}{d\theta_i}, \tag{1.27}$$

accounts for the AP correction to linear order, with the derivatives of k and μ computed from Eq. (1.25).

The Fisher element for parameters θ_i, θ_j is then calculated as[89]

$$F_{ij} = \sum_z \int k^2 \mathrm{d}k \int \mathrm{d}\mu \frac{V(z)}{2(2\pi)^2} \left(\frac{\overline{n}\tilde{P}_g}{\overline{n}\tilde{P}_g + 1}\right)^2 \left(\frac{\partial \log \tilde{P}_g}{\partial \theta_i}\right)\left(\frac{\partial \log \tilde{P}_g}{\partial \theta_j}\right), \tag{1.28}$$

where $V(z)$ is the comoving volume for each redshift bin summed over, and $\overline{n}(z)$ is the comoving number density of tracers, given by $\overline{n}(z) = \Delta z f_{\mathrm{sky}} V^{-1}(z)\, dN/(dz\, d\mathrm{deg}^2)$, where the last factor is reported for each survey in Table 1.1. The integral over μ goes from -1 to 1, and over wavenumbers from $k_{\rm min} = \pi V(z)^{-1/3}$ to $k_{\rm max} = 0.2h\,{\rm Mpc}^{-1}$, which is mildly in the nonlinear regime[64]. While at higher z the fluctuations are smaller and, thus, we could reach higher $k_{\rm max}$ while linear, the biasing of galaxies becomes more complex, so we fix $k_{\rm max}$ for all z. We expect that non-gaussianities in the likelihood will affect constraints on cosmological parameters, but we do not model those effects in this work[102].

1.4 Results

In this section we discuss our cosmological constraints for a LiMR. We will perform two parallel analyses. First, we will show the reach of different combinations of datasets by forecasting $\sigma(g_X)$ for a Weyl (neutrino-like) relic of different masses and temperatures, covering the entire range of interest. Then, we will focus on the minimal case (that with $T_X^{(0)} = 0.91$ K) for the four relic types we consider, and find more precisely above which mass m_X they can be ruled out.

1.4.1 Full Parameter Space

We will start with a Weyl relic, and cover a broad range of cases, where in each case we will assume that there exists a LiMR in our universe with mass m_X and temperature today $T_X^{(0)}$, and forecast how well g_X can be measured as a measure of how significant a detection would be.

We scan through a range of LiMR masses m_X from 10^{-2} eV, as all lighter relics behave identically, up to $\sim$10 eV, where the relic abundance overcomes that of all DM. As for their temperature, we cover $T_X^{(0)} = [0.91 - 1.50]$ K, where the lower limit is as found in Section 1.2, and the upper limit saturates the current 95% C.L. *Planck* + BOSS DR12 BAO limit on $N_{\rm eff}$[17].

First, as a test, we forecast the errors on $N_{\rm eff}$ by looking at our lightest relic ($m_X = 0.01$ eV) as a proxy of the massless case, and translating the forecasted error $\sigma(g_X)$ in the degrees of freedom into

$$\sigma(N_{\text{eff}}) = \frac{\sigma(g_X)}{g_\nu} \left(\frac{T_X^{(0)}}{T_\nu^{(0)}} \right)^4 . \tag{1.29}$$

For reference, we have confirmed that assuming lower values of m_X result in the same forecasts for N_{eff}. This result is largely independent of the chosen $T_X^{(0)}$, so we will show forecasts for a Weyl fermion with $T_X^{(0)} = 0.91$ K.

Beginning with the CMB, the *Planck*-only forecast gives $\sigma(g_X) = 8.11$ corresponding to $\sigma(N_{\text{eff}}) = 0.19$ which is in agreement with the *Planck* value of $\sigma(N_{\text{eff}})$ in non-photon radiation density when allowing extra relativistic degrees of freedom Ref.[17]. Likewise, the CMB-S4-only forecast yields $\sigma(N_{\text{eff}}) = 0.040$. This is to be compared with the value of $\sigma(N_{\text{eff}}) = 0.035$ reported in Ref.[2] for the same combination of resolution and sensitivity. The $\sim 10\%$ difference is due to the delensing of T and E modes[40,126] that is performed in Ref.[2] but not in our forecasts. This is because we are chiefly interested in more massive relics, for which the phase shift is not the main cosmological signature.

In both cases, as well as the ones below, we account for a noted degeneracy with Σm_ν by marginalizing over the neutrino mass in our forecasts. Adding LSS data only improves these results, as we show in Table 1.2. In particular, we find that adding BOSS to *Planck* gives $\sigma(N_{\text{eff}}) = 0.14$; substituting DESI for BOSS yields $\sigma(N_{\text{eff}}) = 0.06$. Looking to the future, Euclid and CMB-S4 will lower this constraint to $\sigma(N_{\text{eff}}) = 0.02$.

We now move to non-zero masses, and provide marginalized posteriors from forecasts for a 0.91 K (minimum temperature) Weyl relic at different masses in Fig. 1.3. We only show the 2D contours between g_X and other cosmological parameters; for the full triangle plots at fixed mass $m_X = 0.01$ eV, see Appendix 1.8. The combination of information from the CMB and LSS can be seen to significantly improve constraints by breaking parameter degeneracies present in the individual datasets.

$\sigma(N_{\rm eff})$	CMB Only	BOSS	DESI	*Euclid*
LSS Only		0.92 (0.84)	0.29 (0.25)	0.20 (0.13)
Planck	0.19 (0.19)	0.14 (0.08)	0.06 (0.04)	0.06 (0.04)
CMB-S4	0.04 (0.04)	0.04 (0.03)	0.03 (0.02)	0.02 (0.02)

Table 1.2: Forecasted 1σ errors on $N_{\rm eff}$ from different combinations of experiments. Numbers in parenthesis assume fixed total neutrino mass, whereas the rest are marginalized over neutrino masses.

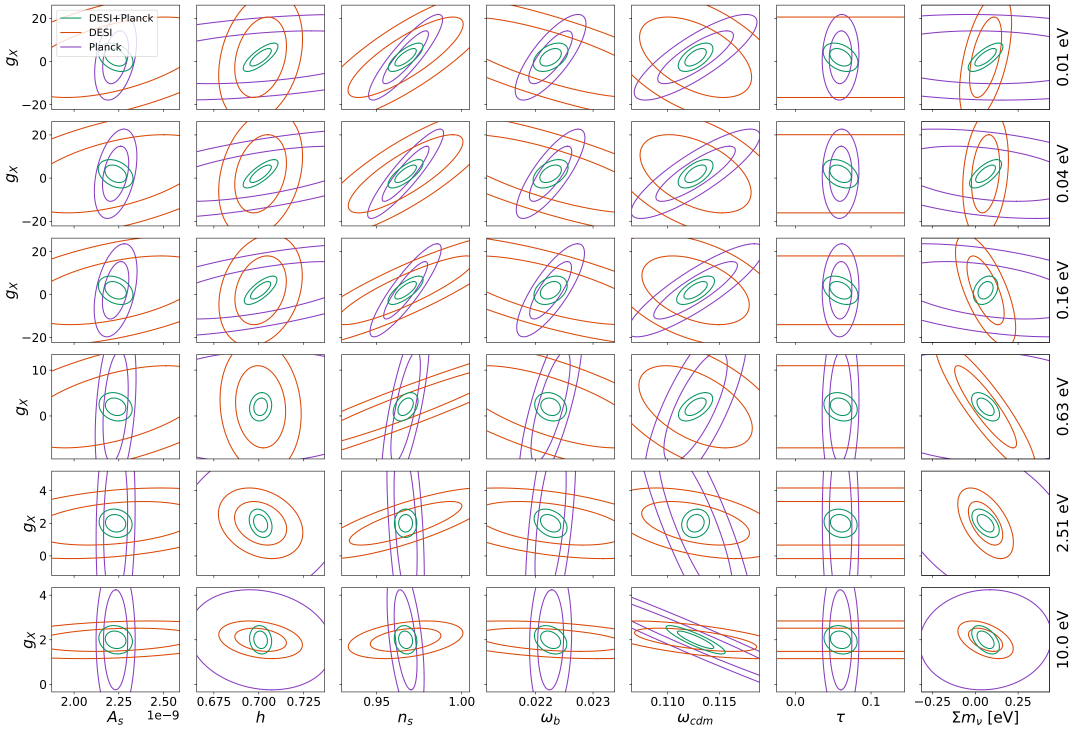

Figure 1.3: 68% C.L. and 95% C.L. projected confidence ellipses for each of the parameters we marginalize over, as well as the LiMR number g_X of degrees of freedom, for DESI (red), *Planck* (purple), and their combination (green). Each row has a different fiducial relic mass, denoted on the right, all with an assumed temperature $T_X^{(0)} = 0.91$ K at $z = 0$. Note that we also marginalize over the unknown neutrino mass, which loosens our constraints by as much as 143% for LSS-only information, 64% for CMB-only information, and 81% for combined LSS and CMB information.

Interestingly, the degeneracy directions change with LiMR mass. As an example, the degeneracy line for g_X and $\omega_{\rm cdm}$ for CMB data changes direction as the LiMR becomes more massive, and starts behaving as matter instead of radiation at recombination. The LSS degeneracy line, however, stays relatively stable, improving the CMB result by different amount at each mass.

The result described above indicates that combining CMB and LSS information is critical for an optimal constraint of LiMRs. We confirm this in Fig. 1.4, where we show the forecasted error in g_X for CMB and LSS data on their own, as well as together, which dramatically improves the constraints. For the rest of this work we will consider different combinations of CMB and LSS surveys together.

We now forecast to which level of significance different LiMR can be constrained, under three different survey combinations. The first is what would be realizable by current data, where we assume galaxy data from BOSS and *Planck* for the CMB. We show the forecasted $\sigma(g_X)$ in Fig. 1.5, which clearly shows that LiMRs with larger $T_X^{(0)}$ and m_X are more readily observable. However, to observe (or rule out) a LiMR at 3σ it has to be relatively heavy ($m_X \gtrsim$ few eV), as we will see below. Note that in this figure we show results for $T_X^{(0)} < 0.91$ K, as for instance a scalar at that minimum temperature would be equivalent to a Weyl fermion with $T_X^{(0)} = 0.79$ K, as we will discuss below.

The second case we consider is the near-future one, where we add DESI data to *Planck* . We show the forecasted constraints on g_X for this combination in Fig. 1.6, which are clearly improved with respect to the results shown in Fig. 1.5. In this case one can rule out relics of any mass with $T_X^{(0)} = 1.4$ K at 3σ. More interestingly, we see that masses above 1 eV would be ruled out, even for the lowest possible relic temperature of $T_X^{(0)} = 0.91$ K.

The final case we consider is more futuristic, and adds CMB-S4 data to DESI. We show the results in Fig. 1.7, which further improves the prospects for detecting light relics. In this case even relics at low temperatures can be ruled out at 3σ confidence for masses above 0.78 eV, whereas minimum-temperature massless Weyl relics can only be found at 0.5σ confidence.

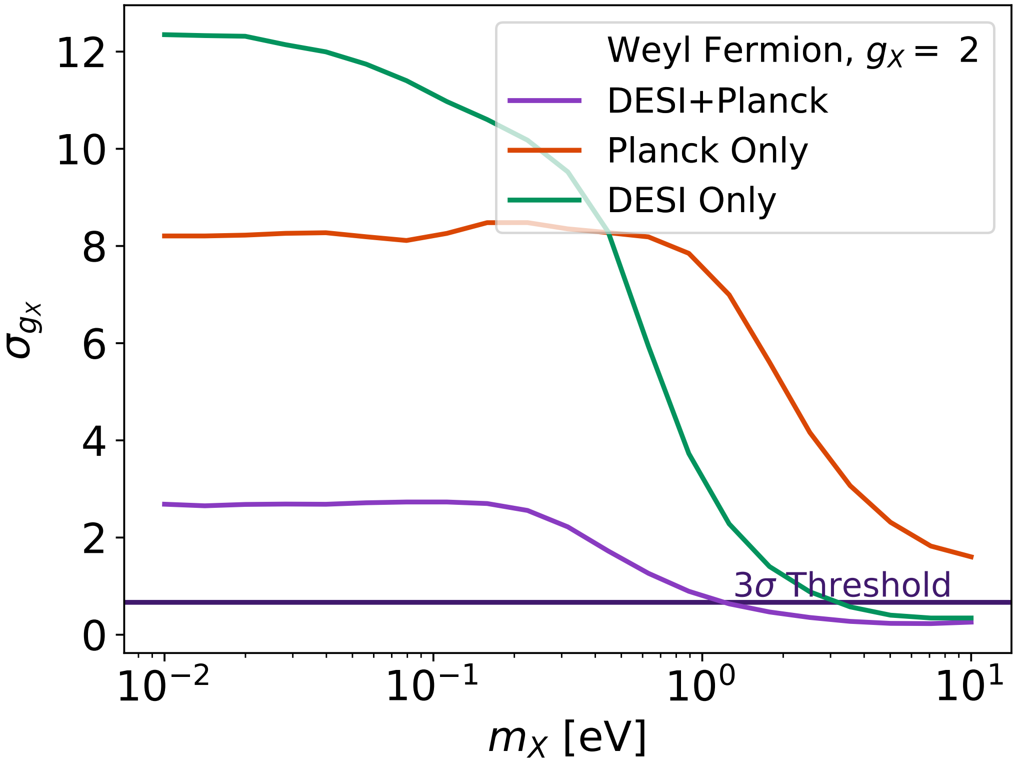

Figure 1.4: Improvement of Weyl relic measurements by addition of LSS data with DESI and *Planck* constraints. The relic is fixed at its minimum possible temperature, $T_X^{(0)} = 0.91$ K. As shown, the joint constraints are much stronger than the LSS or CMB alone.

1.4.2 MINIMUM TEMPERATURE

While the figures discussed above covered a broad range of temperatures and masses, they all assumed a Weyl relic. Here we extend our results to other types of relics, focusing on the minimum temperature of $T_X^{(0)} = 0.91$ K, corresponding to the earliest decoupling from the SM plasma. We divide our results into fermionic and bosonic relics. The cumulative results of our forecast for each type of particle are tabulated in Table 1.3.

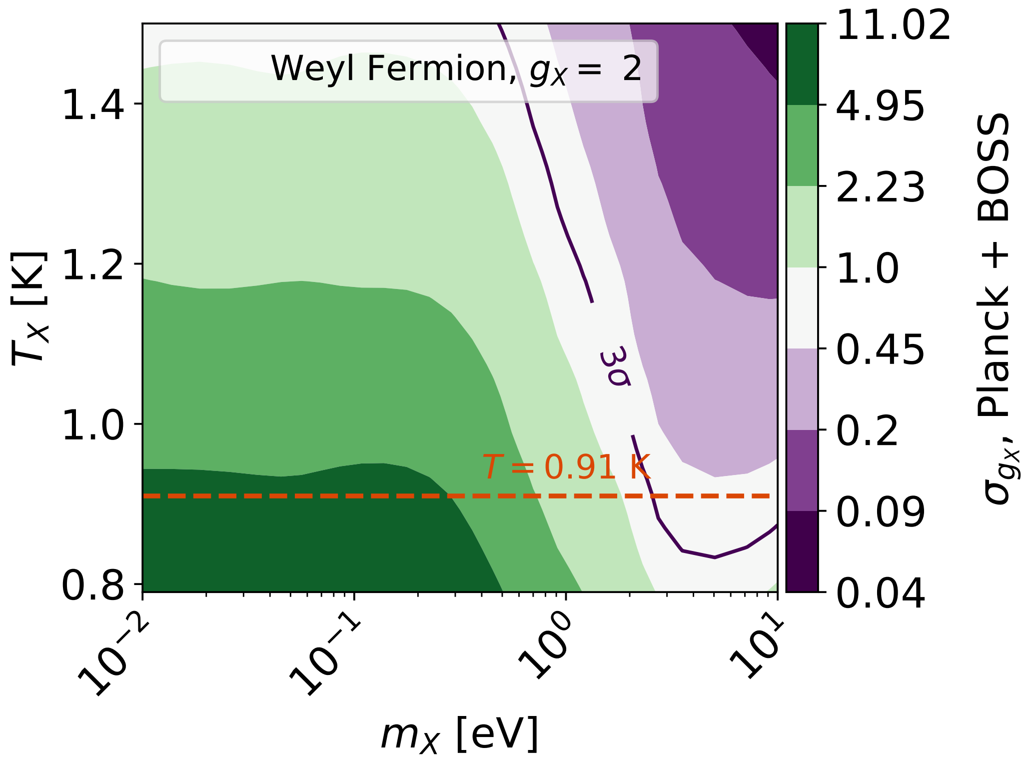

Figure 1.5: Forecasted errors on g_X for a Weyl (neutrino-like) relic of different fiducial masses and temperatures, in all cases with fiducial $g_X = 2$, assuming BOSS+*Planck* data. The region of parameter space measurable at the 3σ-level lays rightward of the purple solid line, and the dashed red line shows the minimum temperature expected for a relic.

FERMIONIC RELICS

We start with a massive Weyl fermion with $T_X^{(0)} = 0.91$ K, for which we show our forecasts on $\sigma(g_X)$ for various combinations of galaxy surveys and CMB experiments in Fig. 1.8a, with a finer mass resolution than the results above. We report the minimum relic masses that are observable at 3σ significance, both with (and without) marginalizing over the neutrino masses, as a test of how degenerate LiMRs are with the total neutrino mass. The combination of presently available *Planck* and BOSS datasets are forecasted to observe or rule out LiMRs above 2.85 (2.47) eV at 3σ signifi-

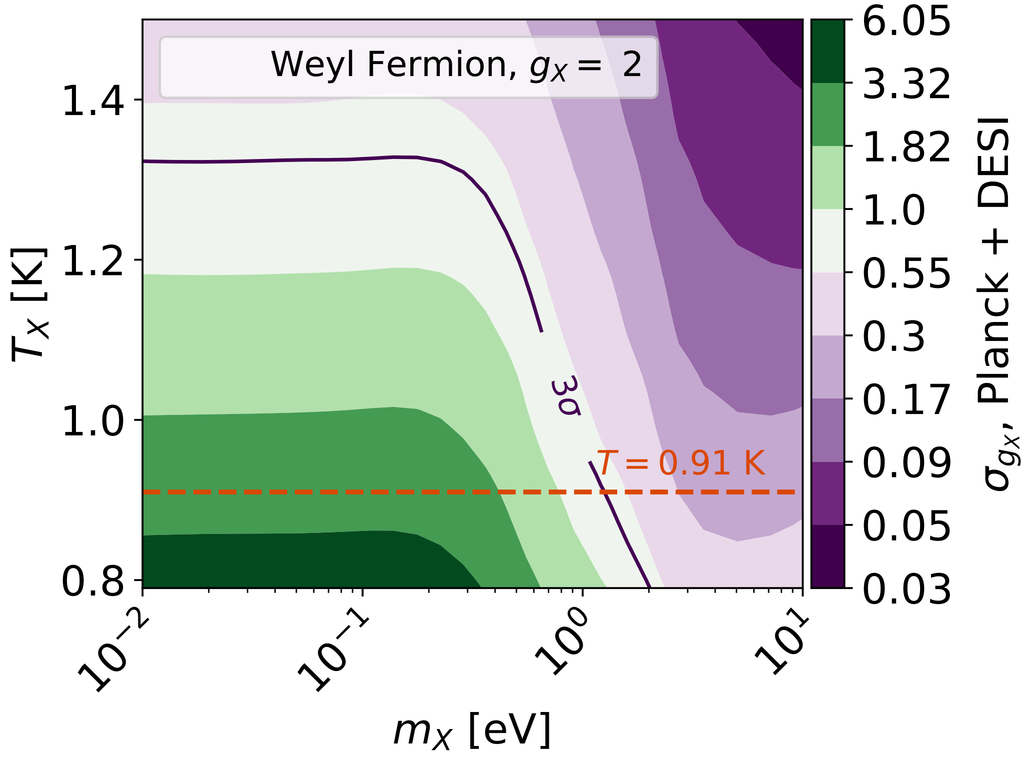

Figure 1.6: Same as Fig. 1.5 for DESI + *Planck*.

cance. For *Planck* and DESI, this is lowered to LiMRs with masses above 1.20 (1.00) eV. This result should motivate an analysis using presently available datasets. For the futuristic combination of CMB-S4 and *Euclid* datasets, we show that LiMR masses above 0.63 (0.59) eV can be observed or ruled out at 3σ significance.

As an example of the physical implications of these constraints, let us apply to them to the ($s = 3/2$) gravitino, which is related to the scale of SUSY breaking in some models. The gravitino is cosmologically equivalent to the neutrino-like Weyl relic that we have studied, as only the $s = 1/2$ modes are thermalized with the SM plasma in the early universe[161], and are expected to have the lowest relic temperature of 0.91 K. This has allowed previous work to constrain the gravitino mass

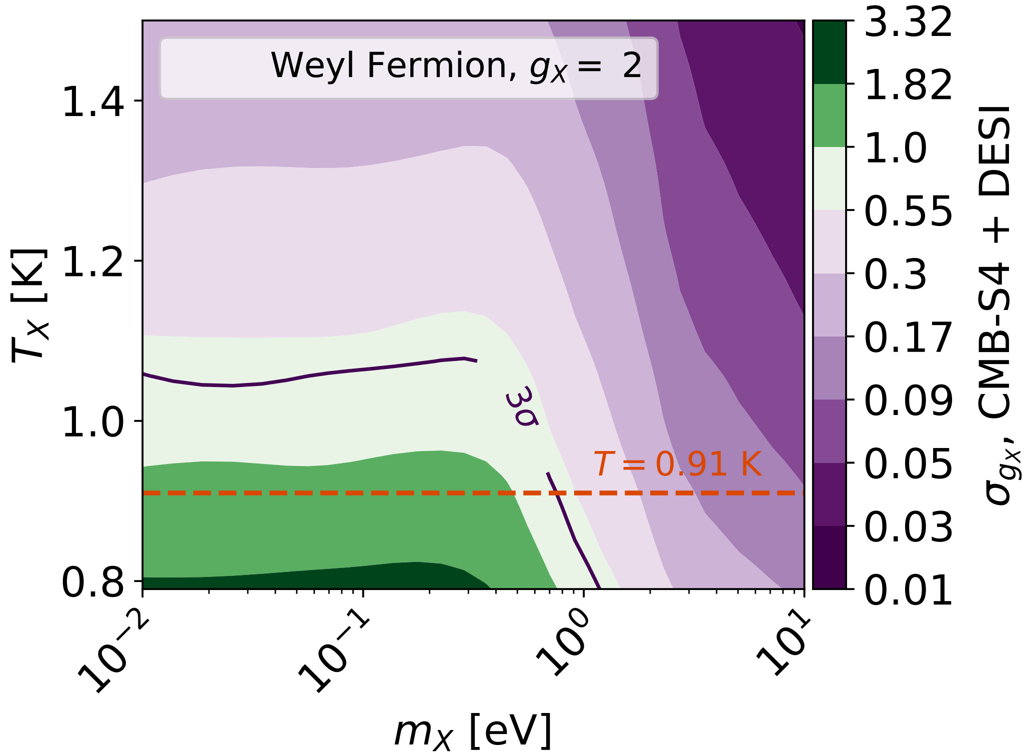

Figure 1.7: Same as Fig. 1.5 for DESI + CMB-S4.

by requiring that their abundance does not overcome that of the cosmological dark matter[151]. Our forecast above shows that current data is sensitive to gravitinos heavier than $m_X = 2.85$ eV, which is around the benchmark of some models of SUSY breaking[110,109], and a factor of a few better than the best limits currently available[199,161]. Upcoming data from CMB-S4 combined with *Euclid* is expected to further detect such gravitino population masses above 0.63 eV. Under the assumption that a cosmological gravitino population no longer exchanges entropy after decoupling from the SM bath, we can relate constraints on m_X to bounds on the SUSY breaking scale $\Lambda_{\rm SUSY} \sim \sqrt{m_X M_{\rm Pl}}$[76,144]. Our forecasted *Planck* and BOSS dataset translates to an upper bound $\Lambda_{\rm SUSY} \lesssim 80$ TeV, whereas the CMB-S4 and *Euclid* datasets lower this to $\Lambda_{\rm SUSY} \lesssim 50$ TeV. These

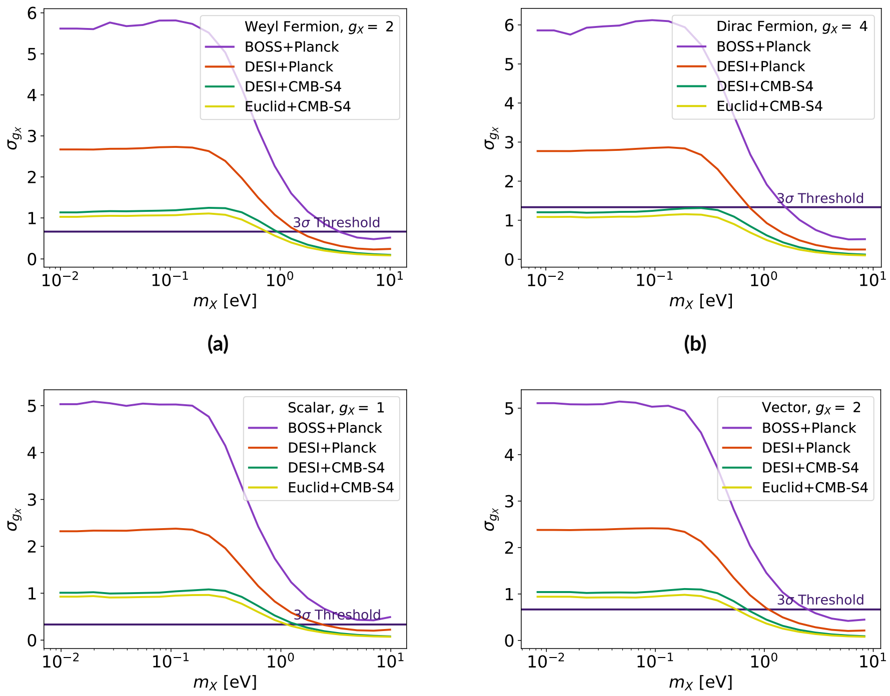

(a)

(b)

(c)

(d)

Figure 1.8: Forecasted error on the relic degrees of freedom for a neutrino-like Weyl fermion (with fiducial $g_X = 2$, top left), a Dirac fermion ($g_X = 4$, top right), a real scalar ($g_X = 1$, bottom left), and a vector particle ($g_X = 2$, bottom right), all at their minimum temperature $T_X = 0.91$ K, for various combinations of CMB + LSS experiments. The horizontal line denotes the uncertainty required to detect each relic at 3σ.

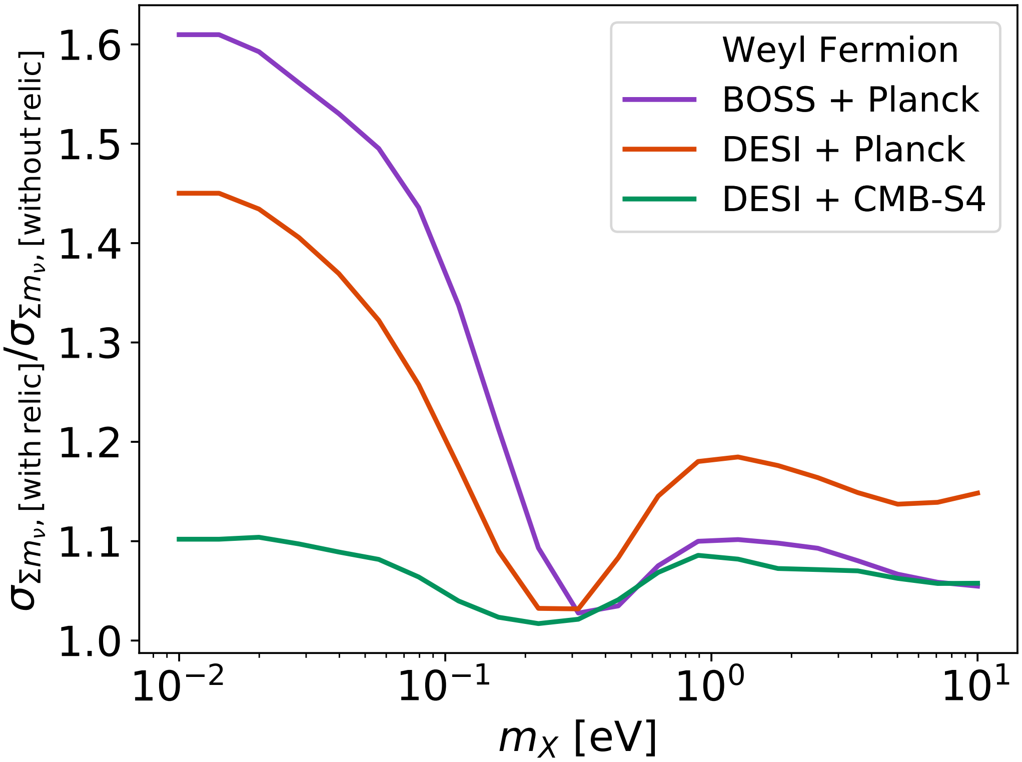

Figure 1.9: Forecasted DESI + CMB-S4 uncertainty on the sum $\sum m_\nu$ of neutrino masses, when it is jointly searched for with a relic of mass m_X and temperature $T_X = 0.91$ K. The degeneracy is minimized at $\sim \mathcal{O}(0.3\mathrm{eV})$ for all particle types, although the constraints on neutrino masses using CMB data from *Planck* are always expected to weaken by $\sim 10\%$, if a new light relic is present.

projections are interestingly complementary to the energy range that will be reached by the proposed $\mathcal{O}(100\ \mathrm{TeV})$ particle collider, showing the promise of our approach.

We also consider a Dirac fermion, with $g_X = 4$ and mass m_X. In terms of the equivalent Weyl fermion, this corresponds to a temperature $T_W^{\mathrm{eq}} = 1.08$ K and mass $m_W^{\mathrm{eq}} = 1.19\, m_X$. In Fig. 1.8b, we show that the combined *Planck* and BOSS datasets are forecasted to observe or rule out such particles above 1.30 (1.12) eV at 3σ significance. For *Planck* and DESI, the 3σ constraint is lowered to 0.61 (0.52) eV. Interestingly, CMB-S4 data will enable the parameter space of Dirac fermions with any mass to be observed or ruled out at 3σ significance when combined with LSS data from DESI.

Bosonic Relics

We now move to bosonic degrees of freedom. First, we study a minimum-temperature real scalar, with $s = 0, g_X = 1$, and mass m_X. This is equivalent to a Weyl relic with $T_W^{\rm eq} = 0.79$ K and $m_W^{\rm eq} = 1.01 m_X$. We show in Fig. 1.8c that, while the combination of presently available *Planck* and BOSS datasets cannot constrain scalar relics at the 3σ significance, DESI and *Planck* can jointly rule out scalars with masses above 1.96 (1.61) eV. Further, the combination of CMB-S4 with either the DESI or *Euclid* datasets can observe or rule-out real scalar bosonic relics above 1.14 (1.06) and 0.93 (0.87) eV, respectively.

Second, we consider a massive vector, with $s = 1$ and $g_X = 2$. This massive vector is equivalent to a Weyl relic with $T_W^{\rm eq} = 0.94$ K and $m_W^{\rm eq} = 1.21 m_X$. In Fig. 1.8d we show that the combination of *Planck* and BOSS datasets can observe or rule-out massive vector bosonic relics above 2.05 (1.79) eV, whereas substituting BOSS for DESI improves this number to 0.90 (0.75) eV. Combining the CMB-S4 and *Euclid* datasets further improves this to 0.47 (0.44) eV.

1.4.3 Neutrino-mass forecasts

We have detailed in each previous subsection the constraints with and without marginalizing over neutrino masses to emphasize the importance of this step, as it is seen to affect results noticeably when LSS information is being considered. We note that DESI is particularly sensitive to the marginalization or fixing of $\sum m_\nu$. This is due to its chosen bias prescription, which does not include a parameter to marginalize over the redshift dependence of the bias, as opposed to BOSS and *Euclid*. This underscores the sensitivity of our results to the details of the bias prescription, which is further

	CMB Only	BOSS	DESI	*Euclid*
Scalar m_X[eV]				
LSS Only		- (-)	4.98 (4.54)	3.24 (3.22)
Planck	- (-)	- (-)	1.96 (1.61)	1.31 (1.16)
CMB-S4	1.48 (1.44)	1.41 (1.31)	1.14 (1.06)	0.93 (0.87)
Weyl Fermion m_X[eV]				
LSS Only		- (-)	3.13 (2.78)	2.42 (2.41)
Planck	- (-)	2.85 (2.47)	1.20 (1.00)	0.87 (0.78)
CMB-S4	1.03 (1.02)	0.98 (0.91)	0.78 (0.71)	0.63 (0.59)
Vector m_X[eV]				
LSS Only		- (-)	2.41 (2.08)	1.88 (1.88)
Planck	- (-)	2.05 (1.79)	0.90 (0.75)	0.65 (0.60)
CMB-S4	0.81 (0.78)	0.75 (0.70)	0.58 (0.54)	0.47 (0.44)
Dirac Fermion m_X[eV]				
LSS Only		4.06 (3.72)	1.82 (1.36)	1.50 (1.50)
Planck	- (-)	1.30 (1.12)	0.61 (0.52)	0.45 (0.43)
CMB-S4	0.56 (0.55)	0.51 (0.48)	**All (All)**	**All (All)**

Table 1.3: Minimum mass at which a LiMR (scalar boson, Weyl fermion, vector boson or Dirac fermion, from top to bottom) can be observed or ruled out at 3σ significance. Also reported in parentheses is the result with fixed $\sum m_\nu$ (to its fiducial value). A“—” sign corresponds to no masses within the 3σ constraint. “**All**” corresponds to all LiMR masses analyzed being within the 3σ constraint.

explored in our companion paper[205].

As a consequence of our analysis, we can also forecast how much neutrino-mass measurements would be affected by the presence of a LiMR, given the degeneracies between $\sum m_\nu$ and g_X shown in Fig. 1.3. We show in Fig. 1.9 the relative increment in the error of the sum $\sum m_\nu$ of neutrino masses when marginalizing over a relic of varying mass. For reference, we forecast $\sigma(\sum m_\nu)$ to be 61.1×10^{-3} eV for BOSS and *Planck*, 28.2×10^{-3} eV for DESI and *Planck*, and 24.1×10^{-3} eV for DESI and CMB-S4, with a fiducial at the (normal-hierarchy) minimum $\sum m_\nu = 60 \times 10^{-3}$ eV and no other relics. The degradation in the expected errors ranges from 10% for heavy relics and futuristic data (DESI+S4), to nearly 100% for lower masses and current or upcoming data. (BOSS/DESI+*Planck*). Note that for relics of $m_X \approx 0.3$ eV the degradation minimizes in all survey specifications. This mass corresponds to relics that become non-relativistic around the time of recombination. In essence, heavier relics produce suppression in the matter fluctuations, whereas lighter relics chiefly affect CMB and LSS observables through their change in $N_{\rm eff}$. We encourage the reader to see our companion paper[205] for in-depth neutrino forecasts without relics.

1.5 CONCLUSIONS

In this work we have studied how well current and upcoming cosmological surveys can detect light (but massive) relics (LiMRs), focusing on the 10^{-2} eV to 10^{1} eV mass range. These particles become non-relativistic before $z = 0$, and thus affect the formation of structures in the universe. By combining information from the CMB and the LSS we have shown that a large swath of the 2D-parameter space (of relic mass and temperature) will be probed by upcoming surveys.

There is a minimum temperature that any relic that was in thermal equilibrium with the Standard Model should have, $T_X^{(0)} = 0.91$ K. Interestingly, we find that Weyl, vectors, and Dirac relics with this temperature, and masses above ≈ 1 eV, can be observed or ruled out at the 3σ significance using the presently available combination of *Planck* and BOSS datasets. Looking slightly to the future, the *Planck* and DESI datasets will improve these constraints, and reduce the minimum mass allowed for LiMRs by roughly 50%. The more futuristic *Euclid* and CMB-S4 datasets will present an 80% improvement and, in the case of Dirac fermions, fully cover the parameter space. If the sum of neutrino masses, $\sum m_\nu$, can be learned independently of CMB and LSS surveys, the effect of fixing the $\sum m_\nu$ parameter manifests as an approximate 20% improvement on these constraints. This could be accomplished, for example, by KATRIN which currently sets the leading upper bound on the effective electron neutrino mass of 1.1 eV, independently of cosmology[21]. We emphasize that the effect of marginalizing $\sum m_\nu$ significantly weakens the 3σ constraints for some of the cases reported, suggesting that it is important to account for $\sum m_\nu$ in any search for LiMRs. While the need to properly account for $\sum m_\nu$ has been discussed in previous work[43,100,41,81,89,37], our analysis, which does so for massive but light relics, is unprecedented.

This result is particularly interesting for the case of the gravitino. Since the gravitino would have a cosmological imprint identical to a Weyl fermion, we have shown that *Planck* and BOSS can observe or rule out gravitinos heavier than 2.85 eV. If a gravitino, or any other LiMR, were detected, then their parameters (i.e., mass and temperature) could also be measured, as suggested in Ref.[36].

In summary, while light relics are commonly assumed to be nearly massless — and constrained through $N_{\rm eff}$ — here we have shown that relics with masses on the 10^{-2} eV to 10^{1} eV scale can be constrained with cosmological data. These constraints are broadly expected to apply to the full range of allowed relic masses, from effectively massless to saturating the DM abundance. This complements current efforts in the search of relics, allowing many new routes for finding physics beyond the Standard Model.

1.6 MCMC VALIDATION OF FISHER FORECASTS

In this Appendix we show a comparison of our Fisher formalism and an MCMC analysis of the same mock data to confirm our Fisher analysis throughout the main text. In Fig. 1.10 we show the MCMC (solid) and Fisher forecasted (dotted) marginalized posteriors for cosmological parameters and nuisance parameters (including the neutrino mass $\sum m_\nu$), assuming CMB-S4 + DESI data. This Figure shows that the predicted errors agree remarkably well between our Fisher-matrix approach and the full MCMC of mock data.

Moreover, we show posteriors for models with and without the growth induced scale-dependent modification to the bias (as described in our companion paper[205]), which we termed GISDB. The MCMC results are from Ref.[205], and the Fishers are calculated here. The non-GISDB Fisher ellipses are centered on the corresponding MCMC maximum likelihood point. The GISDB ones, however, are shifted by[153]

$$\delta\theta_i = (F^{-1})_{ij} D_j, \tag{1.30}$$

in each parameter θ_i, where we have defined

$$\begin{aligned} D_j = \sum_z \int k^2 \mathrm{d}k \int \mathrm{d}\mu \frac{V(z)}{2(2\pi)^2} \left(\frac{\partial \log \tilde{P}_g(k,\mu)}{\partial \theta_j} \right) \\ \left(\tilde{P}_{g,\mathrm{GISDB}}(k,\mu) - \tilde{P}_{g,\mathrm{no\ GISDB}}(k,\mu) \right) \left(\frac{\bar{n}\tilde{P}_g}{\bar{n}\tilde{P}_g + 1} \right)^2, \end{aligned} \tag{1.31}$$

and the GISDB Fisher ellipses are computed centered on the shifted best-fit. As shown, the good cohesion between the Fisher and MCMC analyses of the data, particularly in the inclusion of the GISDB effect, demonstrates that the considered effects are well-approximated by the linearity of the Fisher approach, and thus validates the constraints we present on additional light relics.

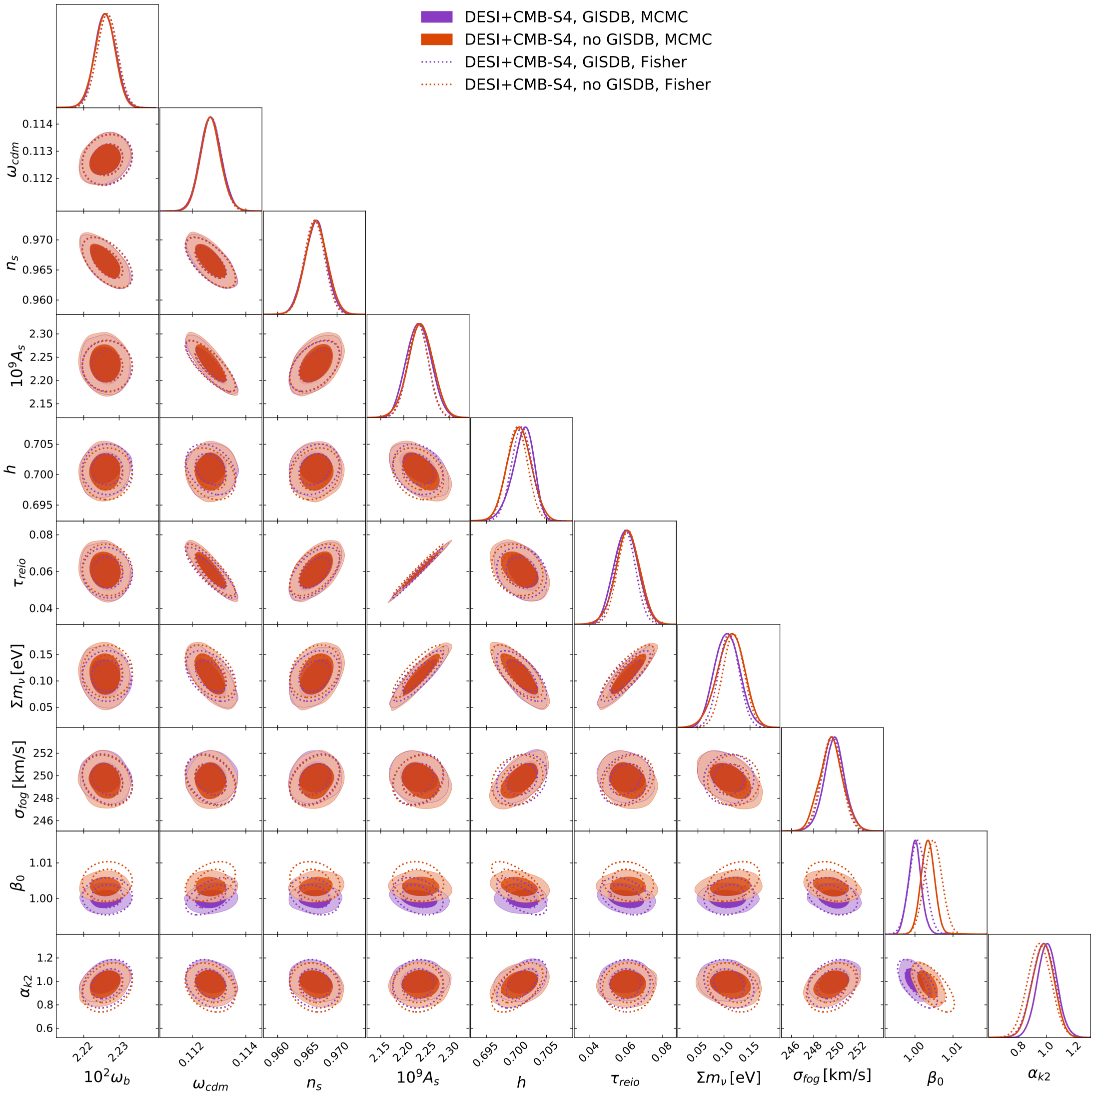

Figure 1.10: MCMC and Fisher forecasted marginalized posteriors for cosmological parameters and nuisance parameters for a joint DESI + CMB-S4 analysis. The degenerate hierarchy is assumed with a total mass of $\sum m_\nu = 0.1$ eV. Models with and without the bias step (GISDB) are considered. As shown, the good consistency between MCMC and Fisher results, particularly the reproduced shift in parameters upon turning off GISDB, demonstrates that the effects we consider are well-captured at linear order and validates our results regarding the detectability of LiMRs.

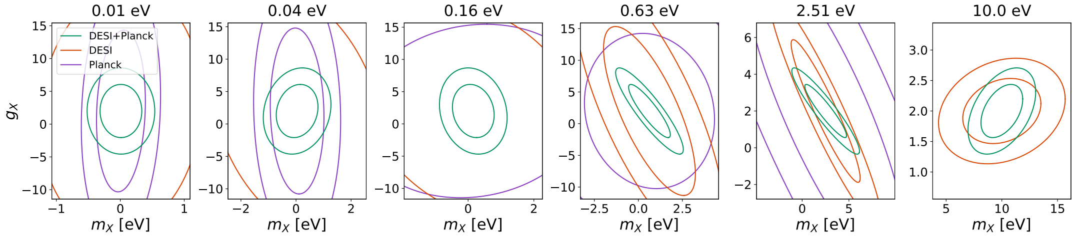

Figure 1.11: Fisher-matrix forecasted marginalized posteriors for the parameters g_X and m_X. In this forecast, the LiMR mass has been allowed to vary in addition to its degrees of freedom. We present the marginalized posterior contours for five choices of the fiducial LiMR mass: 10^{-2} eV, $10^{-1.4}$ eV, $10^{-0.8}$ eV, $10^{-0.2}$ eV, $10^{0.4}$ eV and 10^{1} eV. As shown, the degeneracy lines are driven by the relative orthogonality of CMB information at low masses, and by strong degeneracy in the LSS data at intermediate to high masses.

1.7 Marginalization Over the Relic Mass

Throughout the main text, the LiMR mass has been held fixed. In this appendix, we allow the LiMR mass to vary in the forecasts to study what effect this has on the LiMR constraints presented earlier, as well as to study how well a prospective LiMR detection could constrain its properties.

For all combinations of LiMR species, galaxy surveys and CMB experiments studied in this work, we find that marginalizing over the LiMR mass m_X weakens the constraint on the relic degrees of freedom g_X, as expected. This effect is most exaggerated in the cases where the constraint is dominantly set by LSS information. In a joint *Planck*-BOSS analysis, high-mass relics (with $m_X \geq 0.2$ eV) see the g_X constraint weakened by nearly a factor of 2. In cases where CMB information dominates, however, such as when adding CMB-S4 to BOSS, the g_X constraint is weakened by no more than 6%. Adding Planck information to DESI, the higher-mass region sees the g_X constraint weakened by no more than a factor of 2. Adding CMB-S4 to DESI, the g_X constraint is weakened by no more than 25%.

In Figs. 1.11 and 1.12, we illustrate this effect, assuming that a Weyl fermion with fiducial $T_X = 0.91$ K and different values of m_X is observed using different combinations of galaxy and CMB surveys. The broadening of the error bars is primarily driven by the LSS information and, as a con-

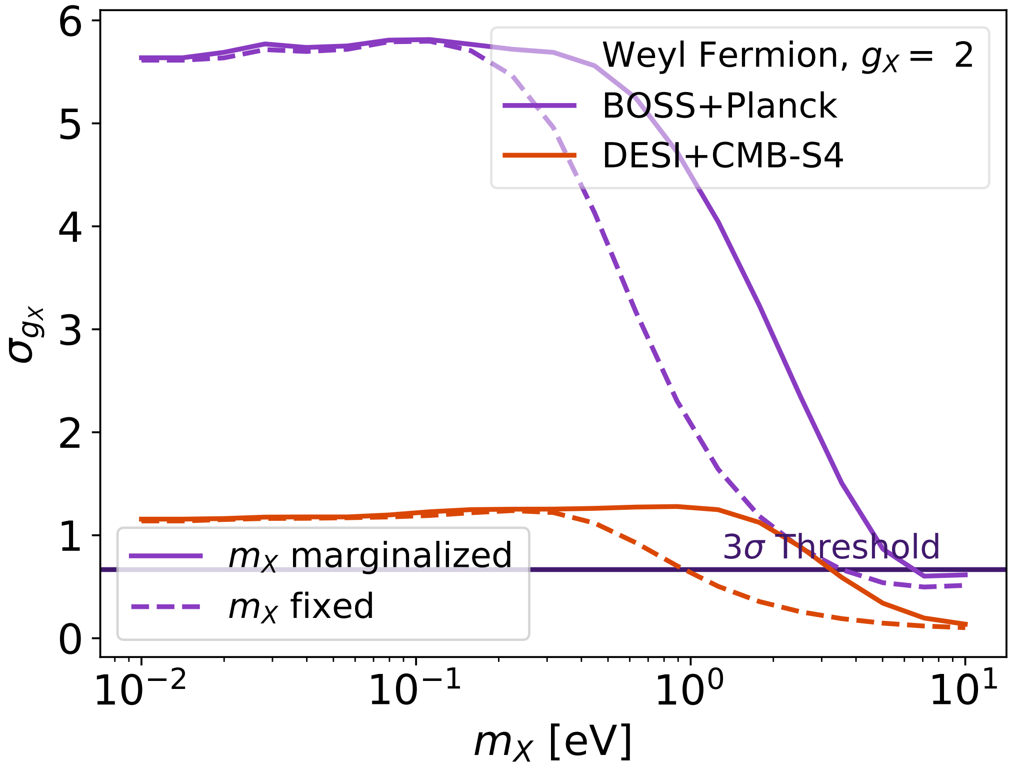

Figure 1.12: Forecasted sensitivity on the relic degrees of freedom g_X for a Weyl fermion with (solid) and without (dashed) marginalization over relic mass m_X for combinations of datasets BOSS + *Planck* and DESI + CMB-S4. As expected, there is little degeneracy in the low-mass limit, where the relic mainly contributes as $N_{\rm eff}$; the $g_X \quad - \quad m_X$ degeneracy enters most at intermediate masses.

sequence, the biggest shift in constraints is observed for datasets that are primarily or exclusively constrained by the galaxy surveys.

We see in Fig. 1.11 that the Planck constraint monotonically weakens with increasing fiducial relic mass. This can be explained by the decreasing effect of a relic on the radiation energy density ρ_r which the CMB is primarily sensitive to. At low masses, the Planck dataset demonstrates an orthogonal relationship between the relic mass and degrees of freedom. Considering that at low masses, changes in mass will modify the weak lensing signal of the CMB and produce no change in ρ_r yet small changes in g_X will produce directly proportional changes in ρ_r we expect a nearly orthogonal relationship between these two parameters at low masses primarily governed by Eq. (1.4), as the CMB signal is dominated by changes to ρ_r. However, as the fiducial relic mass is increased, and the relic effect on ρ_r at recombination becomes smaller, the CMB becomes sensitive to the relic primarily through its effect on the weak lensing signal and the governing relationship changes to Eq. (1.6) which is directly proportional to the product of m_X and g_X. Thus, these two parameters are expected to develop an anti-correlation at high masses in the CMB dataset, which is indeed what we observe in Fig. 1.11 at higher masses.

Now we consider how the degeneracy direction in the $m_X - g_X$ plane varies at different relic masses for the LSS datasets. Here the scales affected by the relic, as governed by Eq. (1.5), and by the magnitude of the effect, as determined by Eq. (1.6), control the effect on the LSS signal. At low masses, the contribution of the relic to ω_M is small and the relic will primarily affect the LSS signal through its free-streaming scale, which is independent of g_X. So at small relic masses, we expect m_X and g_X to be approximately orthogonal. As the fiducial relic mass is increased, the contribution of the relic to P_M and hence to the LSS signal increases and is again proportional to the product of m_X and g_X. So with increasing relic mass, we generally expect an anti-correlation to develop between the relic mass and degrees of freedom. We again see this to be the case in Fig. 1.11.

As discussed above, allowing the relic mass to vary modifies the constraints of the LSS and CMB datasets such that the accuracy of those constraints is generally less affected for lower mass relics. As the relic occupies a greater portion of Ω_M, it becomes more important to simultaneously vary the relic mass and degrees of freedom. We emphasize that for a fixed relic abundance, there is a degeneracy between the relic parameters m_X, T_X, g_X according to Eq. (1.6). This allows us to translate constraints on any two of these parameters into constraints on the third parameter. Where we have allowed the relic mass and degrees of freedom to vary, the resulting constraints can be translated to errors on the temperature. We also bring attention to the fact that marginalizing over the relic mass is only valid in the neighborhood of parameter space around each fiducial choice, and not over the entire parameter space of masses permitted.

1.8 SAMPLING OF FULL MODEL POSTERIOR FORECASTS

Datasets with different parameter degeneracies can powerfully constrain parameters when combined. To illustrate this complementary effect between CMB and LSS surveys, we present a sampling of fully marginalized posteriors in Fig. 1.13 for a Weyl (neutrino-like) relic with temperature 0.91 K and mass 0.01 eV. In each figure, we present constraints using only DESI (red), only *Planck* (violet), and the joint dataset (green).

As in the case of the LiMR parameter g_X (number of degrees of freedom) discussed in the main text, the addition of LSS information to CMB data will generally break degeneracies between parameters. As an interesting example, we observe that the LSS provides a measurement of $\omega_{\rm cdm}$ that is very close to orthogonal from the CMB one, breaking degeneracies with A_s, n_s and g_X for very light relic masses. DESI information also serves to set the measurements on h and $\sum m_\nu$, which are poorly measured by *Planck* as their effects on the CMB are degenerate. In turn, the LSS by itself is generally ineffective at measuring the other cosmological parameters, and provides no information on $\tau_{\rm reio}$. While, as illustrated in Fig. 1.3, the degeneracies between g_X and other parameters shift significantly between relics of different masses, those between the cosmological parameters themselves remain largely unchanged.

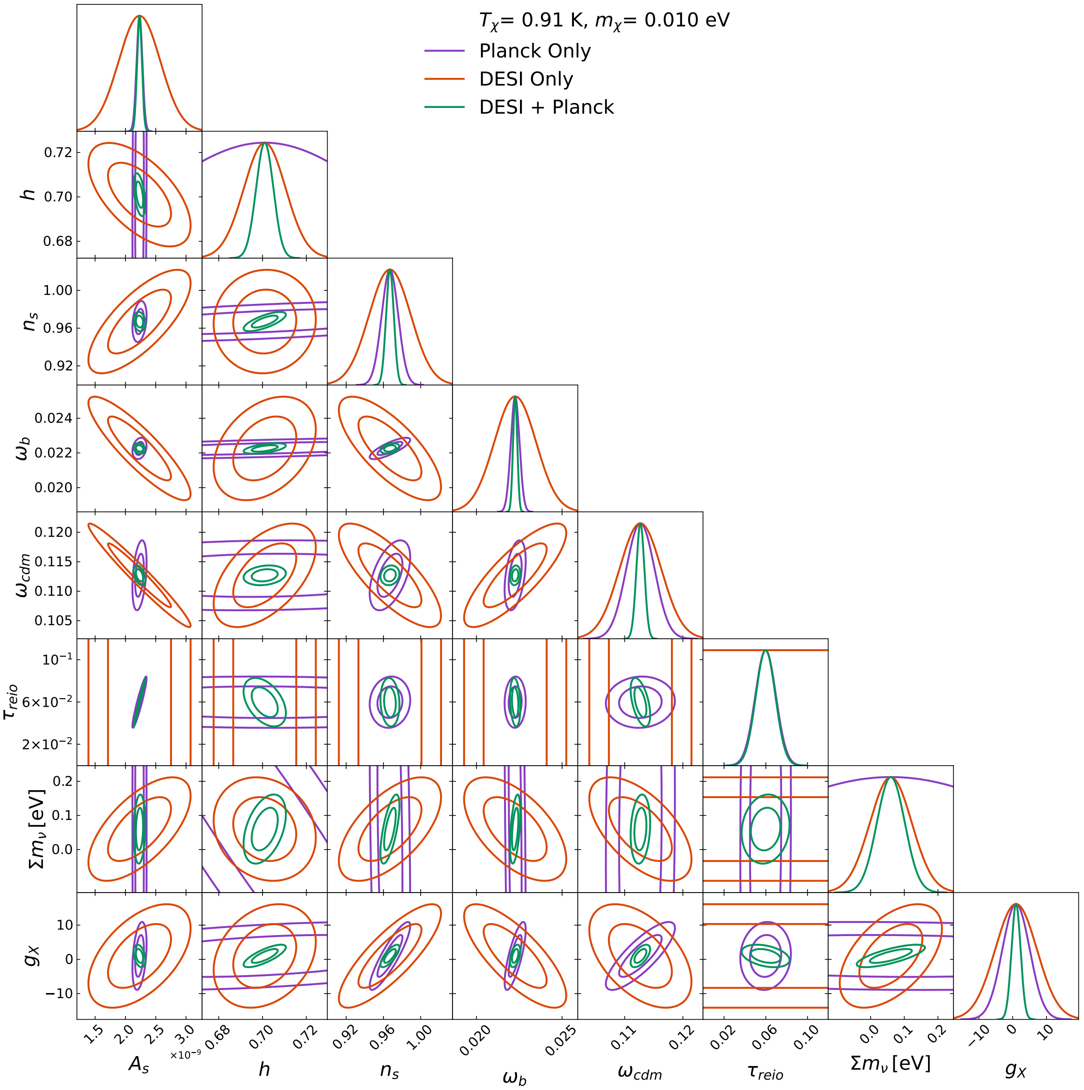

Figure 1.13: 2-dimensional posterior distributions for parameter forecasts using DESI + *Planck* , and each experiment individually. We assume here the presence of a Weyl fermion LiMR ($g_X = 2$) with T_X = 0.91 K and m_X = 0.01 eV. As shown, the complementarity between the two datasets results in marked improvement on the sensitivity to such a relic.

2

Accurately Weighing Neutrinos with Cosmological Surveys

2.1 Introduction

The existence of neutrinos has long been established, but comparatively little is known about them, due to their weak couplings to the visible sector. Although in the Standard Model (SM) neutrinos are massless, compelling evidence of flavor oscillations from solar, atmospheric, and reactor fronts [146,121,159,98,88,138,137] yield measurements of two mass splittings, indicating non-zero masses for at least two of the three neutrino species. The sign of one of the measured mass splittings is yet to be determined, suggesting that neutrinos are ordered in one of two scenarios: the normal hierarchy (NH), where the two lighter neutrinos are closer in mass, or the inverted hierarchy (IH), where the two heavier ones are. Distinguishing between the two neutrino hierarchies, as well as measuring their overall mass scale, are integral steps towards amending the Standard Model via characterizing its least-understood fermions.

Current results from the KATRIN Tritium decay experiment have improved the limits on the mass m_ν of each neutrino species to $m_\nu < 1100$ meV in the quasi-degenerate regime[21], and are expected to constrain each neutrino mass in this regime to 200 meV with upcoming data[162]. This is, however, still far from the minimum (total) masses expected for the NH and IH, of 60 and 100 meV respectively. A diverse range of other particle experiments are also underway aiming to fully characterize the oscillation parameters and determine the mass ordering[29,93,15,12,34,20,11]. At the same time, cosmological data sets provide a powerful tool in the search for massive neutrinos, as these particles are very abundant in our universe, with a density per species today of $n_\nu \sim 100$ cm^{-3}, comparable to that of cosmic microwave background (CMB) photons. This cosmic neutrino background influences the formation of large-scale structure (LSS) in the universe: at least two of these species are non-relativistic at the present day and contribute to the dark matter (DM) content. However, their small masses imply large streaming velocities and induce structure suppression at small scales (see Refs. [129,127] for detailed reviews of these effects).

Indeed, the leading constraints on the sum of neutrino masses are currently obtained with cosmological data[198,127,58,26,68]. The latest 2018 *Planck* data, in conjunction with measurements of the baryon acoustic oscillations (BAO) from the Baryon Oscillation Spectroscopic Survey (BOSS), have been used to constrain $\sum m_\nu \leq 120$ meV at 95% C.L.[17]. This measurement is compatible with both the normal and inverted hierarchies, though the available parameter space for the latter will always be smaller. Upcoming data from the CMB Stage-4 (CMB-S4) experiment, as well as the Dark Energy Spectroscopic Instrument (DESI) and *Euclid* galaxy surveys, will reduce these error bars dramatically, and it is expected that these experiments will measure the sum of neutrino masses at least at the 3σ level[82].

As cosmological data (especially that from large-scale structure experiments) become increasingly precise, it becomes critical to accurately characterize the cosmological effects of neutrinos in the analysis of this data. This is crucial both for the correct characterization of cosmological neutrinos and also for the measurement of relevant parameters of structure formation, such as the amplitude of fluctuations or the intrinsic bias of tracers. Currently, searches for massive neutrinos with cosmological data often make two simplifying assumptions, described below.

The first assumption commonly made is that the three neutrinos have the same mass, a configuration commonly termed the degenerate hierarchy (DH). While this is a valid approximation for neutrino masses much heavier than their splittings ($m_{\nu_i} \gg 50$ meV), it fails for the range of masses still allowed by current data. While the dominant cosmological effect of neutrinos is set by the sum of their masses, the distribution of individual masses has an effect that, although subtle, might be detectable by future surveys[130,174,71,119,103,96,104,198,30]. It is possible that the next-generation measurement of total neutrino mass will simply eliminate the inverted hierarchy by ruling out its minimum mass. However, in the case that this measured sum permits both normal and inverted configurations ($\sum m_\nu \geq 100$ meV), it is also worth investigating whether cosmological data has the power to distinguish between the two.

The second approximation is that the halo bias is unaltered by the presence of light degrees of freedom. However, it was shown in Refs.[134,133,152] that the same scale-dependent growth that gives rise to a suppression in the matter power spectrum in the presence of massive neutrinos induces a scale- and redshift-dependent enhancement to the halo bias as well. This growth-induced scale-dependent bias (GISDB) can partially compensate the effect of neutrinos on the galaxy power spectrum and, opposed to other biasing effects that abound in standard cosmology, its amplitude and shape are determined by the neutrino masses, so it must be properly modeled during searches for these particles.

In this work we include both of these effects for the first time (see for instance Refs.[175,197,135] for previous efforts), and forecast constraints on the neutrino mass from the upcoming DESI[16] and *Euclid*[28] surveys, combined with the CMB-S4 experiment[2]. We study whether the omission of these corrections would bias upcoming results. To find the halo power spectrum in the presence of neutrinos with three different masses we employ the publicly available software `CLASS`[49], which we have modified to include the recently developed `RelicFast`[*] [152] code as a native module. We dub this code `RelicCLASS`[†], and in addition to neutrino masses, it can also be used to search for any other light relic (as we do in our companion paper[75]).

This paper is structured as follows: in § 2.2 we briefly review the effect of neutrinos on the LSS observables. In § 2.3 we explain the datasets we consider, which we employ in § 2.4 to forecast the resulting constraints on neutrino masses. We conclude in § 2.5.

2.2 NEUTRINOS AND THEIR EFFECT ON THE LSS

We begin with an overview of the physics of neutrinos and their effects on the LSS observables, which can be divided in two pieces: the suppression of the matter power spectrum, and a modification to the galaxy-halo bias. Both of these effects are most relevant at scales of $k \sim 0.01\, h/\mathrm{Mpc}$ for currently allowed neutrino masses, making galaxy surveys such as DESI and *Euclid* ideal probes, given their expected low noise at those scales.

The SM contains three species of neutrinos corresponding to the electron, muon, and tau lepton

---

[*] `https://github.com/JulianBMunoz/RelicFast`

[†] `https://github.com/wlxu/RelicClass`

flavors, which mix into three non-degenerate mass eigenstates. Various oscillation experiments have measured two mass splittings to be[137]

$$\begin{aligned}\Delta m_{21}^2 &= 79\,\mathrm{meV}^2\\ |\Delta m_{31}^2| &= 2.2\times 10^3\,\mathrm{meV}^2,\end{aligned} \tag{2.1}$$

where the absolute value on the latter measurement denotes ambiguity to which species is heavier. This results in two possible mass configurations, the normal hierarchy, where $m_3 > m_1$, and the inverted hierarchy, where $m_3 < m_1$. Assuming the lightest neutrino is massless, the NH has a total mass sum of ~ 60 meV and the IH has one of ~ 100 meV[129].

Neutrinos decouple shortly before cosmic electron-positron annihilation, and so their present-day temperature $T_\nu^{(0)}$ is offset from that of the photon bath, $T_\gamma^{(0)}$, by the subsequent entropy injection, yielding $T_\nu^{(0)} = (4/11)^{1/3}T_\gamma^{(0)} = 1.95$ K. Thus, a neutrino of mass m_{ν_i} will become nonrelativistic at $z_{\mathrm{NR}} \sim 500(m_{\nu_i}/100\mathrm{meV})$, so by today we expect at least two of the species to be non-relativistic. In that case, the neutrino abundance Ω_ν today is related to the sum of masses simply as[129]

$$\Omega_\nu h^2 = \sum_i \frac{m_{\nu_i}}{93.2\,\mathrm{eV}}, \tag{2.2}$$

where h is the reduced Hubble parameter. The non-zero temperature of neutrinos allows them to freely stream out of dark-matter structure. This defines a free-streaming scale as the wavenumber $k_{\mathrm{fs},i}$ above which neutrinos behave as hot dark matter, given at late times $z < z_{\mathrm{NR},i}$ by

$$k_{\mathrm{fs},i} = \frac{0.08}{\sqrt{1+z}}\left(\frac{m_{\nu_i}}{100\mathrm{meV}}\right) h\,\mathrm{Mpc}^{-1}, \tag{2.3}$$

assuming matter domination[48,25]. We will focus on neutrinos in this paper, but we note that other light (but not massless) relics produce similar effects, and we search for them in our companion paper[75].

2.2.1 EFFECT ON MATTER FLUCTUATIONS

For this discussion it is informative to differentiate between two different types of matter content in the universe: the component consisting of cold dark matter and baryons (CDM+b), which largely follow each other at late times and actually source the formation of galaxies, and neutrinos, which are generally non-relativistic at the present day but do not cluster at small scales[47,57]. At late times, when all neutrinos are non-relativistic, we can define the matter fluctuations as a sum of these two components:

$$\delta_m = f_{\rm cb}\delta_{\rm cb} + \sum_i f_{\nu_i}\delta_{\nu_i}, \tag{2.4}$$

where δ_i is the overdensity in the i component, $f_i = \Omega_i/\Omega_m$ its fraction of the total matter abundance Ω_m, and we define $f_\nu = \sum_i f_{\nu_i}$, where this sum includes all massive neutrinos.

On large scales ($k \ll k_{{\rm fs},i}$), neutrinos ν_i will follow CDM+b fluctuations, so $\delta_\nu \approx \delta_{cb}$, whereas on small scales ($k \gg k_{{\rm fs},i}$) they will freely stream out of matter potential wells and their fluctuations will be suppressed, following $\delta_\nu \propto \delta_{cb}k^{-2}$ [129]. This affects structure formation in two main ways. First, δ_m is suppressed by a factor of $(1-f_\nu)$ with respect to δ_{cb}, as the larger the neutrino abundance, the smaller the fraction of matter content that contributes to the growth of structure. Second, and more important, the absence of small-scale neutrino fluctuations slows the growth of the CDM+b component at large k. If the neutrinos become non-relativistic after matter-radiation equality, this produces an additional scale-dependent suppression on δ_{cb} of roughly $(1-3f_\nu)$. The result in the linear approximation is a total suppression of the matter power spectrum of $(1-8f_\nu)$ for $f_\nu \ll 1$[129].

Both of these effects become present at $k \sim k_{{\rm fs},i}$ for each species, so not only do neutrino masses determine the overall amount of suppression, but also the location in the power spectrum where said suppression occurs. The effect is dominantly determined by the total f_ν, and thus the total neutrino mass $\sum m_\nu$. However, hierarchical neutrinos will each modify the power spectrum at slightly different free-streaming scales due to their individual masses. For instance, for the same total neutrino mass of $\sum m_\nu = 100$ meV, the normal hierarchy suppression effects turn on at $k_{\mathrm{fs},i} \sim 0.02\,h$ Mpc^{-1} (for the lighter neutrinos) and $k_{\mathrm{fs},i} \sim 0.04\,h\,\mathrm{Mpc}^{-1}$ (for the heavier), while the inverted hierarchy suppression effects turn on at $k_{\mathrm{fs},i} \sim 0.005h\,\mathrm{Mpc}^{-1}$ and $k_{\mathrm{fs},i} \sim 0.03h\mathrm{Mpc}^{-1}$, respectively.

Furthermore, although the difference is subtle, the amplitude of the small-scale suppression of $P_{cb}(k)$ for cosmologies with fixed total f_ν is dependent on the epochs $z_{\mathrm{NR},i}$ where neutrinos begin to behave as matter, with larger suppression for heavier individual species as the growth of fluctuations is slowed from an earlier time. Thus, the scale-dependent suppression is most prominent in the limit where all the mass is carried by one neutrino, and least prominent for the case of three degenerate neutrino; and in general, the inverted scenario will generate more suppression for the same total neutrino mass than the normal one.

In addition to these scale-dependent effects, the inclusion of massive neutrinos while holding fixed the baryon and DM abundances ($\omega_b,\ \omega_{\mathrm{cdm}}$), and the Hubble parameter (h) forces a shift in the dark-energy abundance (Ω_Λ), the effect of which is an overall suppression of the amplitude of fluctuations at all scales. However, this effect is less important for our analysis here, as it can be mimicked by a compensating shift in the amplitude A_s of fluctuations, the Hubble parameter, or the overall halo bias.

We illustrate the effects described above in Figure 2.1, where we incorporate massive neutrinos into the cosmology with various masses and fixed hierarchy (upper panel) and with various hierarchies and fixed total mass (lower panel). We investigate the suppression of each with respect to a cosmology with massless neutrinos, holding fixed the other cosmological parameters $\{\omega_b, \omega_{\mathrm{cdm}}, h, A_s, n_s, \tau_{\mathrm{reio}}\}$ at values listed in Table 2.2 (where n_s is the tilt of primordial fluctuations and τ_{reio} is the optical depth to reionization). Note that the suppression is larger for heavier neutrinos and occurs at smaller scales, and also it is larger for the inverted hierarchy than the normal hierarchy, though only becoming apparent at comparatively small scales.

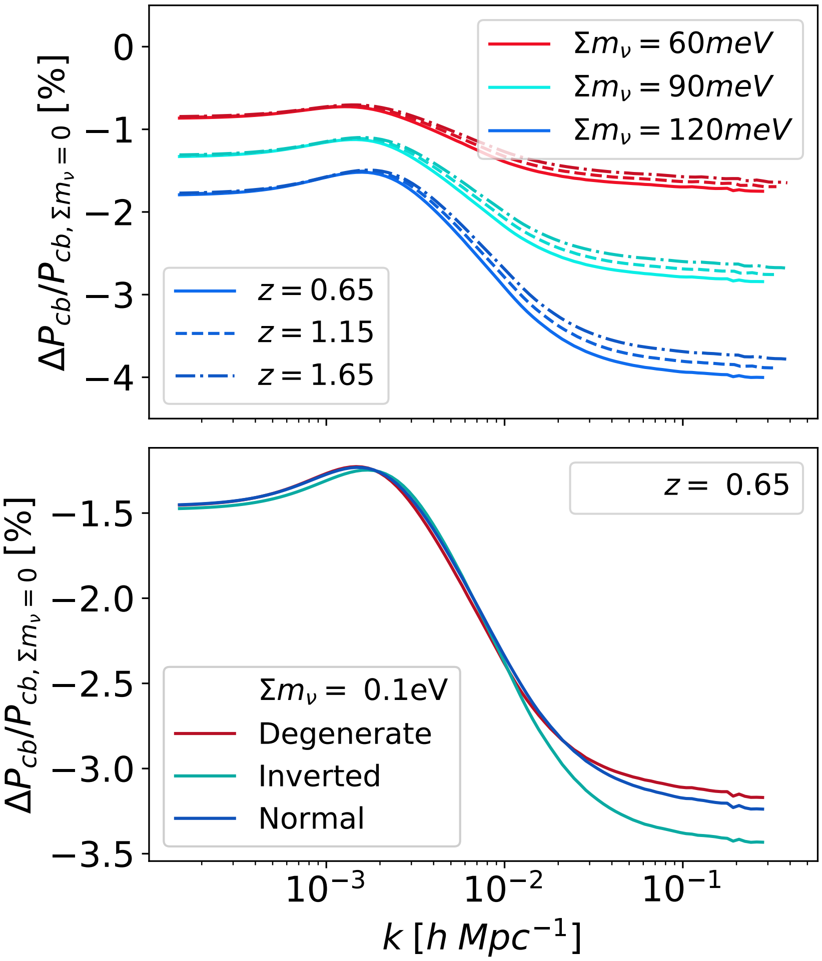

Figure 2.1: Percent differences in CDM+baryon power spectra compared to a massless neutrino cosmology; for different total neutrino masses assuming the degenerate hierarchy (upper panel) and for different hierarchies assuming a total neutrino mass of 100 meV (lower panel). In each case we fix the cosmological parameters $\{\omega_b, \omega_{\rm cdm}, \Omega_m, A_s, n_s\}$, varying h. As shown, the primary effect of massive neutrinos is a suppression of amplitude at small scales – the change in amplitude at large scales is attributed to varying values of h. Note that both the total mass and individual neutrino masses affect the amount of suppression and scale at which it turns on, though the latter effect is subdominant. In addition, the amount of suppression is redshift dependent, with a larger spread at small scales for larger neutrino masses.

2.2.2 Effect on the Bias – `RelicCLASS`

We can observe the neutrino-induced suppression in δ_m by directly measuring the matter power spectrum, for example through weak lensing of the CMB or galaxies[123,10]. Most measurements are, however, of galaxy distributions, which trace the underlying matter fluctuations. In this case it is not enough to study how neutrinos affect the matter power spectrum, but rather it is necessary to find how they change the relation between the halo and the matter overdensities, the galaxy bias.

Due to the nonlinearity of halo formation, the overdensity of haloes traces that of the matter, albeit with a rescaling referred to as the bias. In this work we will always refer to the bias with respect to the CDM+b field, to avoid spurious scale-dependencies due to the non-clustering nature of neutrinos[47,200,58,65]. In that case, the halo fluctuation (without redshift-space distortions) is given by

$$\delta_h(k,z) = b_1(k,z)\delta_{cb}(k,z) \tag{2.5}$$

to linear order, where b_1 is the Eulerian bias, which can be written in terms of the linear Lagrangian bias b^L as

$$b_1(k,z) = 1 + b^L(k,z), \tag{2.6}$$

where the Lagrangian bias is also defined with respect to the cb fluid.

In previous LSS searches for neutrino masses it was typically assumed that either b_1, or its equivalent with respect to all matter, was constant at all scales. Nonetheless, neutrinos produce a scale- and redshift-dependent growth in the CDM+b fluctuations, due to their free-streaming nature. This effect cannot be simply included through transfer functions, due to the non-local temporal nature of the halo-formation process[188,187,133,134], and the process of halo collapse has to be modeled.

We use the publicly available code `RelicFast`, which solves for the spherical collapse of haloes including the effect of neutrinos. In Ref.[152] we found that, while the overall value of the bias is very sensitive to the astrophysics of specific tracers, the scale-dependence of the Lagrangian bias is impervious to those effects within our model.

The result of this correction is a scale-dependent step in the Eulerian bias as shown in Figure 2.2, which we term the *growth-induced scale-dependent bias* (GISDB). We parametrize the bias through

$$b^L(k, z) = \overline{b^L}(z) f(k, z), \tag{2.7}$$

where $f(k, z)$ is numerically computed with `RelicFast` to properly account for the effect of neutrinos in the halo bias, and $\overline{b^L}(z)$ denotes the scale-independent magnitude of the Lagrangian bias, which we will marginalize over. We remind the reader that $f(k)$ can be approximated in terms of $\tanh[\log(k)]$, as it is roughly a step function in log-k space[152]. Nonetheless, we choose to use the full shape of the function, in order to fully capture its physical effect. Additionally, as the amplitude $\overline{b^L}(z)$ of the bias depends very sensitively on the properties of the haloes studied, we will marginalize over it as a free parameter. On the other hand, the scale-dependent behavior, parametrized through $f(k, z)$, is largely independent of those factors[152,133].

The specifications of the `RelicCLASS` code is as follows: We have modified the publicly available Boltzmann solver `CLASS` to include the effect of neutrinos in the halo bias computed by `RelicFast` within `CLASS`, which directly outputs both Eulerian and Lagrangian scale-dependent biases for the input cosmologies. Here, `RelicFast` is included as a `CLASS` module, executed after the `Lensing` module, which outputs the realistic scale-dependent Lagrangian bias for requested ranges of redshifts and halo masses. This output is accessible from the python wrapper to facilitate interfacing with `MontePython` and other codes that take `CLASS` outputs. Inverted, normal and degenerate neutrino scenarios with lightest neutrino mass (m_0 in the code), which sets the full spectrum, are accepted as specifications at the input level. Aside from neutrinos, `RelicCLASS` can also be used to model other scale-dependent effects on the growth function, such as those induced by other light (but massive) degrees of freedom.

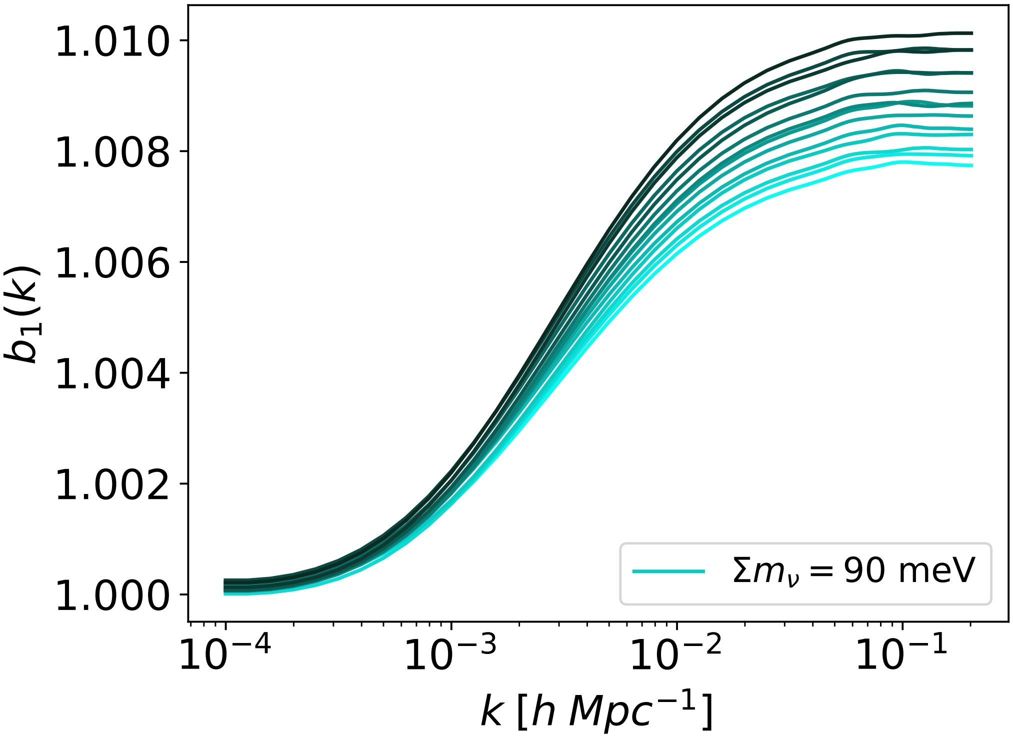

Figure 2.2: The growth-induced scale-dependent bias (GISDB) for redshifts from 0.65 (lightest) to 1.65 (darkest) with massive neutrinos. The total neutrino mass is set at 90 meV and the degenerate scenario is assumed. As shown, the growth-induced scale-dependent bias (GISDB) is both scale- and redshift-dependent.

2.3 DATASETS

In this section we describe the data sets used in our analysis, code specifications, and discuss details of likelihoods and nuisance parametrizations. We use mock data from CMB-S4 as well as either DESI or *Euclid* for the LSS component.

z	0.65	0.75	0.85	0.95	1.05	1.15	1.25	1.35	1.45	1.55	1.65	1.75	1.85	1.95
$\frac{\mathrm{d}N_{ELG}}{\mathrm{d}z\ \mathrm{ddeg}^2}$ [DESI]	309	2269	1923	2094	1441	1353	1337	523	466	329	126	0	0	0
$\frac{dN_{H\alpha}}{\mathrm{d}z\ \mathrm{ddeg}^2}$ [*Euclid*]	2434	4365	4729	4826	4729	4508	4270	3721	3104	2309	1514	1475	894	498

Table 2.1: Forecasted number of ELGs measurable by DESI and Hα emitters measurable by *Euclid* per redshift per deg^2 at each redshift bin z, taken from Refs.[16,28].

2.3.1 GALAXY DATA

For the LSS component we will use a modified version of the basic `pk` likelihood implemented in `Montepython v3`[53] adapted to mock data from the upcoming DESI[16] and *Euclid*[28] surveys, assuming their most abundant tracers. For this analysis we consider emission line galaxies (ELGs) for DESI, and assume the baseline survey covering 14,000 deg^2. Conversely, for *Euclid* we study Hα emitters and assume the reference efficiency given in Ref.[28] with 15,000 deg^2 coverage. The projected galaxy number densities achievable by the DESI and *Euclid* surveys are given in Table 2.1. Somewhat more optimistic constraints could potentially be achieved through multi-tracer techniques[51]. Nonetheless, our goal in this work is to determine whether the inclusion of different neutrino-induced effects would bias the results from upcoming surveys, so we will limit ourselves to the case of one tracer per survey.

The shot noise spectrum is given simply by the inverse of the observed galaxy density,

$$\overline{n}_g^{-1}(z) = \frac{dV(z)}{dN_g(z)}, \tag{2.8}$$

where $dV(z)$ is the comoving volume of the shell at redshift z in the fiducial cosmology, and $dN_g(z)$ is the total number of tracer galaxies within the shell, as computed in Table 2.1.

Galaxies are located in the line-of-sight $\hat{n}$ direction at their measured redshift z. Gravitational attraction of galaxies into clusters, as well as bulk velocities of the clusters themselves, give rise to distortions to the inferred 3D positions of galaxies, which are usually termed redshift-space distortions (RSD)[120]. To linear order we can relate the redshift-space galaxy power spectrum to the CDM+b one as

$$P_g(k, z, \mu) = \left[b_1(k, z) + f_{cb}(k)\mu^2\right]^2 P_{cb}(k, z), \tag{2.9}$$

where $\mu = \hat{k} \cdot \hat{n}$, and we have defined the growth factor of CDM+b fluctuations as

$$f_{cb}(k, z) = -\frac{d \log \sqrt{P_{cb}(k, z)}}{d \log z}, \tag{2.10}$$

which we compute numerically using `CLASS`.

Additionally, the non-linear integrated effect of RSD (usually referred to as the Finger-of-God effect), as well as the intrinsic redshift uncertainty of the galaxy, can be encoded as a multiplicative damping term[55]

$$\begin{aligned} \tilde{P}_g(k, z, \mu) &= P_g(k, z, \mu) \exp\left[-\frac{\mu^2 k^2 \sigma_v^2(z)}{H^2(z)}\right], \\ \text{with} \quad \sigma_v &= (1+z)\sqrt{\sigma_0^2 + \frac{\sigma_{\rm fog}^2(1+z)}{2}}, \end{aligned} \tag{2.11}$$

where $\sigma_0 = 10^{-3}$ accounts for the resolution limits of DESI[16] and *Euclid*[28], and $\sigma_{\rm fog}$ is related to the velocities of galaxies and is treated as a nuisance parameter in this work.

The physics of galaxy formation is known to produce an additional scale-dependence to the bias term, proportional to k^2 at high k. While unrelated to the effect of neutrinos on the LSS, this k^2-term is properly accounted for here following the formulation of Ref.[149]. We follow their prescription in writing the total Eulerian bias as

$$b_1(k, z) = 1 + b^L(k, z) + \alpha_2 k^2, \tag{2.12}$$

where α_2 is a free parameter that we marginalize over. At large scales ($k \to 0$), where both the growth-induced and k^2-terms are negligible, the bias is scale-independent and we choose a fiducial value that matches the simulations of each specific tracer. For the DESI survey of ELGs we use the parametrization $b_1(z) = \beta_0/D(z)$, where $D(z)$ is the growth function, and for the *Euclid* survey of Hα emitters we use $b_1(z) = \beta_0(1 + z)^{0.5\beta_1}$, following the prescriptions of their respective Science Books[16,28]. The nuisance parameter β_0 rescales the overall bias, and β_1 parametrizes any uncertainty in redshift dependence of the bias. We will take as fiducial $\beta_{0,\mathrm{DESI}} = 1.0,\ \beta_{0,Euclid} = 1.7,\ \beta_1 = 1.0$, consistent with recent results from simulations[63].

Figure 2.3 shows the percent differences in $\tilde{P}_g(k, z, \mu)$ with respect to a fiducial scenario of inverted neutrino hierarchy with $\sum m_\nu = 100$ meV, upon changing the hierarchy (while fixing the total neutrino mass) and switching off the GISDB. Other cosmological parameters are held fixed, and the shaded regions represent the shot noises expected from DESI and *Euclid*. This figure shows that the effect of the neutrino hierarchies, as well as the GISDB, is at the 0.5% level. However, while the hierarchies affect the ratio differently at different scales, the GISDB acts as an overall change in normalization at $k \gtrsim 10^{-3}\ h/\mathrm{Mpc}$.

We also account for the Alcock-Paczynski effect[24,33,128], which concerns the cosmology-dependence of inferring distance from angular- and redshift-space measurements, by multiplying the power

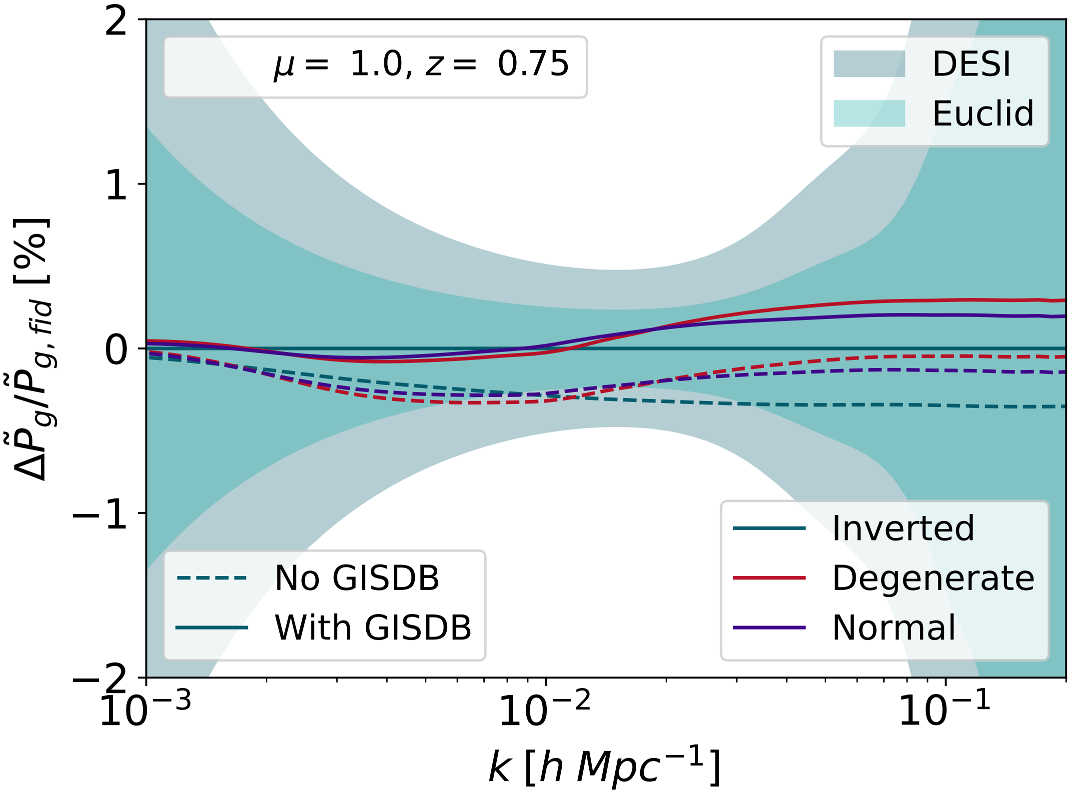

Figure 2.3: Percent differences in galaxy power spectra $\tilde{P}_g(k, z, \mu)$ between the various neutrino hierarchies (at fixed $\sum m_\nu = 100$ meV), as well as with and without the GISDB, compared to a fiducial case of inverted hierarchy with GISDB, at $z = 0.75$. The shot noises associated with DESI and *Euclid* are shown as the shaded areas. Here the cosmological parameters $\{\omega_b, \omega_{\rm cdm}, h, A_s, n_s, \sum m_\nu\}$ as well as all bias and RSD nuisance parameters are held fixed.

spectrum and the shot noise at each bin by $H(z)/D_A^2(z)$, dividing by the same quantity evaluated at our fiducial cosmology, as well as writing the inferred k, μ with respect to the fiducial k, μ by the relation

$$\frac{k(z, \mu_{\rm fid})}{k_{\rm fid}} = \left[(1-\mu_{\rm fid}^2)\frac{D_{A,\rm fid}^2(z)}{D_A^2(z)} + \mu_{\rm fid}^2 \frac{H^2(z)}{H_{\rm fid}^2(z)}\right]^{1/2}, \tag{2.13}$$

$$\frac{\mu(z)}{\mu_{\rm fid}} = \left[(1-\mu_{\rm fid}^2)\frac{D_{A,\rm fid}^2(z)}{D_A^2(z)}\frac{H_{\rm fid}^2(z)}{H^2(z)} + \mu_{\rm fid}^2\right]^{-1/2}. \tag{2.14}$$

Additionally, we stop our analysis at $k_{\max}(z) = 0.2 \times (1+z)^{2/(2+n_s)} h\,\mathrm{Mpc}^{-1}$, safely within the linear regime.

We compute the log-likelihood as [191]

$$-\log\mathcal{L} = \frac{1}{2}\sum_z \int \mathrm{d}k_{\mathrm{fid}} k_{\mathrm{fid}}^2 \int \mathrm{d}\mu_{\mathrm{fid}} \frac{V_{\mathrm{fid}}(z)}{2(2\pi)^2} \left[\frac{\frac{H(z)}{D_A^2(z)}\tilde{P}_g(k,z,\mu) - \frac{H_{\mathrm{fid}}(z)}{D_{A,\mathrm{fid}}^2(z)}\tilde{P}_{g,\mathrm{fid}}(k_{\mathrm{fid}},z,\mu_{\mathrm{fid}})}{\frac{H(z)}{D_A^2(z)}\tilde{P}_g(k,z,\mu) + \frac{H_{\mathrm{fid}}(z)}{D_{A,\mathrm{fid}}^2}n_g^{-1}(z)}\right]^2 . \tag{2.15}$$

Additional parameters, such as those accounting for non-Poissonian shot noise or theoretical error in this likelihood function, can be considered for more detailed analyses, as in e.g. Ref. [52].

2.3.2 CMB Data

We complement the LSS information from galaxy surveys with mock CMB data from the upcoming CMB-S4 experiment, implemented with `MontePython`'s `Likelihood_mock_cmb`. We model the CMB-S4 simply as a single effective frequency channel, with temperature noise $\Delta_T = 1\,\mu$K-arcmin, and polarization noise $\Delta_P = \sqrt{2}\Delta_T$. We additionally assume a resolution of $\theta_{\mathrm{FWHM}} = 3$ arcmin.

CMB data will not only help break the degeneracies between cosmological parameters, but can also measure the matter power spectrum directly through CMB lensing. This data will, for instance, break the degeneracy between the Hubble parameter and $\sum m_\nu$. We perform iterative delensing as in Ref. [106,160] to obtain the deflection field with nearly optimal noise. Finally, we do not account for modes below $\ell = 30$ from CMB-S4, as it will be ground-based, and instead add a Gaussian prior on τ_{reio} with a width of 0.01 to account for low-ℓ CMB data. This width is reflective of current *Planck* sensitivities [17] but conservative in light of future measurements.

2.4 RESULTS

In this section we perform Markov Chain Monte Carlo (MCMC) analyses on mock data for CMB-S4 added to either DESI or *Euclid* galaxy power spectra, for cosmologies with massive neutrinos. We consider parameters $\{\omega_b, \omega_{\rm cdm}, h, A_s, n_s, \tau_{\rm reio}, \sum m_\nu\}$ as well as nuisance parameters $\{\beta_0, \beta_1, \alpha_{k2}, \sigma_{\rm fog}\}$, and show our fiducial values in Table 2.2. We vary our model in three types of ways. First, we attempt to distinguish between the normal, inverted, and degenerate hierarchies. Second, we study if the recovered parameters would be shifted if the GISDB was not included, both with and without marginalizing over the redshift-dependence uncertainty β_1. Additionally, we omit the $\tau_{\rm reio}$ prior in one case with CMB-S4+*Euclid*, to investigate the importance of additional optical depth information in the presence of the existing CMB lensing and LSS shape information.

These different runs are designed to explore the set of physical effects that will appear in upcoming measurements of neutrino masses. A table of relevant reconstructed parameters and associated best-fit log-likelihoods for selected models is shown in Table 2.3 for runs with *Euclid* data, and in Table 2.4 for those with DESI. We emphasize that in all cases our fiducial model corresponds to the inverted hierarchy with its lightest neutrino taken as massless, and thus represents a plausible model of nature; shifts from the best-fit in other models can be seen as the expected shift one would observe in a realistic analysis.

We find that with CMB-S4+DESI data the total neutrino mass is expected to be measured up to uncertainty of 26 meV, while for CMB-S4+*Euclid* data that would be improved to 20 meV. For $\sum m_\nu = 98$ meV this corresponds to a 4σ and 5σ detection respectively, and the *minimum-mass* scenarios of normal and inverted hierarchy (with total masses 60 and 100 meV) can be distinguished at the 1.5σ and 2σ level. We note that neglecting the nuisance parameter β_1 results in an over-tightening of $\sum m_\nu$ resolution to an uncertainty of 15 meV. Finally, the omission of a $\tau_{\rm reio}$ prior results in a $\sum m_\nu$ uncertainty of 27 meV for the combined CMB-S4+*Euclid* data; equivalently showing that these data are able to measure $\tau_{\rm reio}$ to the 10% level despite the lack of low-ℓ information. Conversely, we find that a stricter $\tau_{\rm reio}$ prior of width 0.006 would tighten e.g. CMB-S4+DESI sensitivities to 20 meV.

Cosmological Parameters							
ω_b	$\omega_{\rm cdm}$	h	A_s	n_s	$\tau_{\rm reio}$	$\sum m_\nu$ [meV]	Hierarchy
2.226e-2	0.1127	0.701	2.23221e-9	0.967	0.0598	98.5	Inverted

Nuisance Parameters				
Euclid		DESI		
β_0	β_1	β_0	α_2	$\sigma_{\rm fog}$ [km/s]
1.7	1.0	1.0	1.0	250

Table 2.2: Fiducial cosmology used in generation of mock data for MCMC analysis. Consistent cosmologies are used for DESI and *Euclid* analyses except for β_0, which is matched to simulation results.

2.4.1 Differentiation of Hierarchy

We first consider whether we can differentiate the neutrino hierarchies, if they had the same $\sum m_\nu$. In Figure 2.4c we show a corner plot comparing posteriors for the three hierarchies (the two physical ones plus the degenerate one), where the underlying fiducial cosmology is IH. Due to the physical lower bounds for the total mass of neutrinos in the IH, the posteriors for that case are notably one-sided, and as a result the mean value of most cosmological parameters for the other two hierarchies are shifted relative to the inverted one. This is because all cosmological parameters other than ω_b exhibit significant degeneracy with $\sum m_\nu$. However, as Table 2.3 shows, these near-future surveys show at most a 1σ preference for the fiducial choice of hierarchy, as far as best-fit likelihoods are concerned. This agrees with the recent Bayesian analysis done in Ref.[30].

Data		Model			Mean and error				
LSS	CMB	H	Nuis.	GISDB	$-2\Delta\log\mathcal{L}$	$\mathcal{M}_\nu$	τ_{reio}	β_0	β_1
Euc.	S4 + τ	D	β_0, β_1	Yes	1.3	103.6	5.85e-2	1.702	1.005
						(20.1)	(5.96e-3)	(2.97e-3)	(3.08e-3)
				No	1.3	104.2	5.97e-2	1.704	1.003
						(21.9)	(6.47e-3)	(3.14e-3)	(3.24e-3)
			β_0	Yes	1.5	102.8	5.93e-2	1.699	-
						(16.5)	(5.124e-3)	(2.71e-3)	
				No	1.9	114.5	6.25e-2	1.707	-
						(15.6)	(4.96e-3)	(2.59e-3)	
		I	β_0, β_1	Yes	0.0	113.0	6.30e-02	1.700	9.99e-1
						$\left(^{+9.06*}_{-0.72}\right)$	(3.34e-03)	(3.07e-03)	(2.64e-03)
		N	β_0, β_1	Yes	0.9	98.90	5.89e-2	1.701	1.00
						(21.3)	(6.18e-3)	(3.13e-3)	(3.09e-3)
	S4	D	β_0, β_1	Yes	1.3	102.9	5.95e-2	1.699	1.001
						(27.5)	(8.29e-3)	(3.31e-3)	(2.94e-3)

Table 2.3: Comparison of reconstructed mean and error of cosmological and nuisance parameters as well as best-fit log-likelihoods with respect to the fiducial for different models, with *Euclid* mock data. The fiducial for all these cases is the same and is given in Table 2.2, which is exactly recovered by the inverted model tabulated here. The Σm_ν posterior of this model (denoted by an asterisk), is truncated by the prior at its minimum mass and thus is narrower than its normal and degenerate hierarchy counterparts. We use the following symbols: Euclid ("*Euc.*"), CMB-S4 ("S4"), τ_{reio} ("τ"), Hierarchy ("H"), Degenerate ("D"), Inverted ("I"), Normal ("N"), Nuisance Parameters ("Nuis."), Σm_ν [meV] ("$\mathcal{M}_\nu$"). We denote errors on reported values by parentheses.

2.4.2 EFFECT OF GISDB

We then look more in detail into the runs with and without the GISDB. In this case the effects are twofold: a scale-dependent step that counteracts the scale-dependent suppression induced by neutrinos, and a redshift-dependent amplitude of the step that enhances the redshift-dependence induced by neutrinos at the smallest scales. Omission of the GISDB in the analysis analysis is then expected to underpredict $\sum m_\nu$ if the former effect is dominant, and overpredict if the latter effect is. However, the former effect is largely rendered insignificant due to cosmic variance – the scale-dependence of the bias plateaus at scales smaller than $\mathcal{O}(10^{-2} h\mathrm{Mpc}^{-1})$, which is the regime with strongest statistical power. Thus, it is the latter small-scale redshift-dependent effect that becomes

Data		Model			Mean and error			
LSS	CMB	H	Nuis.	GISDB	$-2\Delta \log \mathcal{L}$	M_ν	$\tau_{\rm reio}$	β_0
DESI	S4 + τ	D	β_0	Yes	0.9	107.6	5.99e-2	1.000
						(26.7)	(7.20e-3)	(1.70e-3)
				No	1.1	112.0	6.07e-2	1.003
						(26.1)	(6.93e-3)	(1.73e-3)
		I	β_0	Yes	0.0	107.2	6.16e-2	1.001
						$\left(^{+15.2*}_{-0.42}\right)$	(3.84e-3)	(1.63e-3)
		N	β_0	Yes	1.0	99.7	5.89e-2	1.000
						(28.6)	(6.52e-3)	(1.68e-3)
	Planck†	D	β_0	Yes	-	(27.44)	(8.99e-3)	(7.62e-3)

Table 2.4: Similar to Table 2.3, with DESI used as LSS data. The fiducial for all these cases is the same and given in Table 2.2. Note that as before, the Σm_ν posterior of the inverted model (denoted by an asterisk) is prior-informed and thus narrower than those of other hierarchies. In addition, the cases with *Planck* CMB data (denoted by a dagger) are reported as Fisher forecasts only, without an MCMC analysis. Symbols are defined equivalently as in Table 2.2.

relevant.

The most evident effect of neglecting this growth-induced step is a misreconstruction of the normalization bias, resulting in a shift towards larger $\bar{b}_L(z)$, and in turn the incorrect values of β_0 and its redshift dependence β_1, as seen in Figure 2.5. If the nuisance parameter β_1 was not marginalized over, this would further result in a significant shift of reconstructed cosmological parameters such as A_s, h, and $\tau_{\rm reio}$, notably overestimating the total neutrino mass by $\gtrsim 1\sigma$, as expected. This effect is shown in Figure 2.4a, where we show the ellipses with and without the GISDB, although as opposed to Figure 2.5 we did not marginalize over β_1. As the DESI bias prescription does not include a degree of freedom to vary the redshift dependence, neglecting the GISDB induces a shift in the aforementioned cosmological parameters, albeit at the sub-σ level, due to the smaller signal-to-noise. We illustrate this point in Figure 2.4b. We note that the same shifts can be recovered using a simpler Fisher-matrix formalism, as we show in our companion paper[75].

In general, we expect the DESI and *Euclid* prescriptions for analysis of real data, when collected, to be more sophisticated than those presented in the Science Books. Nonetheless, this is additional reason for the analyses of upcoming e.g. DESI data to marginalize over the redshift dependence of the bias, or parametrize the bias at each redshift bin independently, to avoid cosmological-parameter shifts due to the neutrino GISDB. Note, however, that even when marginalizing over β_1 there is a leftover shift on the scale-independent bias β_0, as seen in Figure 2.5. The value of this parameter affects other observations, such as galaxy high-order functions and cross correlations with other data sets, so if one requires an unbiased estimate of β_0 the full GISDB ought to be accounted for.

2.5 CONCLUSIONS

In this work we presented forecasts on the ability of current and upcoming CMB and LSS experiments to measure neutrino masses, both in total and individually. We included all known linear effects induced by neutrinos in the treatment of galaxy survey data, specifically assessing the impact of the scale-dependent bias induced by the effect of neutrinos in the growth function. We also investigated the effect and detectability of realistic neutrino hierarchies in the analysis of these survey data.

Starting with the different neutrino hierarchies, we have shown that for the data considered, the total neutrino mass is determined up to an uncertainty of $\sigma \sum m_\nu \sim 20$ meV at a fiducial of $\sum m_\nu = 98$ meV, the minimum-mass scenario of the inverted hierarchy. While this is a 5σ detection away from 0, the minimum-mass scenario of the normal hierarchy ($\sum m_\nu = 60$ meV) is excluded only at the 2σ level. Furthermore, for a fixed total neutrino mass of 98 meV, a different choice of hierarchy constitutes a difference of $\sim 0.3\%$ in the power spectrum amplitude, and we show that this is expected to result in a 1σ shift in inferred cosmological parameters. As such, more advanced surveys are necessary to definitively distinguish between the two hierarchies, particularly in the case that the total neutrino mass is the same.

As for the growth-induced scale-dependent bias (GISDB), we find that, while cosmic variance limitations render the data insensitive to the scale-dependent shape of the halo bias for allowable neutrino masses, upcoming surveys are expected to be highly sensitive to the redshift dependence of the GISDB at small scales. In joint analyses of CMB-S4 data with large-scale surveys such as DESI or *Euclid*, not including this GISDB step can result in a $\mathcal{O}(1\sigma)$ over-prediction of total neutrino masses, as well as similarly shifted reconstructions for degenerate parameters such as h and $\tau_{\rm reio}$. If one marginalizes over the redshift dependence of the bias, these shifts can be removed, although the resulting analysis will retain a $\mathcal{O}(1\sigma)$ shift in the magnitude and redshift dependence of the Eulerian bias of the relevant tracers. Thus it is imperative to include these effects in order to accurately recover galaxy and cosmology parameters simultaneously, or otherwise marginalize over bias redshift dependence if the latter is prioritized.

[h]

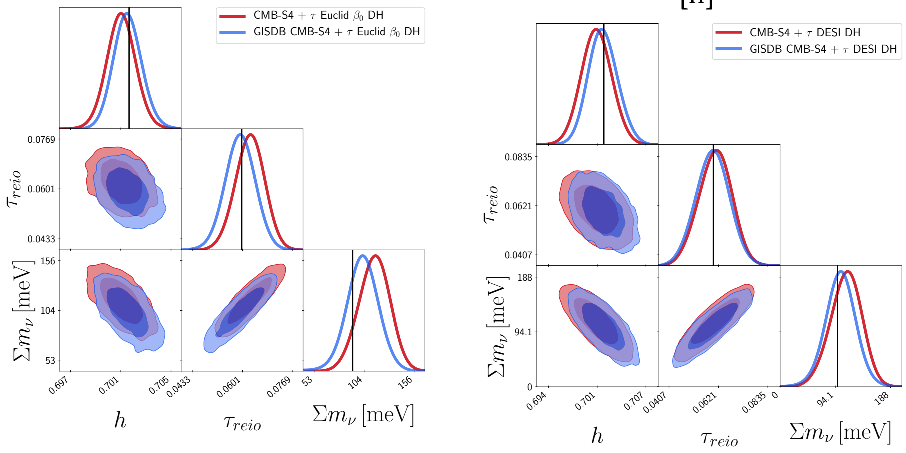

(a) Posteriors for CMB-S4+*Euclid* with a prior on τ_{reio} with a width of 0.01, fixing nuisance parameter β_1, assuming degenerate hierarchy (DH) with and without GISDB. Not marginalizing over β_1 produces a shift in key cosmological parameters due to the GISDB.

(b) Posteriors for CMB-S4+DESI with a prior on τ_{reio} with a width of 0.01, assuming degenerate hierarchy (DH) with and without GISDB. The DESI prescription of not marginalizing over the redshift dependence of the bias gives rise to small parameters shifts.

[h]

[h]

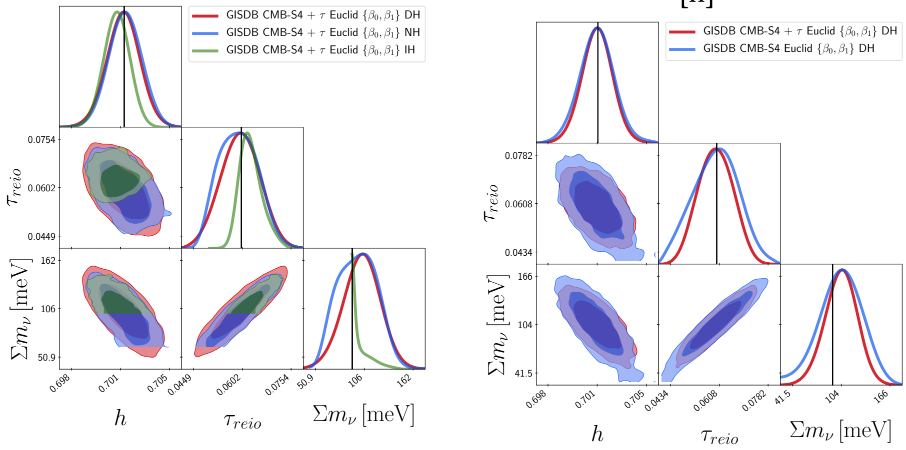

(c) Posteriors for CMB-S4+*Euclid* with a prior on τ_{reio} with a width of 0.01, assuming Degenerate (DH), Inverted (IH) and Normal (NH) hierarchies.

(d) Posteriors for CMB-S4+*Euclid*, assuming Degenerate (DH), with and without a prior on τ_{reio}.

Figure 2.4: MCMC Posteriors for h, τ_{reio}, and $\sum m_\nu$ for CMB-S4 and various LSS experiments. For each, the fiducial cosmology has an Inverted hierarchy in the minimum-mass configuration, with total neutrino mass 98 meV.

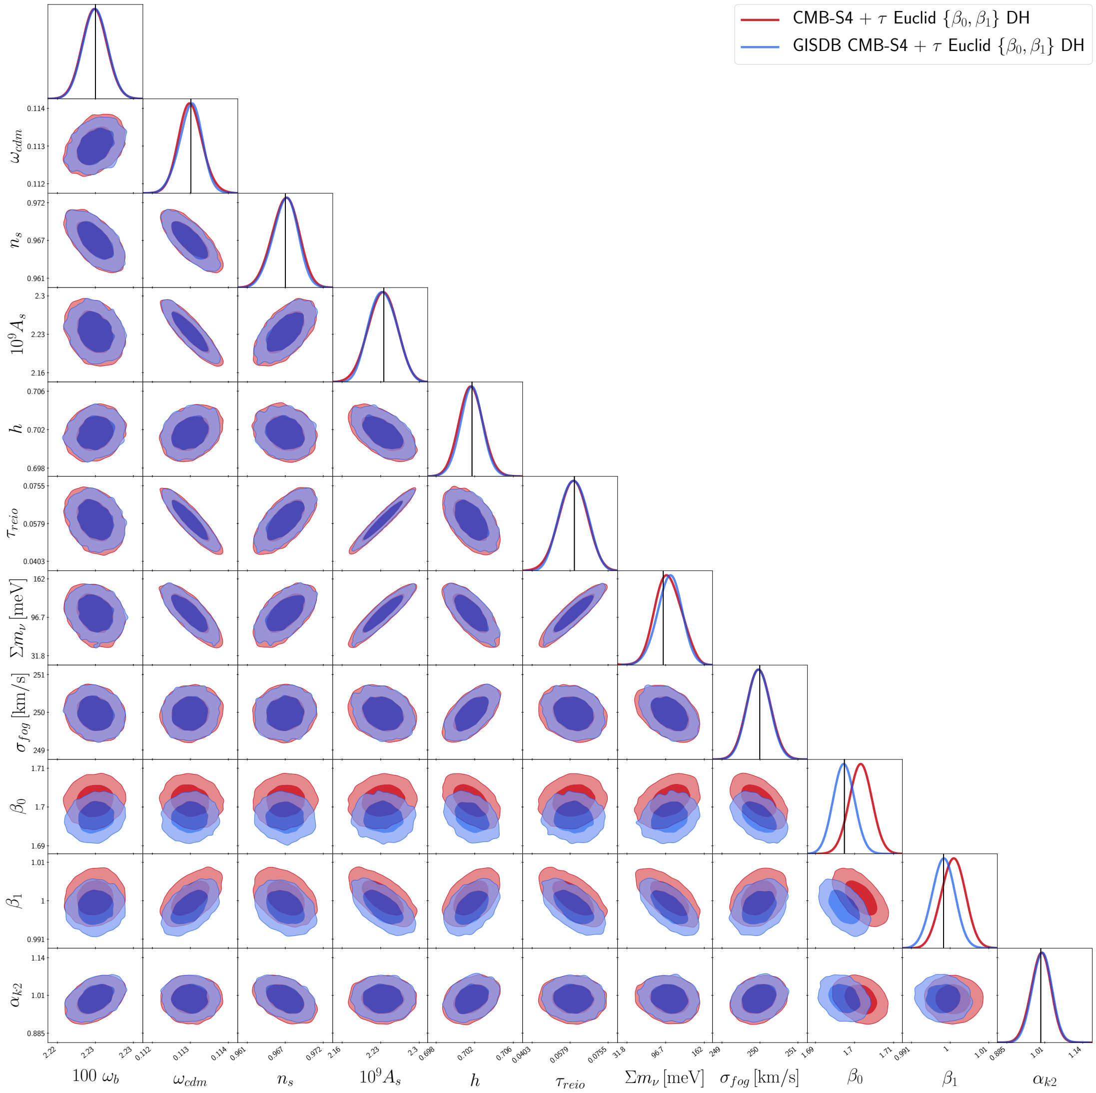

Figure 2.5: Posteriors for CMB-S4+*Euclid* with a prior on $\tau_{\rm reio}$ with a width of 0.01, assuming degenerate hierarchy (DH) with and without GISDB, and assuming the IH with total neutrino mass 98 meV as fiducial.

3

Cosmologically Measuring Ultralight Axion Dark Matter

3.1 Introduction

Axions and, more generally, axion-like particles (ALPs) constitute an intriguing possibility for the particle makeup of dark matter (DM), or at least a subcomponent thereof. Motivated in the QCD sector as a solution to the Strong-CP problem, and in the general case as a consequence of gravitationally complete models of quantum mechanics [202,32], these particles may also produce phenomenologically interesting deviations from the standard picture of cold dark matter cosmology (ΛCDM). This is because independent of any couplings assumed under a specific particle model, the axion or ALP (henceforth, simply "axion") has a purely gravitational signature that can be detected cosmologically.

Ultra-light and extremely feebly coupled axions can acquire a matter abundance via vacuum misalignment, where the particle can be thought of as an initially slowly-rolling field that transitions its equation of state when it begins oscillating at the minimum of its potential. This transition occurs when the Hubble rate is of order the axion mass $H(a_{\rm osc}) \approx m_\varphi$, occurring at later redshifts for lighter masses. For an axion produced via the misalignment mechanism, the relic abundance today $\omega_{\varphi,0}$ is determined by the axion mass m_φ and the initial "misaligned" field value $\varphi_{\rm I}$. For axions with $m_\varphi \gtrsim 10^{-33}$ eV, this transition has occurred before today, and its relic abundance will contribute to the presently observed DM abundance. Additionally, ultra-light axions impede the growth of perturbations at small scales through an effect which has been called "quantum pressure" as a result of their macroscopic de Broglie wavelengths, $\sim$ 10 Mpc at $m_\varphi \sim 10^{-30}$ eV. These two effects are known to change the ratios of cosmological species and the rate of universe expansion throughout cosmic history, as well as suppress the growth of matter perturbations at small scales relative to a ΛCDM cosmology [114].

Upcoming Cosmic Microwave Background (CMB) observations and galaxy surveys will offer the ability to interrogate these processes at a much more detailed level. In order to exploit these substantially larger data sets to search for axions, its effects on structure formation and other cosmological observables must be modeled to adequate precision. While the effect of ultra-light axions on the linear theory of structure formation is well understood, less extensively studied is how these particles influence the non-linear process of structure formation at late times. The halo bias is a key parameter in modeling this late-time physics, quantifying to leading order the relationship between galaxy density perturbations, which we directly observe, and the underlying (predominantly dark) matter density perturbations, which we would like to learn about. While previous work [74] has demonstrated how axions affect the *redshift* dependence of the halo bias, this work presents the first consideration of how ultra-light axions affect the *scale* dependence of the halo bias.

In this work we will specifically consider ultra-light axion dark matter with masses from $m_\varphi = 10^{-20}$ eV to $m_\varphi \approx 10^{-33}$ eV. While other thermal and non-thermal axion production channels are possible [105], we will only consider the vacuum misalignment production mechanism. The high end of this mass range is best constrained by small-scale Lyman-α measurements, which rule out an ultra-light axion composing all of the DM for masses below 2×10^{-20}eV [182,183]. For masses below 10^{-24} eV, measurements of the CMB and large-scale structure (LSS) constrains the axion to only comprise a percent level fraction of the total CDM [108]. Combined CMB and galaxy-clustering analyses set the lowest bounds to date on the allowed fractional abundance of ultra-light axion dark matter, with the most stringent constraint set at 1.4% for a 10^{-28} eV axion [*] [107,170,125,114,27].

There are also a number of complementary and similarly model-independent analyses in the literature that have leveraged other data sets to constrain axion dark matter, such as galaxy cluster counts [79]; probes using 21cm tomography [111,32,132]; line intensity mapping [39] ; weak-lensing shear [74]; kinetic Sunyaev-Zeldovich effects [85]; and gravitational waves [99]. If we also assume some axion coupling to photons, there are a number of astrophysical and laboratory data sets that can be utilized to constrain the properties of the axion [142,14]. Similarly, model-dependent cosmological probes include modifications to the Lyman-α effective opacity [185] and rotations of the CMB polarization angle [32].

Nevertheless, galaxy surveys in particular provide powerful information on the large scale struc-

---

[*] It should be noted that in all these cases, a model with only a single axion is considered. It is plausible that these constraints would change in models containing a spectrum of axions.

ture of the universe, and a precise map of way dark matter fluctuations grow and cluster may well be our only path to discovering these elusive axion populations. Modeling the halo bias accurately is consequently critical to extracting this information faithfully, in order to match the increasingly precise experimental data promised by near-future surveys. To facilitate this work, we develop a modified version of the `RelicFast`[†] code, which we call `RelAxiFast`[‡]. In its un-modified form, `RelicFast` is able to quickly compute the halo bias to linear order in the presence of thermal relics, including massive neutrinos. `RelAxiFast` extends the functionality of `RelicFast` to additionally compute the linear halo bias in the presence of an ultra-light axion with time varying equation of state. We make `RelAxiFast` publicly available for use, including scripts to reproduce all results presented in this work.

We review the process of how an ultra-light axion affects the matter power spectrum in §3.2, beginning with a phenomenological description of the axion in §3.2.1, and then showing how the axion field dynamics effect the background evolution §3.2.2 and evolution of perturbations at different scales §3.2.3. We describe the code `axionCAMB` in §3.2.4, which we use to compute the background and perturbation evolution. We go on to describe how we model the axion in the non-linear dynamics of halo collapse in §3.3. We first describe modifications to the dynamics of halo collapse in §3.3.1 and present our calculation of the halo bias in §3.3.3. We describe the `RelAxiFast` code which we use to evaluate the halo collapse process in §3.3.2. We present our results and conclude in §3.4.

[†] github.com/JulianBMunoz/RelicFast

[‡] github.com/ndeporzio/RelAxiFast

3.2 AXION COSMOLOGY

3.2.1 AXION PHENOMENOLOGY

The QCD sector of the SM allows the incorporation of a CP-violating interaction, but this term, while generically expected to be sizable, is experimentally constrained to be exceedingly small[35]. An elegant solution to this "Strong CP-Problem" is found in the introduction of a periodic scalar known as the QCD axion, which dynamically sets the size of the CP-violation. This sets the scales at which the QCD axion is expected to interact, and correspondingly bounds its mass at $m_\varphi \gtrsim 6 \times 10^{-6}$ eV, so as to avoid overclosing the Universe[193,167,202,203,172,80,5].

However, periodic axion-*like* particles with analogous couplings to gauge fields are known to arise in various string theory scenarios; for example, as Kaluza-Klein zero modes of antisymmetric tensor fields under manifold compactification[204,101,192]. These particles are not required to solve the Strong CP problem, and as a result, ALPs have a less-constrained parameter space compared to the QCD axion and broader phenomenological possibilities. Throughout this work, we will take the "axion" to simply mean any axion-like scalar with a periodic potential.

We assume an axion φ has an effective Lagrangian of the form

$$\mathcal{L} = -\frac{1}{2}(\partial\varphi)^2 - \Lambda^4 U(\varphi/f), \tag{3.1}$$

where $U(\varphi/f)$ can be any dimensionless potential minimized at $\varphi = 0$ and periodic under $\varphi \to \varphi + 2\pi f$. f is known as the decay constant of the field, and Λ is the symmetry breaking scale when mapped onto a low energy theory.

We apply a sinusoidal approximation to the true potential,

$$U\left(\frac{\varphi}{f}\right) \approx 1 - \cos\left(\frac{\varphi}{f}\right), \tag{3.2}$$

Phenomenologically, choices of potentials with an exponent greater than 1 have been considered, for example, in the context of early dark energy (EDE) models[171,190,18]; however, axions in these scenarios ultimately do not behave as CDM at late times. In this work, we are interested in axions that may contribute to the relic CDM density, and so restrict the analysis to the potential in Eq. 3.2. The behavior of Eq. 3.1 expanded around the potential minimum is then described by the effective Lagrangian

$$\mathcal{L} \approx -\frac{1}{2}(\partial\varphi)^2 - \frac{1}{2}m_\varphi^2\varphi^2 \,, \tag{3.3}$$

which has canonical mass term

$$m_\varphi \equiv \frac{\Lambda^2}{f} \,. \tag{3.4}$$

determined by the axion decay constant f and symmetry breaking scale Λ.

3.2.2 Background Evolution

The Klein-Gordon equation of motion corresponding to Eq. 3.3 governs the evolution of the homogeneous free-axion field

$$\ddot{\varphi} + 2\mathcal{H}\dot{\varphi} + m_\varphi^2 a^2 \varphi = 0 \,, \tag{3.5}$$

where overdots indicate derivatives with respect to conformal time η, and where the conformal Hubble parameter is $\mathcal{H} \equiv aH$. Notice in Eq. 3.5 the presence of a Hubble friction term that impedes the growth of the axion field. When the φ and $\dot{\varphi}$ terms become similar, the field begins to oscillate around its minimum, occurring when

$$m_\varphi \approx 3H(a) \,, \tag{3.6}$$

with a time scale of oscillations given by

$$\partial t_{\rm osc} \approx m_{\varphi}^{-1}\,. \tag{3.7}$$

The stress-energy tensor for an arbitrary scalar field is given by

$$T^{\alpha}_{\beta} = g^{\alpha\nu}\frac{\partial\varphi}{x^{\nu}}\frac{\partial\varphi}{x^{\beta}} - g^{\alpha}_{\beta}\left(\frac{1}{2}g^{\mu\nu}\frac{\partial\varphi}{\partial x^{\mu}}\frac{\partial\varphi}{\partial x^{\nu}} + U(\varphi)\right)\,. \tag{3.8}$$

Combining the free-axion potential of Eq. 3.3 and the Friedmann-Lemaître-Robertson-Walker metric (FLRW) with Eq. 3.8, the background energy density and pressure are then given by

$$\rho_{\varphi}(a) = -T^{0}_{\ 0} = \frac{1}{2}a^{-2}\dot{\varphi}^{2} + \frac{1}{2}m_{\varphi}^{2}\varphi^{2}\,, \tag{3.9}$$

$$P_{\varphi}(a) = T^{i}_{\ i} = \frac{1}{2}a^{-2}\dot{\varphi}^{2} - \frac{1}{2}m_{\varphi}^{2}\varphi^{2}. \tag{3.10}$$

Given some initial choice of φ and $\dot{\varphi}$, we can use Eq. 3.5, Eq. 3.9, and Eq. 3.10 to calculate the background values of $\rho_{\varphi}(a)$, $P_{\varphi}(a)$ and the equation of state

$$w_{\varphi} \equiv \frac{P_{\varphi}}{\rho_{\varphi}}\,, \tag{3.11}$$

at all points in time. For the axion, we assume slowly rolling initiation conditions, $\dot{\varphi} \approx 0$, so we have $w_{\varphi} \approx -1$ and the axion field evolves like dark energy (DE). Once the axion field begins to oscillate, its energy density redshifts as ordinary matter[196], such that the relic density is approximately given by

$$\omega_{\varphi,0} = h^{2}\Omega_{\varphi,0} = \left(\frac{\rho_{\varphi,\rm osc}}{\rho_{\rm crit}}\right)a_{\rm osc}^{3}, \tag{3.12}$$

where $\rho_{\rm crit} = 3H_0^2M_{\rm P}^2$ is the critical energy density to produce a flat background, $M_{\rm P} \equiv (8\pi G)^{-1/2}$

is the reduced Plack mass, $H_0 \equiv H(a = 1)$ is the value of the Hubble parameter today, and where we have used the definitions: $\omega_i \equiv h^2 \Omega_i$, $\Omega_i \equiv \rho_i / \rho_{\text{crit}}$, and $h \equiv H(a = 1)/100[\text{km/s/Mpc}]$.

This non-thermal production mechanism is known as the vacuum misalignment mechanism - the field being initially misaligned with the minimum of its potential evolves as to realign with, and then oscillate about that minimum.[§] In this work, we only consider the case of axions produced by the misalignment mechanism and with a symmetry breaking scale greater than the scale corresponding to the end of inflation. The effects of other non-thermal production channels, such as production via topological defects formed by symmetry breaking after inflation, are beyond the scope of this work. Likewise, as we have assumed no SM couplings to the axions considered here, any thermal production can be safely ignored.

Throughout this work, we assume a flat background cosmology with $\omega_{\text{b},0} = 0.02226$, $\omega_{\text{d},0} \equiv \omega_{\text{cdm},0} + \omega_{\varphi,0} = 0.1127$, $n_s = 0.96659$, $A_s = 2.2321 \times 10^{-9}$, $N_{\text{eff}} = 3.046$, and including three massless neutrinos. Note that we consider the effects of introducing massive neutrinos in Appendix §3.5. Here, $\omega_{\text{b},0}$, $\omega_{\text{cdm},0}$ and $\omega_{\varphi,0}$ refer to the energy density today of baryons, CDM, and axions, respectively. The background cosmology will determine $H(a)$, and in turn set the time a_{osc} when $H(a)$ falls below $m_\varphi/3$ and the axion field oscillates about the minimum of its potential. Using this cosmology, the oscillation time as a function of axion mass is numerically calculated in Fig. 3.1. Note that these oscillation times are robust to different choices of relic axion abundances between 0% and 10% of $\omega_{\text{d},0}$, down to at least one part in 10^5. During radiation domination, we see that $a_{\text{osc}} \propto m_\varphi^{-1/2}$ and during matter domination $a_{\text{osc}} \propto m_\varphi^{-2/3}$ in accordance with $H(a)$ during those epochs. Whilst disfavored cosmologically[125], larger abundances would produce more substantial shifts to matter-radiation equality and substantially change the $a_{\text{osc}}(m_\varphi)$ relationship shown in Fig. 3.1.

[§] For this reason, the "misalignment mechanism" is sometimes referred to as the "realignment mechanism".

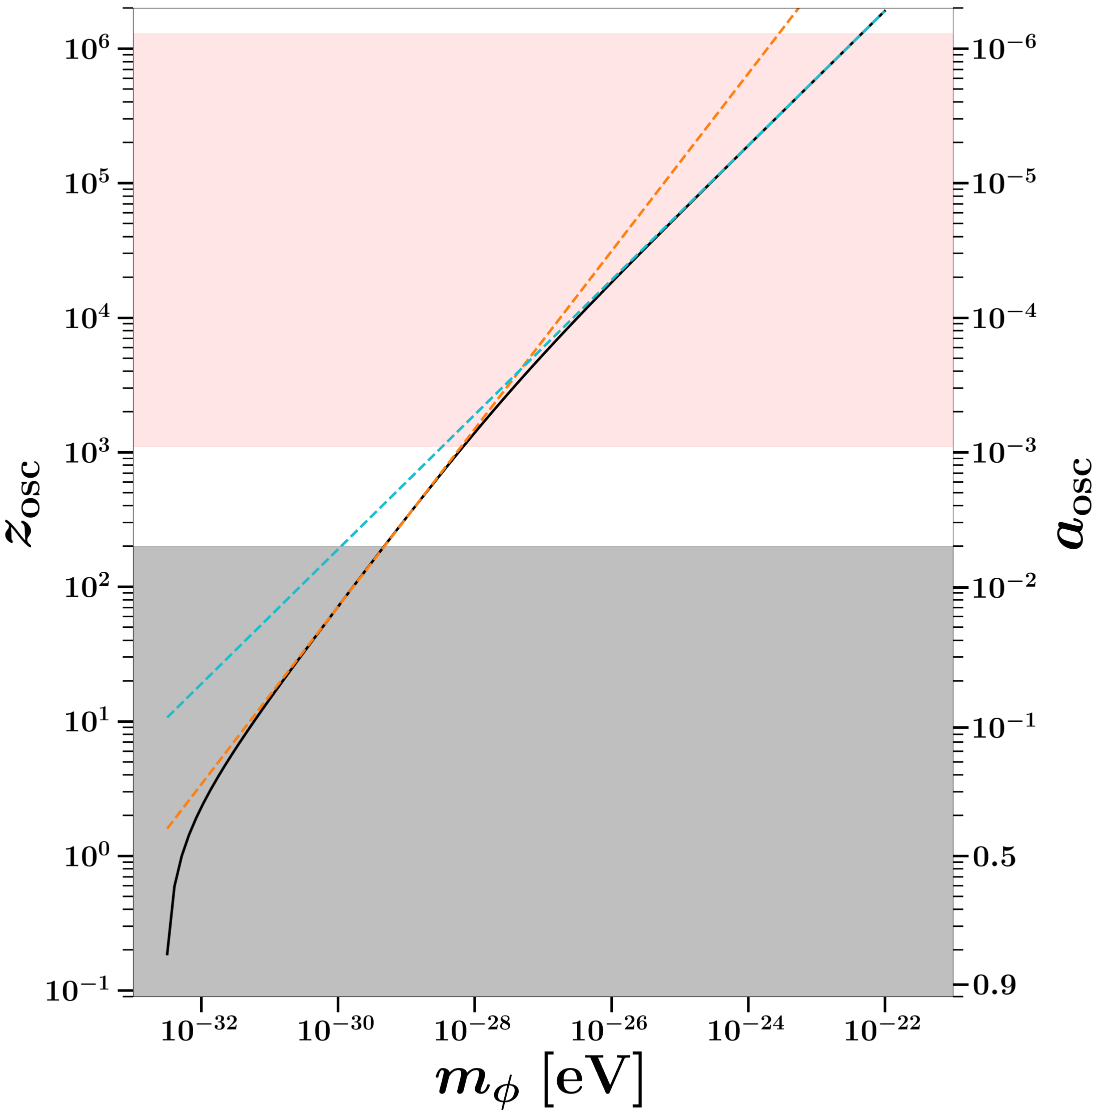

Figure 3.1: The time when the axion field begins to oscillate, $m_\varphi = 3H(a)$, as a function of the axion mass m_φ. The rate of change of the Hubble parameter during matter domination, $\partial \log H / \partial \log a = -3/2$ is shown in orange, and during radiation domination $\partial \log H / \partial \log a = -2$ is shown in blue. The transition between these two regimes occurs at matter-radiation equality. The matching to the matter solution breaks downs at small redshift, when dark energy begins to dominate. The red shaded region is bounded by the redshift of recombination $z \approx 1100$ and the redshift when a CMB mode of $\ell \approx 2500$ crosses the horizon. The grey shaded region is bounded by the largest redshift of halo collapse modeled in this work $z \approx 200$ and today.

Using the $n = 1$ result from [125], we can map the parameters of this ultra-light axion dark matter model (m_φ,f, φ_I) to the cosmological parameters (H_0, a_osc, Ω_φ) as

$$\begin{aligned}\Omega_\varphi(a) &\approx \frac{1}{3}\left(\frac{fm_\varphi}{M_\mathrm{P}H_0}\right)^2 \big(1-\cos\left(\varphi_\mathrm{I}/f\right)\big) \ \times \ \ldots \\ \ldots &\begin{cases}\left(\frac{a_\mathrm{osc}}{a}\right)^3 & , a > a_\mathrm{osc} \\ 1 & , a < a_\mathrm{osc}\,.\end{cases}\end{aligned} \tag{3.13}$$

As shown in Fig. 3.1, the relation between a_osc and m_φ depends on whether the axion field begins to oscillate while the background is radiation dominated or matter dominated. In the case of lighter, sub-dominant axion masses, which begin oscillating during matter domination

$$a_\mathrm{osc} \approx \left(\frac{3H_0}{m_\varphi}\right)^{2/3} (2\Omega_{\mathrm{m},0})^{1/3}\ . \tag{3.14}$$

Alternatively, for heavier, sub-dominant axion masses, which begin oscillating during radiation domination

$$a_\mathrm{osc} \approx \left(\frac{3H_0}{m_\varphi}\right)^{1/2} \left(2\Omega_{\mathrm{m},0}a_\mathrm{eq}\right)^{1/4}\ . \tag{3.15}$$

We can identify which axion mass begins oscillating at the time of matter-radiation equality, again assuming that the axion is a sub-dominant component of the energy density

$$m_{\varphi,\mathrm{eq}} \approx \left(3H_0\sqrt{2\Omega_{\mathrm{m},0}}\right) a_\mathrm{eq}^{-3/2}\ . \tag{3.16}$$

So we can express Eq. 3.13 in terms of the cosmological parameter a_eq instead of the axion parameter

$a_{\rm osc}$

$$\Omega_{\varphi,0} = \mathcal{A} \times \ldots$$

$$\ldots \begin{cases} 6\Omega_{\rm m,0} & , m_\varphi < m_{\varphi,\rm eq} \\ 2^{3/4}\sqrt{\frac{3m_\varphi}{H_0}}\left(\Omega_{\rm m,0}a_{\rm eq}\right)^{3/4} & , m_\varphi > m_{\varphi,\rm eq} \end{cases} \tag{3.17}$$

where we have defined the following dimensionless coefficient which depends only on the axion model parameters

$$\mathcal{A} \equiv \left(\frac{f}{M_{\rm P}}\right)^2 \left(1 - \cos\left(\varphi_{\rm I}/f\right)\right) \ . \tag{3.18}$$

Assuming the axion field always begins in a slowly rolling configuration $\dot{\varphi} = 0$, different choices of the initial value of the field $\varphi_{\rm I}$ will produce different values of the relic axion relic density today $\omega_{\varphi,0}$. An axion described in this manner will have a mass determined by its symmetry breaking scale and decay constant $m_\varphi(f_\varphi, \Lambda_\varphi)$ and a relic abundance determined by all three model parameters $\omega_\varphi(f_\varphi, \Lambda_\varphi, \varphi_i)$. Without a means of independently determining the symmetry breaking scale, we can only assume a lower bound on Λ_φ set by our choice of the Hubble expansion rate during inflation and the assumption the symmetry breaking occurs before the end of inflation.

The cosmological evolution of such an axion will have an oscillation redshift set exclusively by $m_\varphi(\Lambda, f)$, so long as the axion abundance does not substantially change the background evolution of $H(a)$. In Fig. 3.2, for a single choice of m_φ the transition from a DE equation of state to a CDM equation of state always occurs at the same redshift for different choices of $\omega_{\varphi,0}$ (all of which are small compared to $\omega_{\rm CDM}$). Conversely, in Fig. 3.3 the transition occurs at later times for lighter axion masses. Notice that if too light of a mass is chosen, the field will need to have much a much smaller primordial abundance in order to recover reasonable relic abundances today. For this reason, any axion produced through the misalignment mechanism with masses below 10^{-33} eV will comprise a negligible fraction of the energy density and have a similarly negligible effect on the cos-

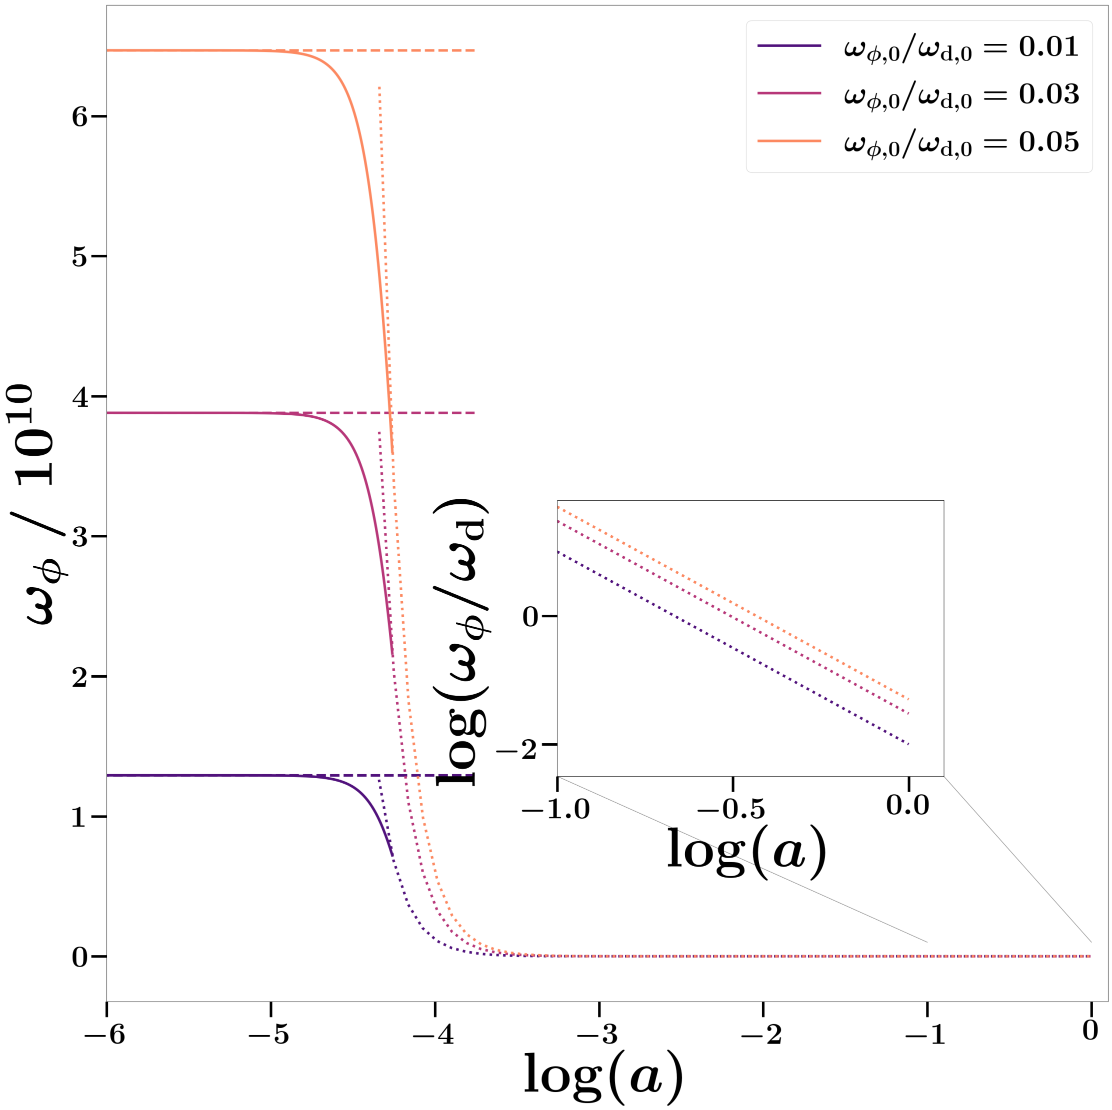

Figure 3.2: The exact evolution of the axion background energy at fixed axion mass $m_\varphi = 10^{-26}$ eV (solid line) compared with a dark energy like evolution with matching early universe abundance (dashed) and dark matter like evolution with matching relic abundance (dotted). The time of the transition from $w_\varphi = -1$ to $w_\varphi = 0$ at a_osc is negligibly changed for different choices of the relic axion abundance. The exact solution is evaluated until $m_\varphi = 3H(a)$, at which point the dark matter like evolution is assumed for the axion field.

mology. In other words, axions with masses below 10^{-33} eV and significant relic abundances cannot be produced by the misalignment mechanism.

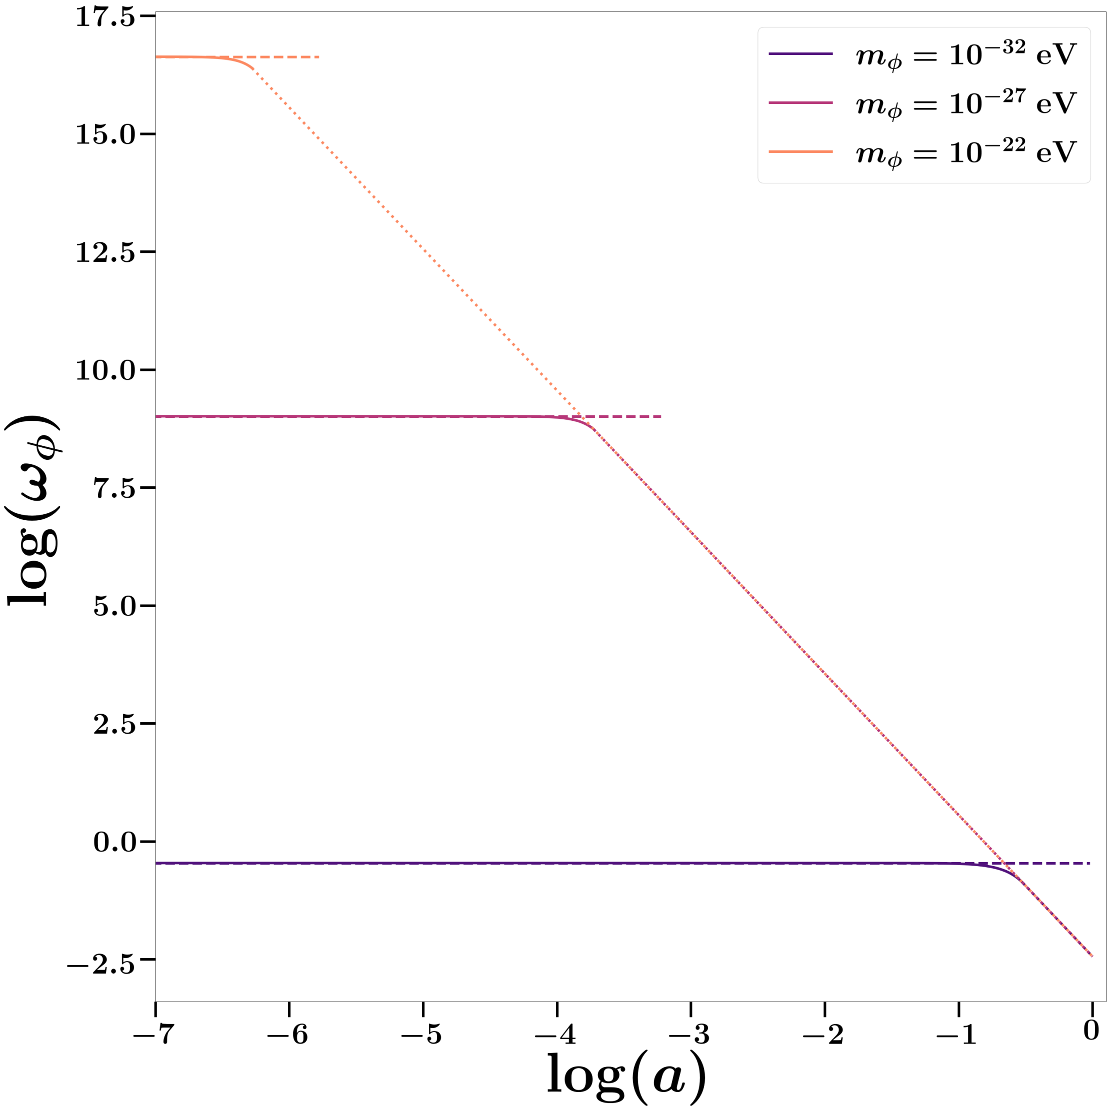

Figure 3.3: The exact evolution of the axion background energy at fixed axion relic abundance ω_φ. Note that a_{osc} occurs at later times for lighter axion masses - H needs more time to evolve to a lower value before lighter fields begin to oscillate. For extremely light masses of $m_\varphi \lesssim 10^{-33}$ eV (the energy scale of H_0), the axion field does not begin oscillating before today and its background energy density is not able to decay to the desired relic abundance.

For the lightest mass in Fig. 3.3, $m_\varphi = 10^{-32}$ eV, notice that the transition of equation of state can occur deep within the redshift regime of halo collapse. This suggests that axions of such masses can leave unique imprints in the halo collapse process, which may have a different sensitivity to the presence and redshift behaviors of the cosmological species than the evolution of the matter perturbations alone.

3.2.3 Perturbation Evolution

To fully utilize cosmological data sets, we are also interested in the evolution of the perturbed (isotropic) axion field around its background value φ_0.

$$\varphi(\tau, k) \equiv \varphi_0(\tau) + \varphi_1(\tau, k) \, . \tag{3.19}$$

In Synchronous gauge, the perturbed field equation of motion corresponding to Eq. 3.5 for the background field becomes

$$\ddot{\varphi}_1 + 2\mathcal{H}\dot{\varphi}_1 + (m_\varphi^2 a^2 + k^2)\varphi_1 = -\frac{1}{2}\dot{\varphi}_0\dot{\beta} \, , \tag{3.20}$$

where β is the trace of the scalar metric perturbation. Likewise, the perturbed stress-energy tensor defines the energy density and pressure of the perturbed field. This system of equations defining the evolution of φ can be expressed equivalently as a generalized dark matter (GDM) fluid with the following equations of motion for the fluid density perturbation δ_φ and velocity perturbation u_φ in Synchronous gauge[113]:

$$\begin{aligned} \dot{\delta}_\varphi = & - k u_\varphi - (1 - w_\varphi)(\dot{\beta}/2) - 3\mathcal{H}(1 - w_\varphi)\delta_\varphi \\ & - 9\mathcal{H}^2(1 - c_{\mathrm{ad}}^2)(u_\varphi/k) \end{aligned} \quad , \tag{3.21}$$

$$\dot{u}_\varphi = 2\mathcal{H}u_\varphi + k\delta_\varphi + 3\mathcal{H}(w_\varphi - c_{\mathrm{ad}}^2)u_\varphi \, , \tag{3.22}$$

where the the adiabatic sound speed is defined in the GDM formalism to be

$$c_{\rm ad}^2 \equiv \frac{\dot{P}}{\dot{\rho}} = w_\varphi - \frac{\dot{w}_\varphi}{3\mathcal{H}(1+w_\varphi)} \,. \tag{3.23}$$

The same physics is determined by either the second order field equation of motion Eq. 3.20 or the set of first order GDM equations of motion Eq. 3.21 and Eq. 3.22. The motivation for working with the GDM equations of motion is to take advantage of approximation schemes developed for the GDM fluid in the fast oscillation regime [170,140,107,114,116,141,166,158,115,163,113]. Recall from Eq. 3.7 that the axion field oscillates on a time scale inversely proportional to its mass m_φ. As H drops below the axion mass, it becomes computationally difficult to evaluate Eq. 3.21 and Eq. 3.22. This is overcome by using the WKB method to find the evolution of the average fluid variables over the oscillation time scale when $a \gg a_{\rm osc}$

$$\dot{\delta}_\varphi = -ku_\varphi - \frac{\dot{\beta}}{2} - 3\mathcal{H}c_{\rm eff}^2\delta_\varphi - 9\mathcal{H}^2c_{\rm eff}^2u_\varphi/k \,, \tag{3.24}$$

$$\dot{u}_\varphi = -\mathcal{H}u_\varphi + c_{\rm eff}^2k\delta_\varphi + 3c_{\rm eff}^2\mathcal{H}u_\varphi \,, \tag{3.25}$$

where the effective sound speed of the fluid is

$$c_{\rm eff}^2 \equiv \frac{\delta P}{\delta \rho} = \frac{(k/k_m)^2}{1+(k/k_m)^2} \,. \tag{3.26}$$

An additional benefit of this formulation is that the effective sound speed $c_{\rm eff}^2$ is a gauge invariant description of perturbation growth, unlike the adiabatic sound speed $c_{\rm ad}^2$, which encodes both spatial and time variations in the equation of state.

To understand the physics described by the effective sound speed of the generalized axion fluid, we first review several relevant scales that are introduced to our cosmology by the presence of a sub-

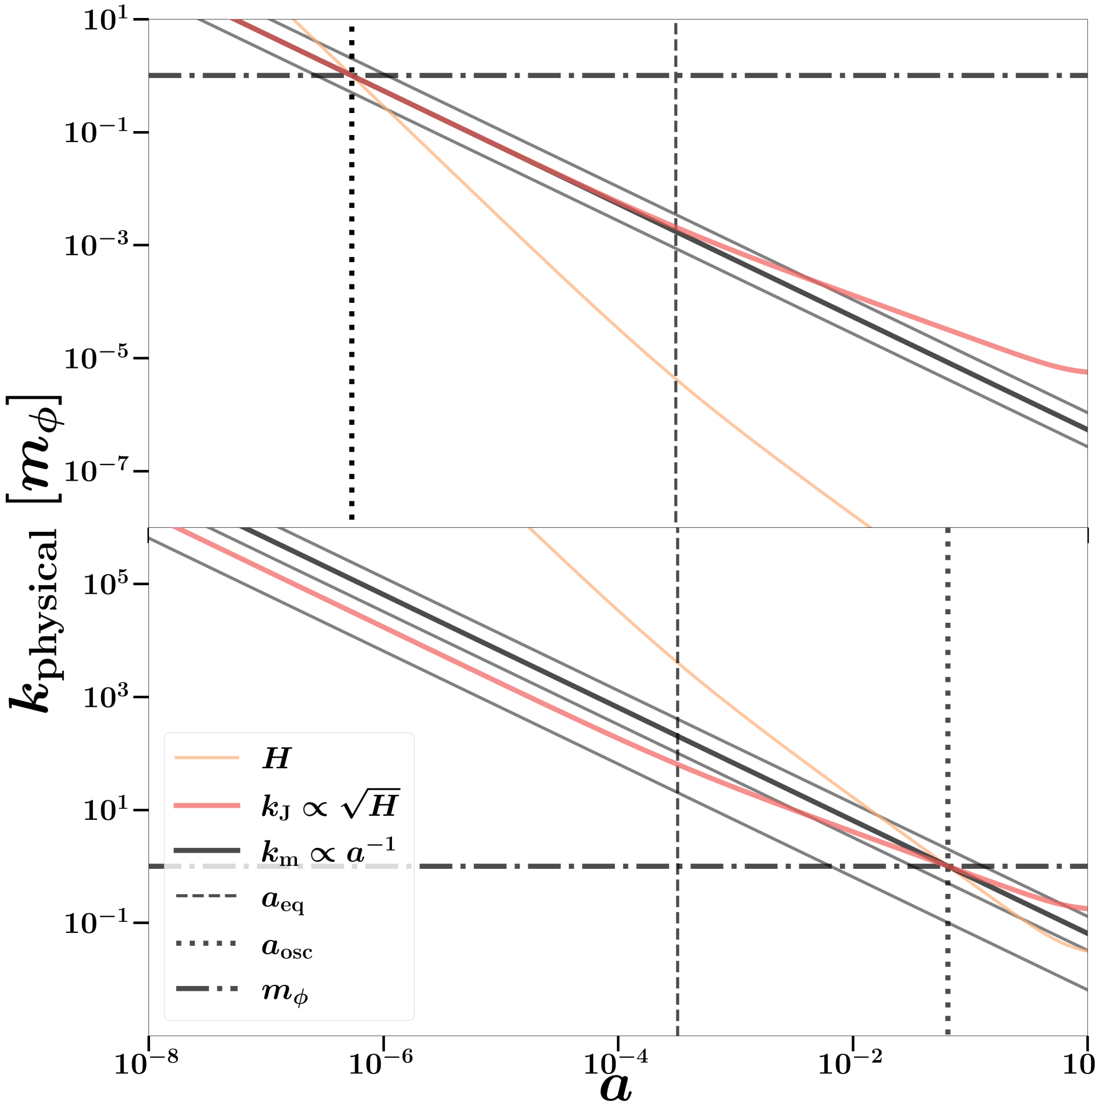

Figure 3.4: Evolution of different physical momentum scales in units of the axion mass as a function of the scale factor for an axion oscillating before matter-radiation equality (top panel, $m_\varphi = 10^{-22}$ eV) and for an axion oscillating after matter-radiation equality (bottom panel, $m_\varphi = 10^{-31}$ eV). Grey lines indicate constant comoving modes, arbitrarily chosen. When the axion field begins to oscillate $a_{\rm osc}$, the Hubble parameter H (we use $m_\varphi = H(a)$ as the oscillation condition here), axion Jeans Scale $k_{\rm J}$ and axion mass scale $k_{\rm m}$ are all equal, by definition. At all times, comoving modes smaller than the axion Jeans scale (grey lines above the $k_{\rm J}$ lines) will be suppressed. Note a slight shift in the location of matter-radiation equality, despite both axions having the same relic abundance.

dominant ultra-light axion. First, the time when a given perturbation mode crosses the horizon $a_{\rm c}$ is set by the evolution of the Hubble parameter. Modes that cross the horizon during different cosmological epochs (e.g. matter domination versus radiation domination) will evolve at different rates. This behavior is largely independent of the physics specific to axions. Each comoving momentum mode k of the axion field will cross the horizon at a different time $a_{\rm c}$ determined by

$$\frac{k(a_{\rm c})}{a_{\rm c}} = H(a_{\rm c})\ . \tag{3.27}$$

However, such a mode can be either relativistic or non-relativistic at the crossing time $a_{\rm c}$ dependent upon the mass of the axion. A given mode will be relativistic in the case where its physical momentum is greater than the axion mass. The time when a mode transitions from being relativistic to being non-relativistic is determined by the condition

$$\frac{k(a_{\rm nr})}{a_{\rm nr}} = m_\varphi\ . \tag{3.28}$$

So we see that a mode which is relativistic at horizon crossing satisfies

$$k_{\rm rel} > m_\varphi a_{\rm c}\ . \tag{3.29}$$

whereas a mode that is non-relativistic at horizon crossing satisfies

$$k_{\rm non-rel} < m_\varphi a_{\rm c}\ . \tag{3.30}$$

Only in the case where a mode is non-relativistic will it contribute to the the CDM perturbations at that particular scale. So we should expect suppression to the matter power spectrum at those modes which are relativistic. In Fig. 3.4, we show two examples of the mode evolution for different axion masses - a heavy mass beginning to oscillate before matter-radiation equality, and a lighter mass beginning to oscillate after. At the level of the background cosmology, the lighter mass will slightly reduce the amount of CDM around the time of matter-radiation equality and thus push $a_{\rm eq}$ to larger values. In Fig. 3.4, this is evidenced by the slight mismatch in the position of the vertical dashed line indicating $a_{\rm eq}$. In each mass case, the axion mass is indicated by the horizontal dot-dashed line. Constant comoving modes are arbitrarily chosen to guide the eye and are indicated by grey lines. In both cases, we see that the largest comoving k mode (highest grey band) first intersects the evolution of the Hubble parameter (orange line) before intersecting the horizontal axion mass line, meaning the mode crosses the horizon while relativistic before becoming non-relativistic. The smallest comoving k mode (lowest grey band) is an example of the opposite - the mode first becomes non-relativistic before crossing the horizon. Those modes which become non-relativistic before crossing the horizon will evolve identically to CDM. We can infer precisely which modes satisfy this condition, and define a corresponding mode $k_{\rm m}$ above which this condition no longer holds. As an example, in the case where we assume the axion begins oscillating well within radiation domination, the time of oscillation is given by

$$a_{\rm osc} \approx \sqrt{\frac{H_0}{m_\varphi}} \left(\Omega_{\rm m} a_{\rm eq}\right)^{1/4} \tag{3.31}$$

so that the modes which are non-relativistic at the time of oscillation are given by

$$k_{\rm non-rel} < k_{\rm m} \equiv m_\varphi a_{\rm osc} = \sqrt{m_\varphi H_0} \left(\Omega_m a_{\rm eq}\right)^{1/4} \ . \tag{3.32}$$

The final relevant scale is the Jeans scale $k_{\rm J}$ - it is the physical scale corresponding to the de Broglie wavelength of axions moving with the Hubble flow. Momentum modes of the axion field become equal with the Jeans momentum of the field when

$$\frac{k_{\mathrm{J}}}{a} \equiv \sqrt{m_\varphi H(a)}\,. \tag{3.33}$$

The Jeans scale characterizes the scale where the axion field pressure matches the competing potential of gravitational collapse, to linear order. Perturbations at length scales greater than the Jeans length will be dominated by gravitational collapse and continue to grow into the non-linear regime and perturbations at smaller lengths will be dominated by the field's quantum pressure and oscillate between collapse and expansion [114]. The physical effect of k_J is encoded in the field (or equivalent GDM fluid) equations of motion, and dynamically manifests the transition in the effective sound speed c^2_{eff} of the axion. In other words, it determines when the sound speed terms become relevant in the evolution of the perturbations.

Returning to Fig. 3.4, we can see how the evolution of the Jeans scale (red line) has a differing affect on large k modes versus small k modes. First, let's consider the case of a large axion mass such that the field begins to oscillate before matter radiation equality a_{eq}. We see that all $k < k_{\mathrm{m}}$ are non-relativistic when crossing the horizon, and are always smaller than k_{J}, thus they act identically to CDM. However, for those modes $k > k_{\mathrm{m}}$, they become non-relativistic after horizon crossing, but remain larger than k_{J} until the scaling relation of $H(a)$ changes after matter-radiation equality. These large k modes will be suppressed relative to the small k modes until they cross k_{J}. If we consider axion masses sufficiently small such that the axion begins oscillating after a_{eq} (bottom panel), an additional possibility emerges - the possibility for a mode $k_{\mathrm{J}} < k < k_{\mathrm{m}}$ before a_{eq}. In this situation we can have a mode that is suppressed by the sound speed while the background is still radiation dominated. However, this situation likely is not of much interest because we note that this situation can only occur before the axion field begins to oscillate. In such a regime, the axion still possesses the equation of state of dark energy and we can exactly evaluate the evolution of the field.

We now know the relevant scales needed to determine when the axion behaves differently from regular CDM. We can use these scales to infer the shapes and sizes of particular features of the matter power spectrum. In the case of ultra-light axions with $m_\varphi < m_{\rm eq}$ eV such that they oscillate after matter-radiation equality, we expect a suppression to small scales beginning at $k_{\rm m}$ and suppression factor that grows with the time between matter-radiation equality and the mode becoming larger than $k_{\rm J}$

$$\frac{P_{\rm d}}{P_{\rm d,\Lambda CDM}} \approx 1 - \frac{3\omega_\varphi}{4\omega_{\rm m}} \log\left(\frac{a_{\rm J}}{a_{\rm eq}}\right) \tag{3.34}$$

where $a_{\rm J}$ is the scale factor when the mode crosses $k_{\rm J}$ [32]. We expect this suppression factor to manifest as a step-like feature in the matter power spectrum and see that it appears when we compute $P_{\rm mm}$ for various cases[¶] in the next section.

3.2.4 Computation with the `axionCAMB` Code

To compute the transfer functions and background evolution of a cosmology consisting of the standard ΛCDM species with an additional (single) axion fluid produced via the Misalignment Mechanism, we use the modified version of the `CAMB` Boltzmann solver known as `axionCAMB` [107]. For a given cosmology and axion mass, the exact field-level equations are used by `axionCAMB` to track the background evolution of $c^2_{\rm eff}$, P_φ,ρ_φ and w_φ at all times. Additionally, when $a < a_{\rm osc}$, the exact GDM fluid equations of motions are used to track the evolution of $\delta\rho$, δp, u, and δ. When $a > a_{osc}$, the time averaged GDM fluid equations of motion are used instead, and the background equation of state for the axion is assumed to be exactly $w_\varphi = 0$. We restrict our analysis to masses which have $a_{\rm osc} < 1$, roughly corresponding to $m_\varphi > 10^{-33}$ eV. We assume that the symmetry breaking scale for the axion field must be higher than that of the end of inflation, which we set to $H_{\rm inf} = 10^{13.7}$ GeV. Eq. 3.9 is used to exactly calculate$\rho_\varphi(a)$ until the point where $m_\varphi = 3H(a)$, after which $w_\varphi = 0$ and the axion energy density exactly scales like matter$\rho_\varphi \propto a^{-3}$. The `axionCAMB` code

---

[¶] we also consider cases of early time oscillation. Need to derive results for those cases as well.

iterates the initial field value $\phi_\phi(z \gg 0)$ until the requested relic density is recovered to precision of 10^{-4}. Because the axion relic always behaves as CDM *today*, we subtract the axion relic density from the CDM relic density when running `axionCAMB` so that the quantity $\omega_\phi + \omega_{\mathrm{CDM}}$ is constant throughout this work.

Using `axionCAMB` we calculate P_{mm} for various axion parameter choices, and identify the mass and abundance dependent suppression to the matter power spectrum matching results previously recovered in the literature [125], as shown in Fig. 3.5. We can also look for the redshift dependent features in the axion suppression to P_{mm} in Fig. 3.6. First, notice that only those masses which begin to oscillate after equality introduce a BAO phase shift, apparent in the appearance of oscillations around BAO scales for those masses. We see that the heaviest mass, which oscillates well before equality, always introduces a step like suppression to the matter power spectrum due to its Jeans Length. Also, that suppression grows with the time since equality. For the lightest masses, these fields will not transition before the onset of matter-radiation equality and so there is an additional effect set by the scale entering the horizon when the field begins to oscillate, k_{m}. Scales larger than k_{m} will be outside the horizon at equality. In this case, part of the CDM which would normally be present during matter dominated perturbation evolution is not. We should expect that anytime oscillation occurs at low redshifts where non-linear structure collapse is occurring, there should be a distinct signature in the structure formation observables, which we consider next.

3.3 STRUCTURE FORMATION WITH ULTRA-LIGHT AXIONS

The presence of an ultra-light axion modifies the process of structure formation throughout cosmic history. These changes can be summarized through three effects. First, the axion modifies the ratios of standard cosmological species to each other. Shifting these abundances presents a new expansion history $H(z)$, changes when periods of equality occur and alters the speed and extent to which over-

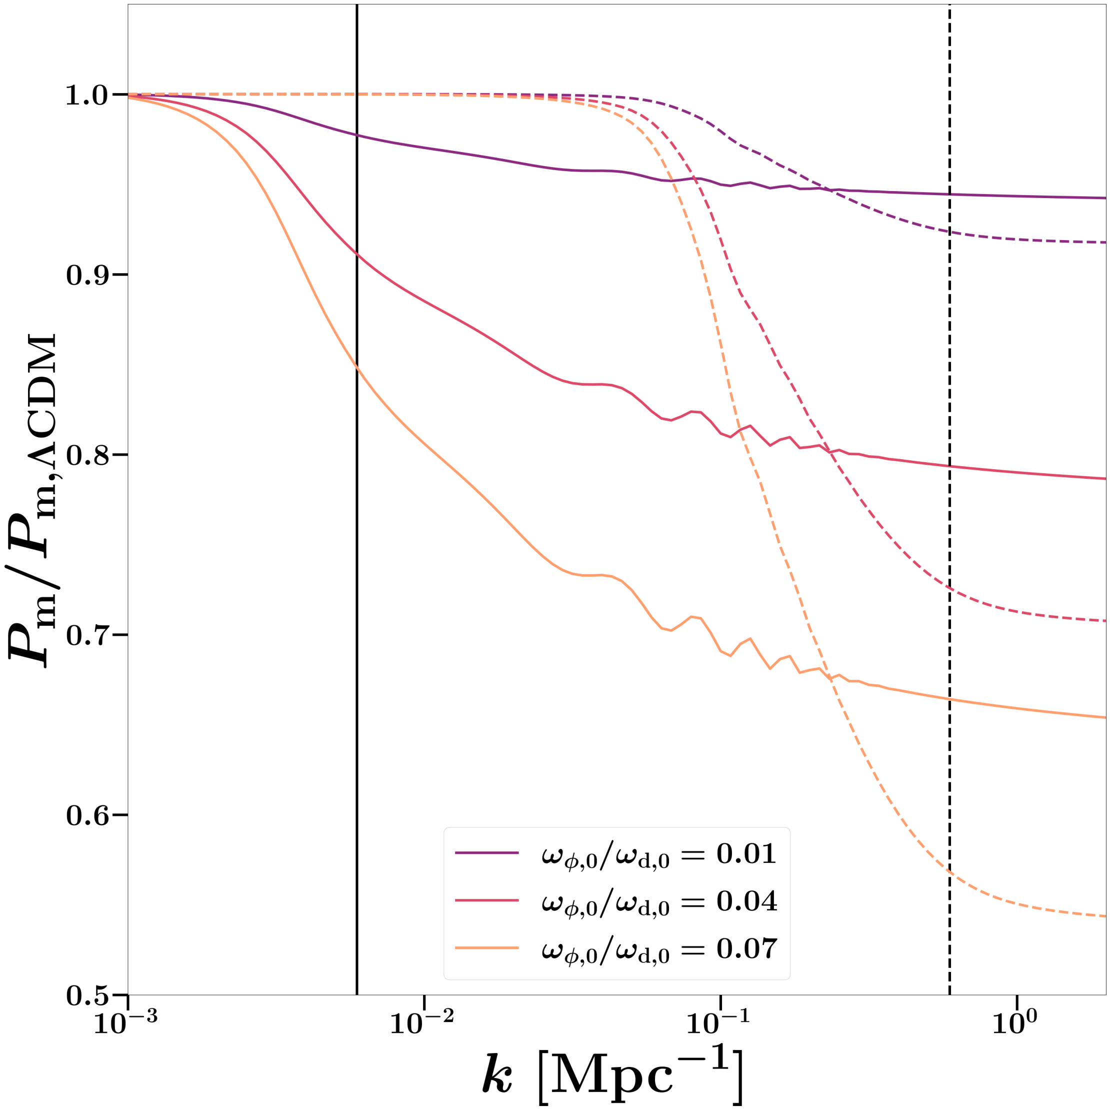

Figure 3.5: The matter power spectrum for three choices of axion abundance normalized to the ΛCDM value. Two choices of axion mass $m_\varphi = 10^{-26}$ eV (dashed line) and $m_\varphi = 10^{-30}$ eV (solid line) are shown with the corresponding Jeans scales indicated by black lines. In the case of the lighter axion mass, the field does not begin oscillating until after matter-radiation equality, so a redshift of matter-radiation equality is slightly shifted in this case, with a corresponding phase shift introduced to the BAO. This shift in the BAO phase is evidenced by the presence of oscillations around BAO scales in the relative matter power spectrum for the lighter axion mass, but not the heavier.

densities and perturbations evolve. These effects were reviewed in §3.2 and are calculated using the `axionCAMB` code. Second, the presence of an axion modifies the non-linear dynamics of halo collapse

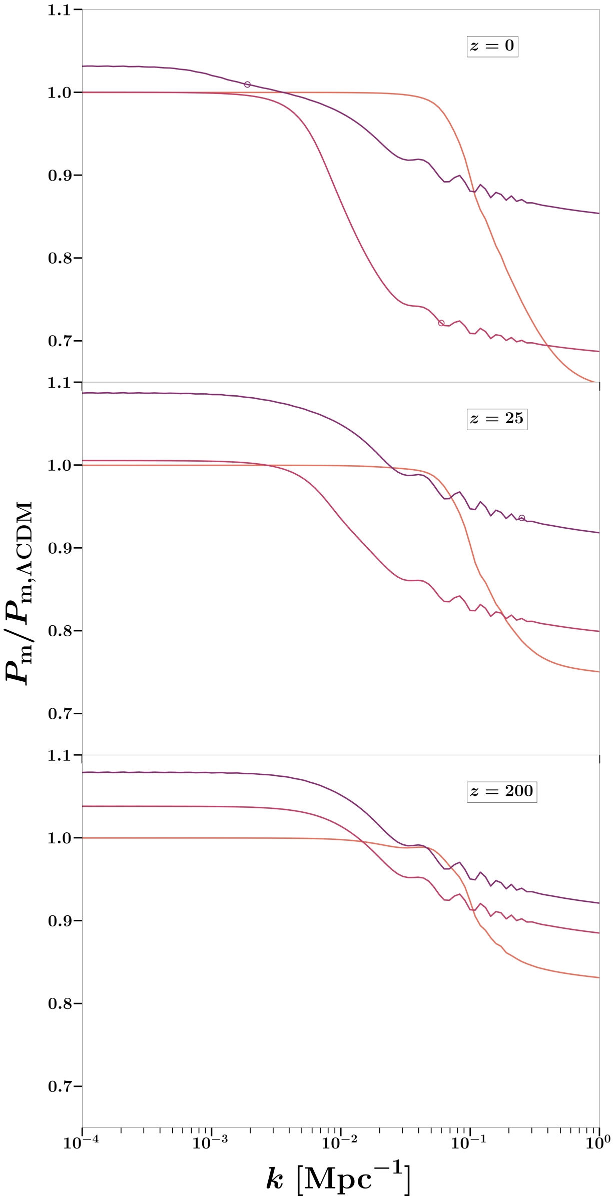

Figure 3.6: The matter power spectrum normalized to the ΛCDM value with $\Omega_\varphi/\Omega_{\rm CDM} = 0.05$ for three choices of redshift. In each case, the Jeans scale for a given axion mass is indicated by a circle. In the high-mass limit, the axion Jeans length becomes cosmologically small and it behaves identically to CDM. At the lowest masses, the Jeans length approaches the horizon size and acts identically to DE. For the three axion masses above, the field begins to oscillate at $z_{\rm osc}(m_\varphi = 10^{-26}\ {\rm eV}) \approx 17,800$, $z_{\rm osc}(m_\varphi = 10^{-29}\ {\rm eV}) \approx 321$, $z_{\rm osc}(m_\varphi = 10^{-32}\ {\rm eV}) \approx 2$.

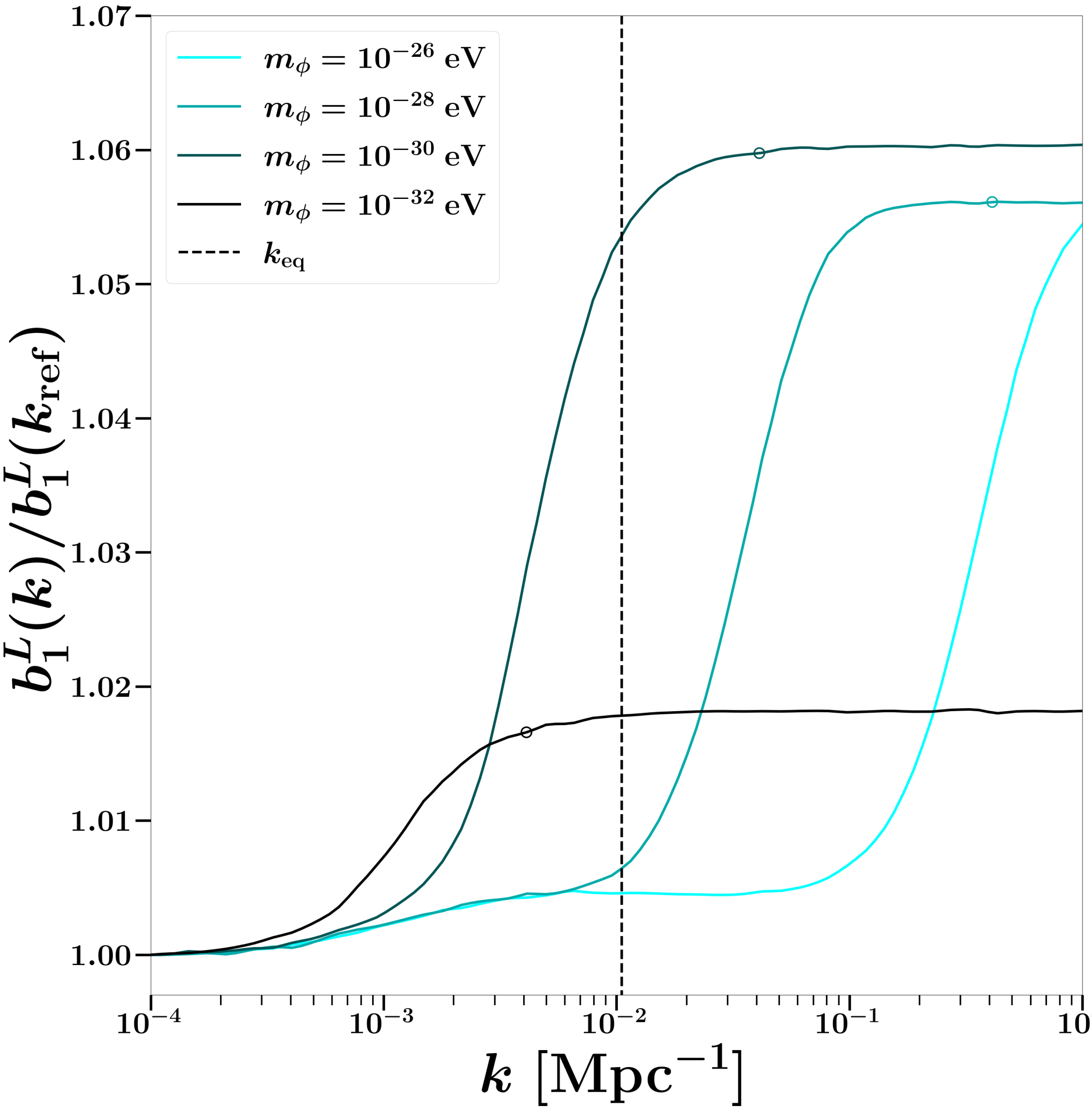

Figure 3.7: The Lagrangian bias normalized to the ΛCDM value for various choices of axion mass. The Jeans scale at $z = 0.65$ is indicated for a given axion mass by a circle. A halo mass of $10^{13}M_\odot$ is assumed. Note the presence of the percent level ΛCDM step in the bias at scales smaller than $k_{\rm eq}$ along with a second, larger step at scales set by the axion mass. In the case of the lightest masses, an apparent attenuation in the step height occurs due to the late oscillation time of the field.

at late times. One driver of this change is the difference in the ratio of clustering to non-clustering matter during collapse. We consider this effect in §3.3.1. Third, an axion will change the critical overdensity necessary for a halo to collapse at a given scale. Along with axion-induced changes to the Halo Mass Function (HMF), this manifests a scale-dependent modification to the halo bias. We explore the axion modifications to the bias in §3.3.3. In all these cases, note that care must be taken when introducing the axion as the expression of the sound speed at the time of field oscillation will depart from approximations which match in the far from oscillation regimes. At the time of this dissertation, the proper method of incorporating this effect is still being considered. As such, results for axion fields which begin oscillating around the redshift at which we initialize the collapse procedure should be considered with limited confidence.

3.3.1 Halo Collapse

To model how the presence of an ultra-light axion affects the structure collapse process, our goal is to parameterize how the halo radius as a function of time $R(t)$ is modified by an ultra-light axion. A complete consideration of the spherical collapse process in the presence of a GDM model is considered in [163]. Under standard spherical collapse, the evolution of a halo is given by

$$\ddot{R}(t) = -\frac{GM}{R^2(t)} - \frac{4\pi G R(t)}{3} \sum_i \left(\rho_i(t) + 3P_i(t)\right) . \tag{3.35}$$

Here, the index summation over i includes all species besides CDM and baryons. For the axion, this means an additional contribution of $\rho_\varphi + 3w_\varphi \rho_\varphi$ to Eq. (3.35). At very early times, $a \ll a_{\rm osc}$, before the axion has started to oscillate, the axion will generate a contribution of $-2\rho_\varphi$. At late times, $a \gg a_{\rm osc}$, for the $n = 1$ axion, this contribution will instead be simply ρ_φ. If the axion transitions its equation of state during spherical collapse, it will leave a distinct impression on $R(t)$ that cannot be exactly duplicated by the standard cosmological fields. Even in the case of generic Light but Massive Relics (LiMRs), amongst which massive neutrinos are a special case, which are known to free stream, the early universe equation of state will look like radiation, not dark energy. At all times, the exact expression for Eq. (3.35) used in calculations is given by

$$\begin{aligned}\ddot{R}(t) = -\frac{GM}{R^2(t)} - \frac{4\pi G R(t)}{3}\Big(&- 2\rho_\Lambda(t) + 2\rho_\gamma(t) \\ &+\rho_\nu(t) + 3w_\nu(t)\rho_\nu(t) + \rho_\varphi(t) + 3w_\varphi(t)\rho_\varphi(t)\Big)\,.\end{aligned} \tag{3.36}$$

Note we include dependence on $w_\nu(t)$ only when we include massive neutrinos in our cosmology. However, in the main text of this work, we only consider the massless neutrino case, in which case the neutrino terms in Eq. 3.36 are set to zero. We consider the case of massive neutrinos, including degeneracies with ultra-light axions, in Appendix §3.5.

At the level of perturbations, the energy densities and pressures are modified to [152]

$$\rho_i(z) = \bar{\rho}_i(z)(1 + \delta_i(z))\,, \tag{3.37}$$

$$P_i(z) = \bar{P}_i(z)\Big(1 + \frac{c^2_{\mathrm{eff},i}(z)}{w_i(z)}\delta_i(z)\Big)\,, \tag{3.38}$$

where overbars denote the background value.

In this work, the conditions for determining whether a halo collapses are as set in the peak-background split [91,189] model of halo collapse. In this framework, δ_i can be expressed in terms of some long δ_L wavelength perturbations we introduce into the system

$$\delta_i \equiv \delta_L \frac{\mathcal{T}_i(k,z)}{\mathcal{T}_c(k,z_{\mathrm{ini}})}\,, \tag{3.39}$$

where $\mathcal{T}_c(k, z_{\mathrm{ini}})$ is the CDM+B transfer function. The evolution of this long wavelength mode throughout collapse is set by ratios of transfer functions between some initial redshift where the collapse process is chosen to begin, z_{ini} and some redshift we chose for collapse z_{coll}. Additionally, a short wavelength perturbation δ_S is introduced at $z_{\rm ini}$ and we are interested in the magnitude of short wavelength perturbation, at a chosen scale, is necessary to produce collapse by $z_{\rm coll}$. We will call the minimum satisfactory small wavelength perturbation, evaluated at the collapse redshift, the critical overdensity $\delta_{\rm crit}$. Likewise, $\delta_{\rm L}$, evaluated at the collapse redshift, will be $\delta_{\rm L,coll}$. To determine these quantities, the halo collapse $R(t)$ is computed between $z_{\rm ini}$ and $z_{\rm coll}$ in the presence of these long and short wavelength perturbations using the specially developed code `RelAxiFAST`. This code is used to compute the value of the critical overdensity for collapse $\delta_{\rm crit}(k)$ at some redshift of collapse $z_{\rm coll}$ in the presence of long δ_L and short δ_S wavelength perturbations at some initial redshift of collapse $z_{\rm ini}$.

3.3.2 COMPUTATION WITH THE `RelAxiFast` CODE

To track the evolution of CDM+baryon halos in the presence of background overdensities of long and short wavelengths, we will use a modified version of the `RelicFast` code. The dynamics of the halo evolution computed by the `RelicFast` code and the initial conditions used for halo evolution are described in detail in [152]. The `RelAxiFast` execution is equivalent up to the following modifications:

1. The background energy densities are modified by the presence of an ultra-light axion with time varying equation of state. This modifies all background quantities, including the energy densities $\rho_i(a)$, the scale factor $a(\eta)$, Hubble parameter $H(a)$, and the transfer functions $\mathcal{T}_i$. The full background cosmology in the presence of an ultralight axion is computed by the `axionCAMB` code at `RelAxiFast` run time and collected for use in the bias calculation. In this sense, `axionCAMB` replaces the role of `CLASS` in the original `RelicFAST` implementation.

2. The perturbation evolution is also modified in the presence of an ultra-light axion. The axion fluid pressure P_φ, effective sound speed $c^2_{\rm eff}$ and oscillation redshift $a_{\rm osc}$ are also calculated

by `axionCAMB` at run time and used by `RelAxiFast` to compute the evolution of the axion fluid perturbations.

3. `RelAxiFast` evaluates Eq. 3.36, a modified version of the spherical collapse equation of motion to account for the presence of a time varying axion fluid. Qualitative changes which occur when the axion field begins to oscillate are accounted for.

With these modifications, `RelAxiFast` quickly ($\mathcal{O}(10)$ seconds) calculates the critical overdensity for collapse, including calculation of the axion fluid background. If the axion background is pre-generated, the bias calculation is much faster (< 1 second). Because of the highly non-linear nature of the collapse, the critical overdensity of collapse is calculated via a shooting method. As input, a collapse redshift is specified by which an initial matter density should form a halo by. Then, for many long wavelength perturbations k_L, `RelAxiFast` calculates the critical overdensity of collapse $\delta_\mathrm{crit}(k_\mathrm{L}, \delta_\mathrm{L}, z_\mathrm{collapse})$ by varying an initial short wavelength overdensity δ_S at an initial redshift of collapse of $z_\mathrm{ini} = 200$, evolving the system forward in time, and varying δ_S until the collapse condition is satisfied at z_collapse. The solutions for δ_crit are then used to compute the halo power spectrum $P_\mathrm{hh}(k)$, which (at a single redshift) relates to the matter power spectrum $P_\mathrm{mm}(k)$ through the scale dependent bias function

$$P_\mathrm{hh}(k) = b_1^2(k) P_\mathrm{mm}(k)\,. \tag{3.40}$$

3.3.3 The Scale Dependent Halo Bias

The halo bias provides a functional description of the likelihood of a halo to collapse in a background matter field. Particularly, it is a measure of the dependence of the halo formation process on different perturbation modes. The halo bias is known[152] to develop a scale dependence in response to different modes growing at different rates (e.g. a step due to the transition from radiation to matter domination in ΛCDM) and we will see similar scale dependencies develop due to late-time transitions in the axion equation of state which we can search for in cosmological data sets. In Fig. 3.8, we see an example of how the matter and halo power spectra, here represented by a suppression factor relative to an equivalent ΛCDM cosmology, can evolve to different values.

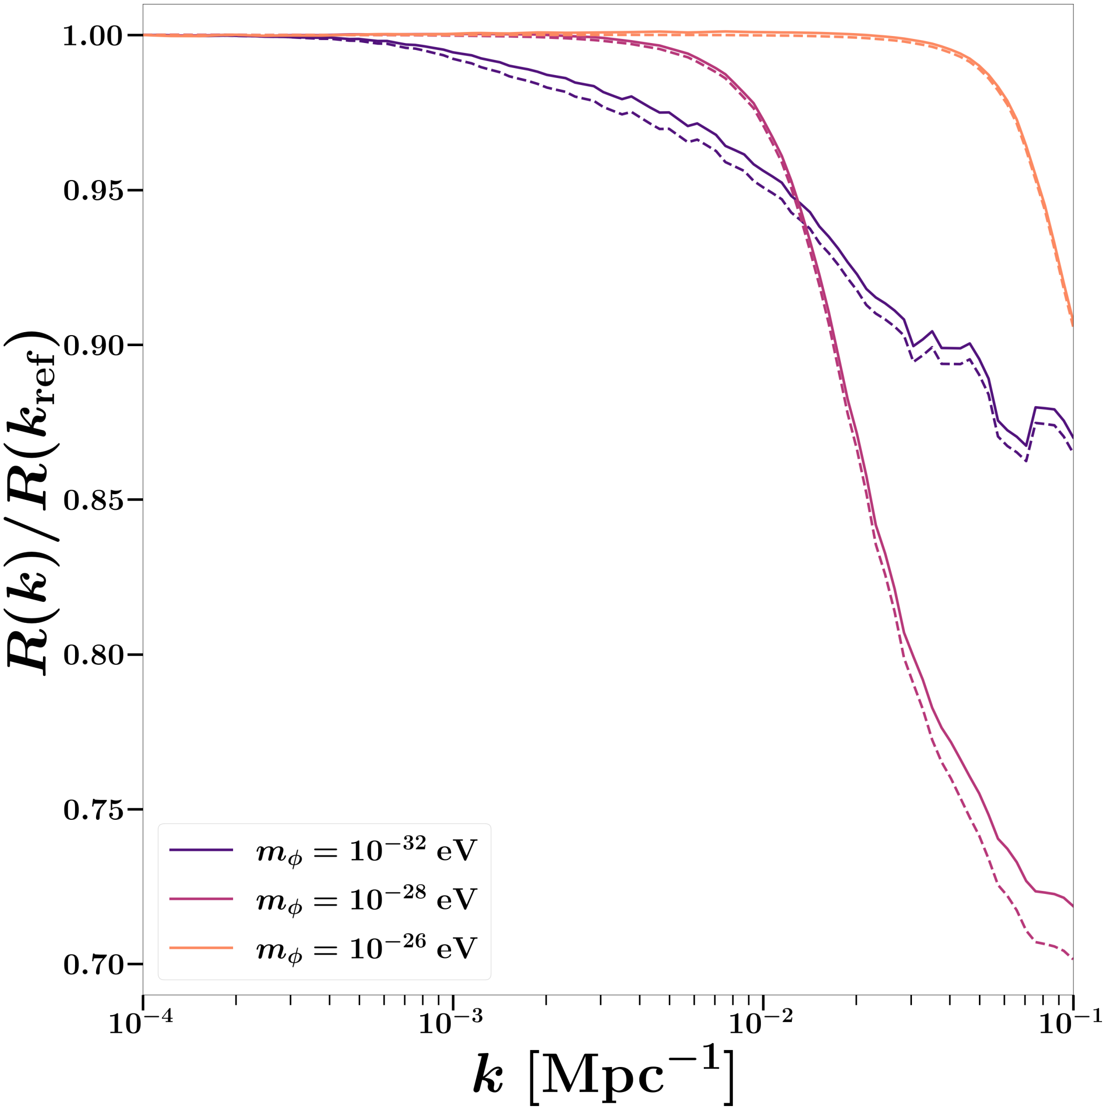

Figure 3.8: Suppression factor normalized at $k = 10^{-4}$ Mpc^{-1} for matter (dashed) and halo (solid) power spectra. This is for halo masses $M = 10^{13} M_\odot$ collapsing at $z = 0.65$.

Prior work has shown that the introduction of light relics to the cosmology introduces a scale-dependence in the halo bias due to their free streaming[115,165,75]. Scale dependence emerges in the halo bias as a reflection of different perturbation modes evolving at different rates over cosmic history[62,187]. While different in nature, the axion introduces a qualitatively similar suppression to the matter power spectrum as do free-streaming relics, and it also introduces distinct scales $k_{\rm J}$ and $k_{\rm m}$ which modify the evolution of perturbation modes, so it is reasonable to consider how they might manifest a scale dependence in the halo bias as well.

At linear order, the Lagrangian bias is defined by

$$b_1^L(k) = \left(\frac{\partial \log n}{\partial \delta_{\rm crit}}\right)\Big|_{\delta_{L,\rm coll=0}} \left(\frac{d\delta_{\rm crit}}{d\delta_{L,\rm coll=0}(k)}\right) , \tag{3.41}$$

where the first term is the Halo Mass Function (HMF) and the second term is calculated by solving for spherical collapse $R(t)$ assuming the peak-background split model. This work is primarily concerned with calculating the scale dependence of the second term using the combined `axionCAMB` + `RelAxiFast` codes. There are few points to mention concerning the first term, the HMF function.

The halo mass function in the CDM paradigm has been extensively modeled and tested to great success leading to various calibrations to the Press-Schechter HMF[173,189,118,195,77]. In contrast, much fewer work has been done to study the HMF for the FDM case[143,124,145,59,186]. In particular,[186] demonstrated a substantial reduction in the number of low mass halos at late redshifts in comparison to standard CDM HMF realizations, while agreeing with the CDM HMF for high masses. With that said, these studies have only modeled ultra-light axions in the realization where they compose all of the CDM. As such, they tend to only consider masses at the higher end of the range we consider here $M_\varphi > 10^{-22}$ eV. For these reasons, we avoid using a special HMF derived from axion simulation. In this work, we use the HMF derived from the `MICE` simulation[67]. Alternatively, there is the possibility of modifying an effective field theory prescription for generalized DM as in[163] or[124]. We do not consider effective descriptions of the ultra-light axion HMF here. In any case, Fig. 3.9 demonstrates that while the Eulerian bias is quite sensitive to different halo masses, the Lagrangian bias, which is the quantity studied throughout this work, is affected relatively little.

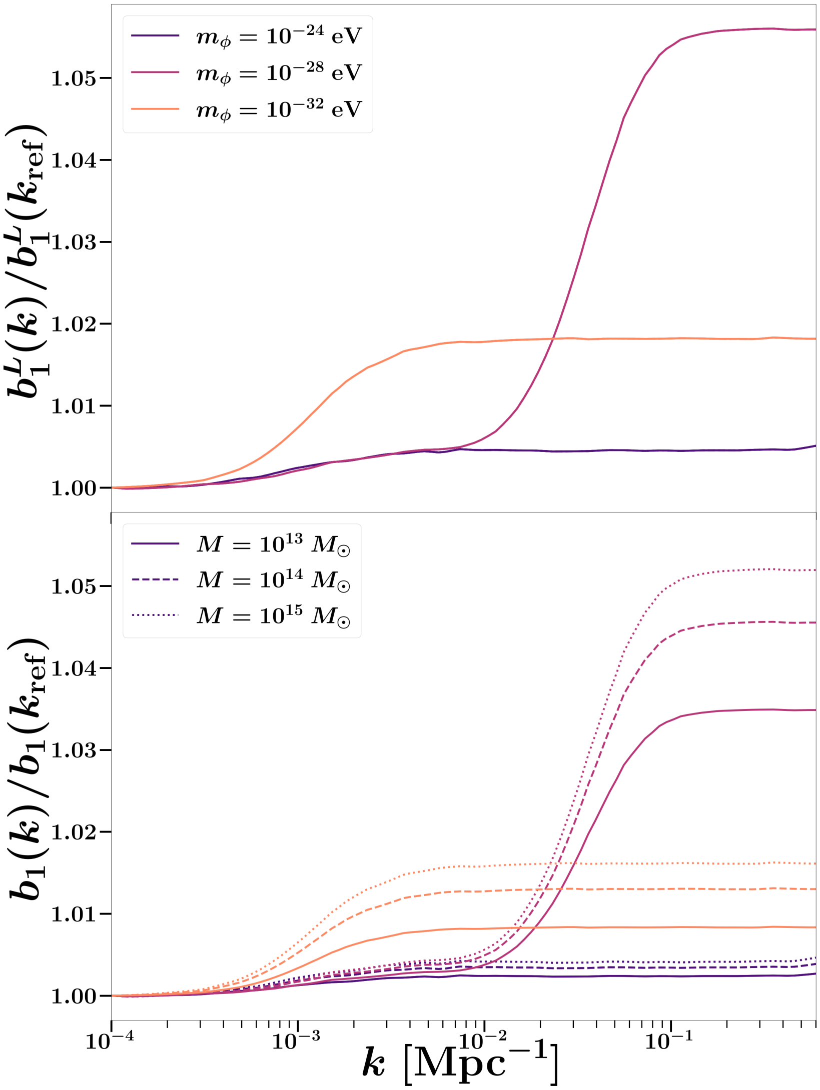

Figure 3.9: The Eulerian bias (top panel) and Lagrangian bias (bottom panel) normalized to the ΛCDM value for various axion mass and halo mass choices. Note that the effect of varying the halo mass is nearly indistinguishable in the Lagrangian bias.

It is known that the ΛCDM cosmology introduces a step-like enhancement to the Lagrangian bias of the form

$$\begin{aligned} R_L^{\Lambda\mathrm{CDM}} &\equiv \frac{b_1^{L,\Lambda\mathrm{CDM}}(k)}{b_1^{L,\Lambda\mathrm{CDM}}(k_{\mathrm{ref}})} \\ &= 1 + \Delta_{\Lambda\mathrm{CDM}} \tanh\left(\alpha k / k_{\mathrm{eq}}\right) \end{aligned} \tag{3.42}$$

with $\Delta_{\Lambda\mathrm{CDM}} = 4.8 \times 10^{-3}$ and $\alpha = 4$. Similarly, light massive thermal relics (LiMRs) - including massive neutrinos - with temperature today $T_X^{(0)}$ produce a similar step-like feature in the bias at the relic's free streaming scale k_{fs}, given during matter domination by

$$k_{\mathrm{fs}} = \frac{0.08}{\sqrt{1+z}} \left(\frac{m_X}{0.1\,\mathrm{eV}}\right) \left(\frac{T_X^{(0)}}{T_\nu^{(0)}}\right)^{-1} h\mathrm{Mpc}^{-1} \tag{3.43}$$

and with an amplitude set by the relic's abundance f_X

$$R_L^X = 1 + \frac{\Delta_X}{2} \left(\tanh\left(\frac{\log q}{\Delta_q}\right) + 1\right) \tag{3.44}$$

with $\Delta_X = 0.6 f_X, q = 5k/k_{\mathrm{fs}}$ and $\Delta_q = 1.6$[152]. In a similar manner, we provide the following fit to the Lagrangian bias response in the presence of ultra-light axions

$$R_L^\varphi = 1 + \Delta_\varphi \tanh\left(\frac{k}{k_{\mathrm{J}}}\right), \tag{3.45}$$

$$\Delta_\varphi \equiv \alpha \left(\frac{\omega_\varphi}{\omega_{\mathrm{d}}}\right) \tanh\left(\left(\frac{a_{\mathrm{eq}}}{a_{\mathrm{osc}}}\right)(1 - a_{\mathrm{osc}})\right). \tag{3.46}$$

Eq. 3.46 represents the amplitude of the bias step induced by the relic, which is proportional to its fractional abundance (the first term) and is attenuated if a mode enters the horizon after a_{eq} due to late oscillation of the axion (the second term). In the large mass limit, Eq. 3.45 should resemble an unattenuated step at the free streaming scale akin to thermal relic dark matter. So we see that the qualitative behavior of ultra-light axions resembles that of thermal relic dark matter up to an additional redshift dependence in the amplitude.

3.4 Results and Discussion

The `RelAxiFast` procedure, combined with the MICE HMF prescription provides a computation of the linear Lagrangian bias. In Fig. 3.7, we compute the Lagrangian bias for various choices of axion mass at $z_{\rm coll} = 0.65$. We observe the presence of a percent level enhancement to the scale dependent halo bias which, appearing at a scale set by the axion mass through the axion Jeans scale, $k_{\rm J}$. Additionally, a redshift dependence to the halo bias is manifest as a result of the ultra-light axion transitioning its equation of state when the axion field begins to oscillate. The effects of the axion beginning to oscillate at time $a_{\rm osc}$ are two-fold. First, this oscillation can happen before or after matter radiation equality $a_{\rm eq}$. Large masses which oscillate before $a_{\rm eq}$ are growth suppressed until the time of matter domination. However, for lighter masses with $a_{\rm osc} > a_{\rm eq}$, modes which enter the horizon before the field begins to oscillate are additionally suppressed. The second consequence is that this redshift dependent suppression is *in addition* to the suppression to modes which are below the Jeans scale of the axion.

In the case of the largest mass axions, the behavior closely resembles that of CDM as they begin oscillating long before structure formation. In this case, each perturbation mode will evolve identically up to a modification to modes entering the horizon before or after matter-radiation equality, which is a purely ΛCDM effect. So we expect no additional step due to the axion Jeans length, for heavy masses, because we are not sensitive to the very small modes effected by the axion field transition. In the case of very light masses, the Jeans wavelength approaches the Hubble scale today and will only affect modes just entering the horizon. Qualitatively, we expect all but the largest modes to evolve similarly to a cosmology which is ΛCDM with a slightly enhanced amount of dark energy and reduced amount of CDM. Again, since nearly all modes are affected identically, we expect no scale dependent effect to emerge in this case either (apart from the ΛCDM step). Another way to realize this is to consider that the only modes affected by very light masses will have only just entered the horizon. As such, there has not been much time for the integrated effect on $R(z)$ to become significant. This offers an explanation for the gradual suppression in the bias amplitude as you move towards the lightest masses.

In Fig. 3.10 we consider the redshift dependence of the bias for choices of the ultra-light axion mass at representative locations in the mass range we consider. Here we see the interactions of the various scales relevant to the axion's quantum pressure and how they are modified by the times of oscillation and mode crossing.

Previous literature has demonstrated that introducing an ultra-light axion field into the λCDM cosmology will generate a suppression in P_{mm} due to the macroscopic wavelength of the fundamental axion field. In §3.2.2 and §3.2.3 we reviewed how this signature manifests in cosmological observables. While modeling the axion in the linearly perturbed universe is a critical first step towards measuring axion signatures in cosmological data sets, it is also important to model the non-perturbative effects of an ultra-light axion in the non-linear structure formation process.

In this work, we presented the first characterization of the ultra-light axion's effect on the linear halo bias. We demonstrated a unique signature in the bias which partially compensates for the suppression in the matter power spectrum at the percent level. The unique relationship between the mass of an ultra-light axion and various factors of suppression to its perturbative modes is not obviously replicated in other cosmological models of physics beyond the Standard Model, making the halo bias a rich source of information for constraining such models. It should be emphasized that the non-linear process of structure formation is distinct from the evolution of the linear cosmological background; structure formation is sensitive to the presence of an ultra-light axion in ways different from the background cosmology and therefore provides information in addition to the

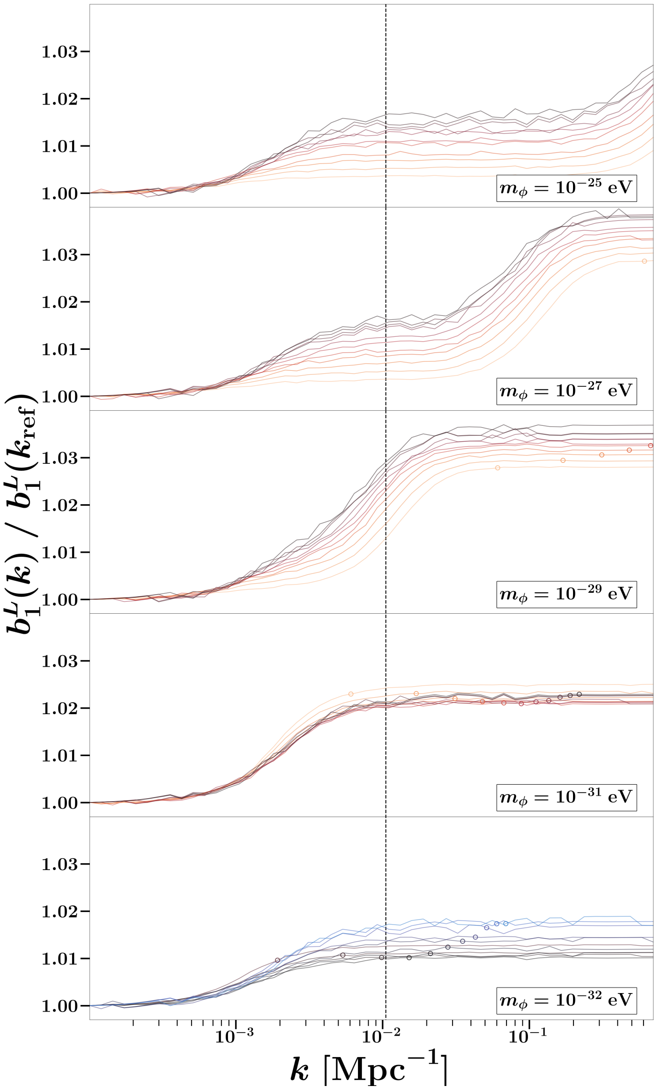

Figure 3.10: Lagrangian bias at various collapse redshifts normalized at $k = 10^{-4}$ Mpc^{-1}. Redshifts linearly span from $z \approx 0$ to $z = 10$. Red (blue) colors are redshifts before (after) oscillation occurs for given mass. The color scale progresses from red (small redshift) to black to blue (high redshifts). Circles indicate $k_{\rm J}$. Note that in the lightest mass cases, the enhancement to the bias at the expected Jean's scale is suppressed relative to the heavier mass cases. This is because even those modes which enter the horizon during matter remain suppressed until the axion field begins oscillating. So only those modes which are larger than the size of the horizon when the field begins oscillating will be enhanced. That is to say, late time oscillation "cools" the size of the bias steps introduced by the axion field.

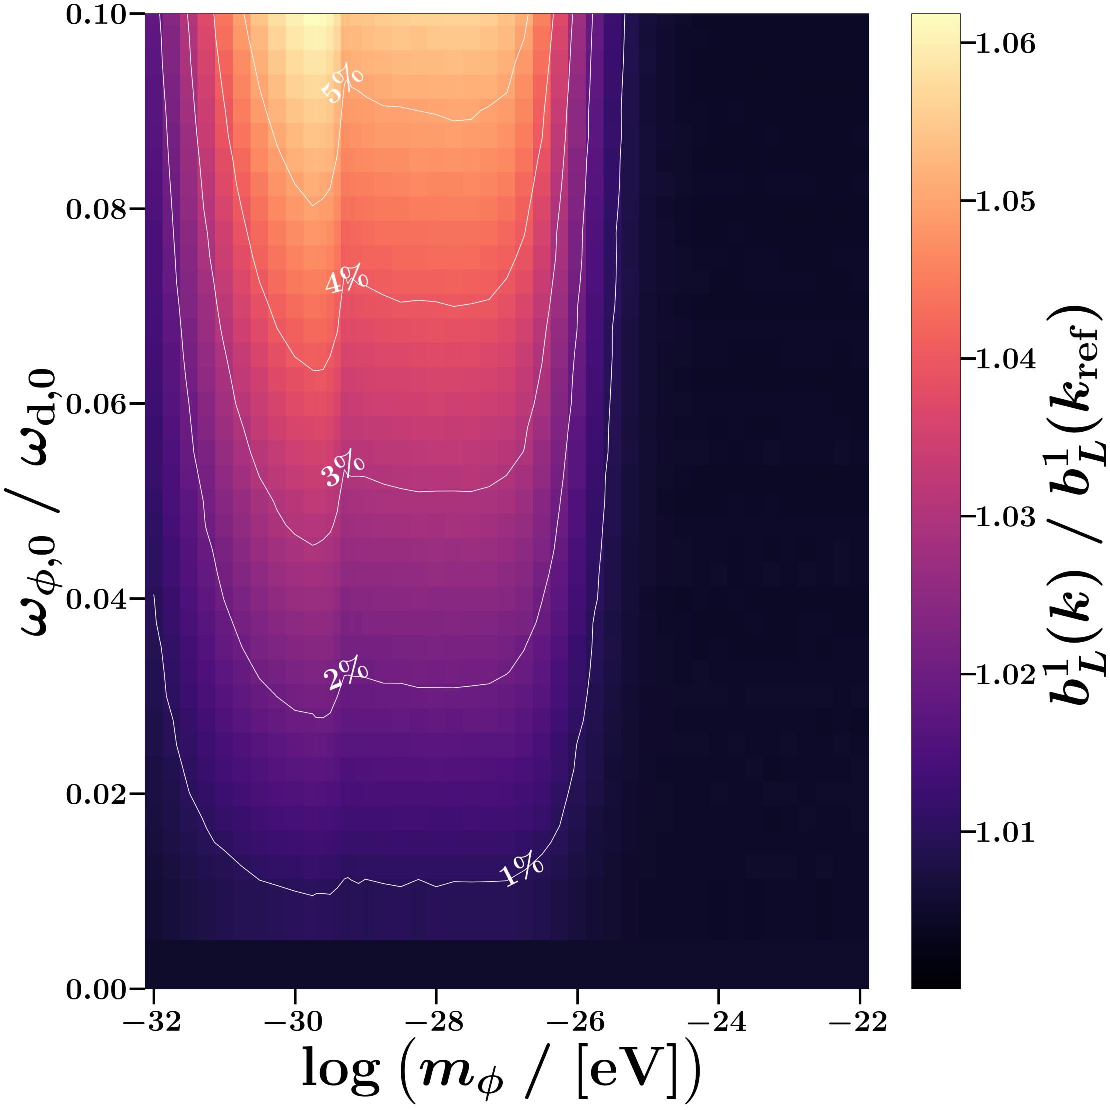

Figure 3.11: The amplitude at $k = 10^{-0.5}$ Mpc^{-1} of the Lagrangian bias normalized at $k_{\mathrm{ref}} = 10^{-4}$ Mpc^{-1} for various axion masses and abundances. Percent level contours are in white. At the time of writing of this dissertation, the origin of the jump corresponding to $m_\varphi = 10^{-29.25}$ eV (the mass beginning to oscillate at the initialization redshift of collapse) is still being explored. Note, for example, that care must be taken with approximate expressions for the sound speed of axion fields which begin oscillating around the redshift at which we begin analyzing collapse.

cosmological signature. Further, modeling ultra-light axions into the structure collapse evolution is a more accurate model even if it does not provide enough distinguishing power to better constrain the ultra-light axion model parameters. This may manifest, for example, as a shift in the reported maximum likelihood values reported by analyses with and without ultra-light axions modeled in the collapse functions. While not pursued in this work, it is interesting to consider the possibility of modeling the clustering of ultra-light axions in spherical collapse processes in future work. Upcoming cosmological probes will reach into higher redshifts where perturbations at smaller scales are still small enough to be linear - the prospect of such probes investigating the scale dependent effects in the halo bias introduced by the lightest axion masses is exciting . In light of previous work on massive neutrinos, and light but massive relics, which share some qualitatively similar effects as the axion but with different fundamental origin, we expect that proper consideration of the axion bias is critical in any constraint of ultra-light axions using galaxy surveys at the risk of mischaracterizing the axion mass otherwise.

3.5 Comparing Massive Neutrinos

Thermal relics are known to introduce a suppression to the matter power spectrum and scale dependent effects in the halo bias as a consequence of their free-streaming while relativistic. As an example, massive neutrinos will have an equation of state that transitions from $w = 1/3$ to $w = 0$ when $m_\nu = T_\nu$. While relativistic, neutrinos will free stream and are known to suppress the matter power spectrum below their corresponding free-streaming scale

$$k_{\mathrm{fs}} = \frac{0.08}{\sqrt{1+z}} \left(\frac{m_{\nu,i}}{100\ \mathrm{meV}} \right) h\,\mathrm{Mpc}^{-1}\,. \tag{3.47}$$

We would like to qualitatively understand how degenerate the imprint on the scale dependent bias is between ultra-light axions and thermal relics such as massive neutrinos. To test this, we compare one cosmology with massive neutrinos and a second cosmology where the neutrinos are made massless, but an axion is introduced with the same relic abundance and axion mass chosen such that k_{J} equals

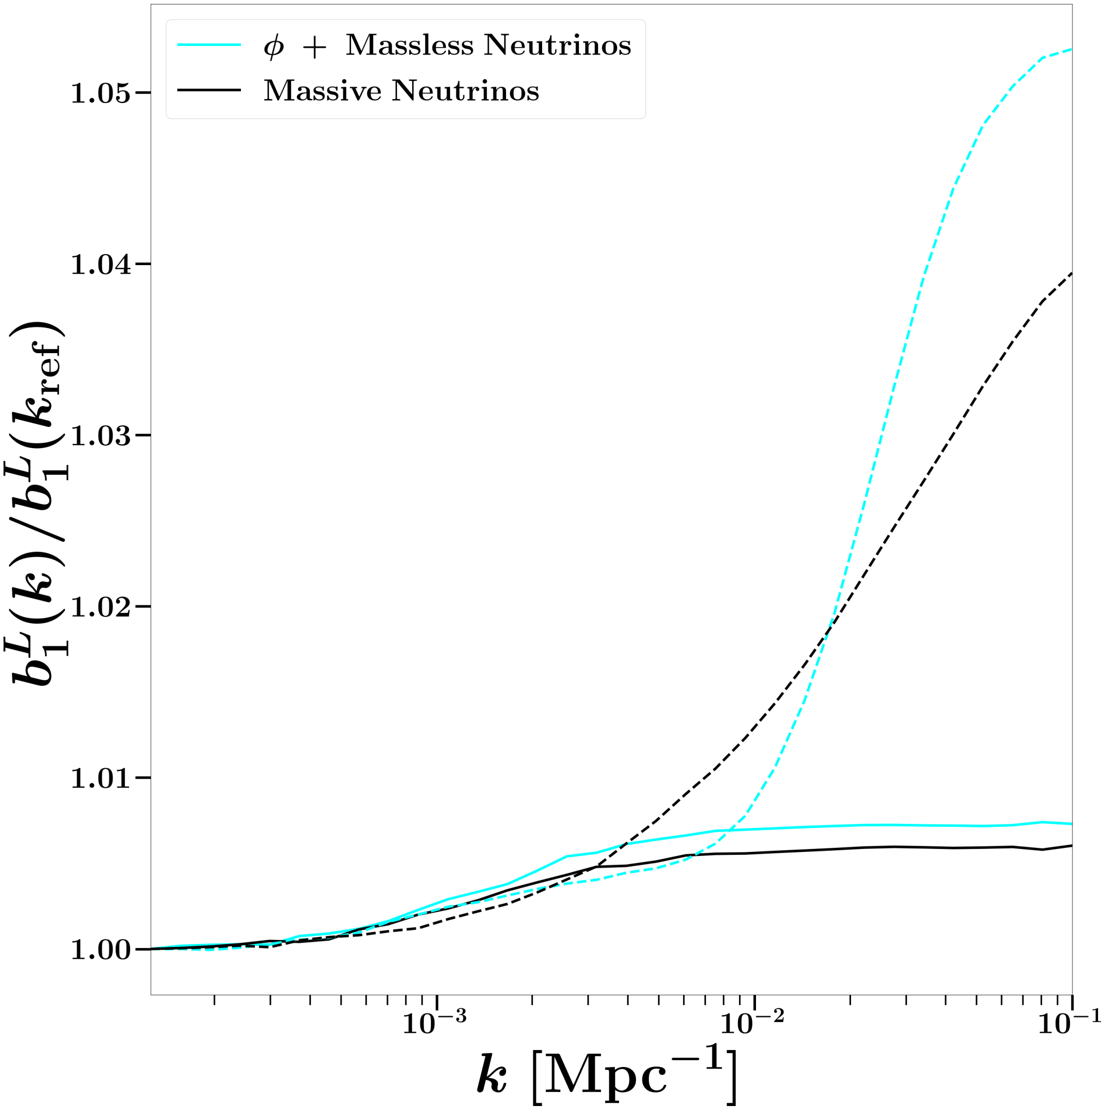

Figure 3.12: Lagrangian bias normalized at $k = 10^{-4}$ Mpc^{-1} for a cosmology with three, degenerate, massive neutrinos with $\Sigma m_\nu = 60$ meV (solid) and $\Sigma m_\nu = 1$ eV (dashed). We compare (in black) with an identical cosmology where the neutrinos are made massless and an axion with the same abundance and $k_{\rm J}$ equal to the neutrino free-streaming scale.

the free-streaming scale of the neutrinos from the first scenario. As an approximation to using either the inverted or normal neutrino hierarchy, we consider a three neutrino degenerate hierarchy with

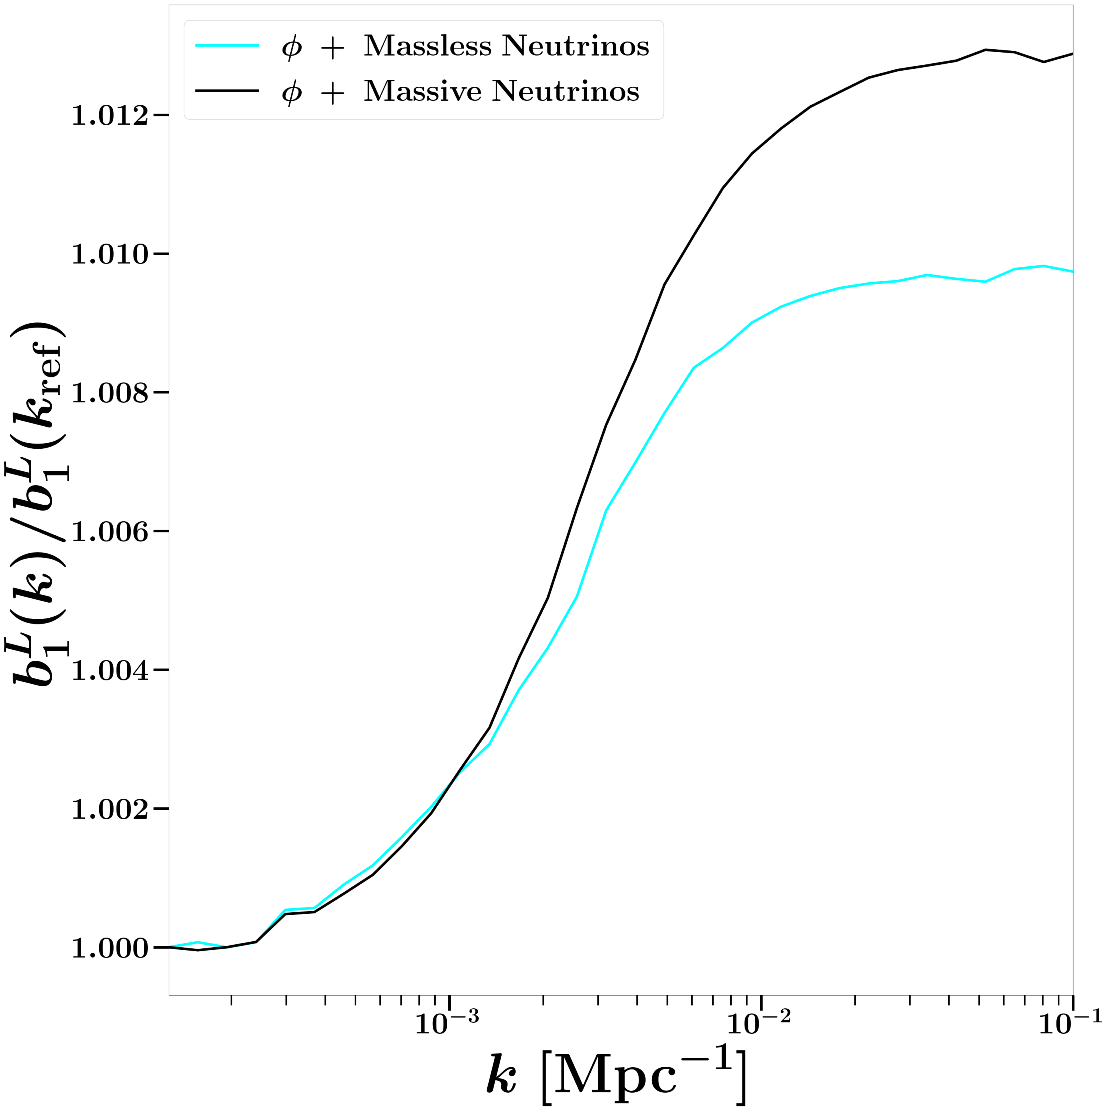

Figure 3.13: Same as Fig. 3.12 except in both cases we have an axion with a relic abundance of 1% of the CDM. We show the effect of replacing massless neutrinos with a massive degenerate hierarchy of three neutrino generations with total mass of 104 eV such that they also have a relic abundance of 1% of the CDM.

$m_{\nu,i} = \frac{1}{3}M_\nu = \frac{1}{3}\sum m_{\nu,i}$ and

$$\omega_\nu = \sum \omega_{\nu,i} = \frac{M_\nu}{93.14\ \text{eV}}\ . \qquad (3.48)$$

In Fig. 3.12, we see that these two signals are not fully degenerate. While the normalized (and absolute) bias is overall larger for the cosmology with an equivalent axion at large k, there is a crossing that occurs at smaller scales for larger neutrino masses. This suggests that scale dependent measurements of the halo bias have the potential to discriminate these two scenarios, and perhaps for more generic light but massive thermal relics (LiMRs). In both cases, a statistical analysis is necessary to determine whether near term surveys would be sensitive to these differences. We also want to verify whether the enhancement to the scale dependent bias introduced individually by massive neutrinos and ultra-light axions combines independently of each other. In Fig. 3.13, we show the case of a cosmology possessing both an ultra-light axion and massive neutrinos, each sector composing 1% of the CDM, and having $k_{\rm fs} = k_{\rm J}$. Note that the presence of both sectors increases the overall bias enhancement, allowing measurements of the halo bias to jointly constrain these two sectors. While axions may possess a wide range of $k_{\rm J}$, laboratory measurements of neutrino masses will set bound their corresponding range of $k_{\rm fs}$, introducing another mechanism by which bias measurements may differentiate these two species.

4

Inferring Black Hole Formation Channels with Eccentricity

4.1 INTRODUCTION

The initial detection of binary black hole pair (BBH) mergers by LIGO/VIRGO [4], have since been complemented by a host of similar detections [8,7] ushering in an era of gravitational wave astronomy and suggesting the existence of a sizeable BBH abundance. The development of our understanding of early Universe physics, the nature of dark matter, and astrophysical black hole production not only stand to benefit from this data, but are challenged by such findings as the existence of black holes of anomalous mass [6]. For these reasons, it is both timely and compelling to consider how we might better infer the formation channels of the black holes we observe in our detectors.

Various properties have been proposed as useful probes of the formation mechanisms of black holes, including black hole spin and mass [156,90,147]. This work is concerned with how the orbital eccentricity of BBHs may offer insight into black hole formation. Prior work has demonstrated how various static and dynamic processs can lead to populations of black holes with different distributions in their orbital eccentricity [157,181,177]. As a BBH approaches merger, orbital eccentricity will be quickly attenuated through the radiation of gravitational waves. While this means we should expect little residual eccentricity at Hz frequencies for stellar mass black holes typical to the detection bands of experiments like LIGO/VIRGO, it allows for considerable eccentricity at lower frequencies further from merger. Planned and proposed experiments in this frequency regime (e.g. LIGO, DECIGO) are well-positioned to measure this eccentricity and potentially lend insight into BBH formation mechanics.

Prior work has considered the observable consequences of eccentricity in mHz frequency detectors[176,179]. Amongst such effects are an increase in the perceived number of events in a given frequency window, and a suppression to the detector Signal-to-Noise Ratio (SNR). While these works have demonstrated the detection response for any single eccentricity of BBH, the effects of a BBH population possessing a distribution in eccentricity have not yet been studied. In this work, we consider the observational consequences of introducing a distribution in eccentricity to a BBH population. We will limit the eccentricity distributions studied to the case single-peaked distributions. For such cases, there are two qualitative changes which occur with respect to a population of fixed eccentricity: a widening about some peak value of eccentricity, and a shifting of the location of the peak value of eccentricity. The result of each effect is studied.

In Section 4.2, we introduce the dynamics of an eccentric BBH. In Section 4.4, we forecast the sensitivity of a LISA-like detector to the presence of particular populations of eccentric BBHs. In Section 4.5, we estimate the extent to which a LISA-like detector can discern populations of BBHs characterized by their eccentricity distributions. We conclude in Section 4.6 and offer on outlook for how gravitational wave detectors in other frequency bands might offer complimentary findings.

4.2 Eccentric Binary Dynamics and Evolution

Only for sufficiently large metric perturbations is a higher-order general relativistic description necessary to model the gravitational signature of binaries in the detectors we consider. For mHz range gravitational wave signals, typical binary systems are far from merger and such considerations are not necessary. To this effect, we shall only consider the quadrupole gravitational wave emission of the signal, modeled to the post-Newtonian order in General Relativity.

We describe a BBH, composed of masses m_1 and m_2, in terms of its total mass

$$m \equiv m_1 + m_2, \tag{4.1}$$

its reduced mass

$$\mu \equiv \frac{m_1 m_2}{m}, \tag{4.2}$$

its chirp mass

$$m_c \equiv \frac{\mu^{3/5}}{m^{2/5}}, \tag{4.3}$$

the semi-major axis of its orbit, a, and the orbital eccentricity, e. A non-circular binary with $e > 0$

will tend to circularize through emission of higher harmonic quadrupole radiation. As the circular orbit is the lowest energy configuration given some a, eccentric binaries will take a longer amount of time to merge in comparison to their circular counterparts. To the post-Newtonian level, these dynamics are described by the Peters' equations [169]

$$\frac{da}{dt} = -\frac{64}{5}\frac{G^3\mu m^2}{c^5 a^3}\frac{1+\frac{73}{24}e^2+\frac{37}{96}e^4}{(1-e^2)^{7/2}}, \tag{4.4}$$

$$\frac{de}{dt} = -\frac{304}{15}\frac{G^3\mu m^2}{c^5 a^4}\frac{e(1+\frac{121}{304}e^2)}{(1-e^2)^{5/2}}\,. \tag{4.5}$$

This set of equations fixes the relationship $a(e)$ in the evolution of a binary system. Ignoring corrections which occur near the merger, they can be used to estimate the time to merger by evolving until a is of order the black hole radius

$$\begin{aligned} t_{\mathrm{merge}} \approx\ & \tfrac{5}{256}\tfrac{c^5}{\mu m^2 G^3} \\ & \times \left(\tfrac{(1+e)^{2\gamma/3}G^{1/3}m^{1/3}}{(1-e^2)(f_p\pi)^{2/3}}\right)^4 (1-e^2)^{7/2}\,. \end{aligned} \tag{4.6}$$

Observing that the e dependent term, and hence the lifetime, in Eq. 4.6 grows monotonically with e and, as noted earlier. Further, we can describe the peak frequency of quadrupole emission by [178,176]

$$f_p \approx \frac{\sqrt{Gm}(1+e)^\gamma}{\pi(a(1-e^2))^{3/2}]}, \tag{4.7}$$

with $\gamma = 1.1954$, which departs from the $e = 0$ relationship $f_p = 2f_{\mathrm{orb}}$ as eccentricity grows. While emission also occurs at higher harmonics of this frequency, in this work we only consider the signal of a single BBH to be comprised of the emission at f_p - a more sophisticated analysis might consider the expected enhancement due to simultaneous observation of the emission at multiple frequencies. Combining Eq. 4.4 and Eq. 4.5 with Eq. 4.7, we see that specifying the mass, frequency, and

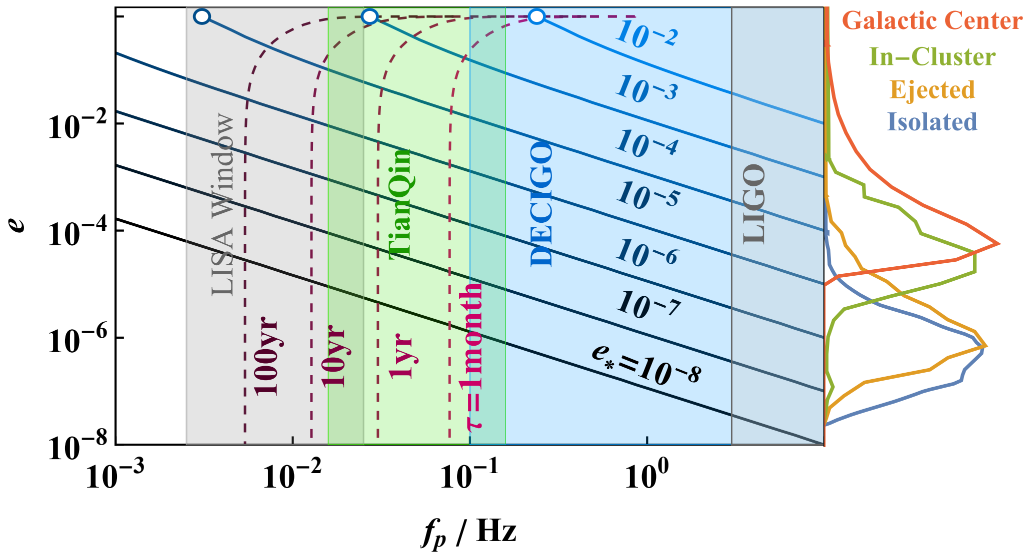

Figure 4.1: Evolution of black hole binary eccentricity as a function of peak quadrupole emission frequency. The evolution for several choices of $e_* \equiv e(f_p = 10\ \mathrm{Hz})$ is shown - for each, the lower limit of f_p is indicated by a circle. Points to the right of the shown magenta lines will merge within the indicated time period. Along the right edge, drawn at $f_p = 10\ \mathrm{Hz}$, are the e_* distributions corresponding to four different formation channels.

eccentricity of a binary at a fixed point in time determines the evolution of all physical properties at all times. Alternatively, we can use these same relations to describe the frequency evolution (i.e. chirping) of the BBH with time

$$\frac{dt}{df_p} = \frac{5c^5}{96\pi^{8/3}}(Gm_c)^{-5/3} f_p^{-11/3} \mathcal{F}(e), \tag{4.8}$$

$$\mathcal{F}(e) \equiv \begin{array}{c} \frac{(1+e)^{8\gamma/3-1/2}}{(1-e)^{3/2}}\Big((1+e)(1+\frac{7}{8}e^2) \\ -\frac{\gamma}{288}e(304+121e^2)\Big)^{-1} \end{array}. \tag{4.9}$$

Noting that $\mathcal{F}(e) \to 1$ as $e \to 0$, we see that the $\mathcal{F}(e)$ acts as a suppression to the chirping (df_p/dt) of a circular BBH pair.

In Fig. 4.1 we demonstrate the evolution of $e(f_p)$ for various choices of $e_* \equiv e(f_p = 10\ \mathrm{Hz})$. Note that for any choice of e_* there will exist a lower bound on f_p corresponding to $e = 1$. As an example, it is worth noting that sufficiently eccentric BBHs with $e_* \gtrsim 10^{-3}$ will never radiate in the LISA window but will still merge in the LIGO window; in such a situation, the difference on BBH counts in LISA versus LIGO could be used to infer the presence of a population of highly eccentric binaries. Likewise, while all BBHs with $e_* \lesssim 10^{-3}$ will produce a signal in LISA and eventually enter the LIGO band, only those which initially appear in the LISA band with sufficiently high frequency will make this transition in a period of time reasonable for a follow up observation in LIGO. The magenta bands in Fig. 4.1 indicate this lower bound on f_p for various choices of follow-up time.

In this work, we are not concerned with the evolution of BBHs at a single e_* but rather for a population distributed at e_*. In Fig. 4.1, the upper limit of f_p is drawn at 10 Hz, the same frequency at which e_* is defined. So we can gain a qualitative understanding of how various e_* populations evolve through a potential detector landscape by drawing the e_* distribution at 10 Hz and tracing lines of constant e_* to the left of the distribution. We show predictions for the e_* distributions generated via different astrophysical mechanisms (Isolated, Ejected, In-Cluster, Galactic Center) [157,181,177]. As an example, we consider the evolution of an Isolated distributions, such as might describe a primordial black hole population existing in a void of large scale structure. In this case, we see that the majority of the population will possess an $e_* \approx 10^{-6}$; will appear in both LISA and LIGO; and will merge within 10 years for binaries appearing in LISA with $f_p \gtrsim 0.02$ Hz.

4.3 Observing Eccentric Binaries

We wish to consider the role that eccentricity plays in modifying the signal entering a gravitational wave detector, as well as in the noise response of the detector to such a signal. We begin by noting there are a number of population characteristics which we assume are independent of the e_* distribution. In the case of a static Universe, we assume that the distribution in spatial position and mass

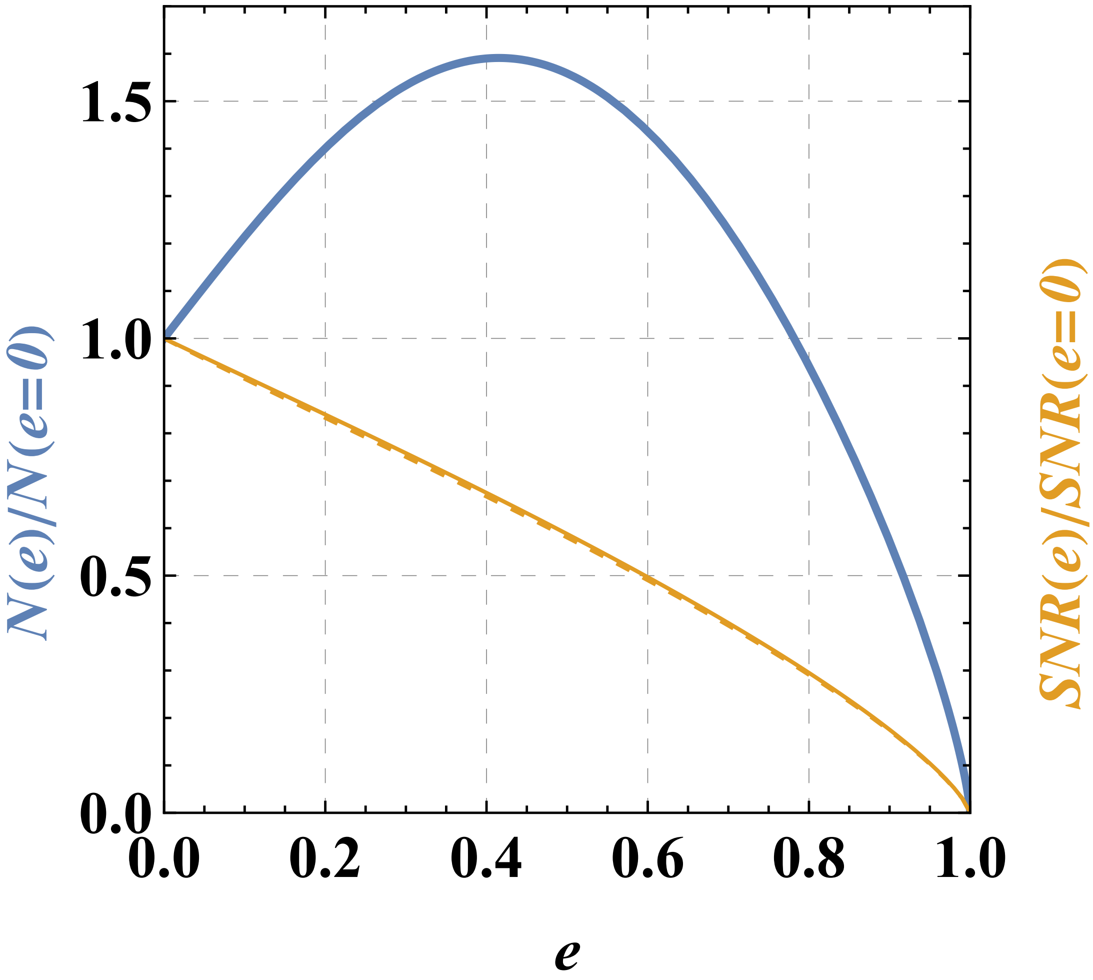

Figure 4.2: The effect on number density and signal-to-noise ratio in LISA as a function of binary eccentricity at fixed frequency. Shown in blue is the enhancement to the observed number density of binaries in LISA relative to circular binaries. Shown in yellow (dashed) is the suppression to the LISA (DECIGO) SNR relative to circular binaries. The nearly identical effect on SNR in LISA/DECIGO indicates that suppression is driven by the change in signal rather than difference in the detector noise strains.

do not correlate with the eccentricity distribution of the population. To this end, we utilize the mass distribution inferred by LIGO to describe BBH populations of any eccentricity [3]

$$p(m_1) \propto m_1^{-2.3}\,. \tag{4.10}$$

Likewise, we assume that binaries are uniformly distributed spatially

$$p(r) \propto 4\pi r^2 \tag{4.11}$$

and that the local merger rate is a constant

$$\mathcal{R} = \frac{dn}{dt} = \text{constant} \tag{4.12}$$

We note, however, that the same mechanisms which generate the eccentricity distributions of interest may also modify these distributions. For example, we do not expect populations produced by astrophysical processes near the galactic center to populate many BBH pairs at extragalactic distances. Such considerations would certainly be necessary to draw conclusions about the underlying formation channels, but as a stepping stone we only consider here the modifications to the signal and noise that eccentricity introduces.

The first place that eccentricity enters the signal is in the f_p distribution of binaries. As we are always concerned with the number of binaries observed over fixed time intervals, the likelihood of a BBH possessing a particular f_p is given by Eq. 4.8

$$p(f_p) \propto \frac{dt}{df_p}\,. \tag{4.13}$$

The only remaining population parameter is e_*, whose distribution $p(e_*)$ will be provided by one of the underlying formation mechanisms shown in Fig. 4.1.

Combining these distributions provides a measure of the expected number of events in a given time interval

$$N_{\text{events}} = \int \mathcal{R} p(m_c) p(r) p(f_p, e(f_p)) p(e_*) dm_c dr df_p de_* \tag{4.14}$$

So we see that eccentricity affects our signal - the event count - in three ways: (1) dt/df_p sees an enhancement due to eccentricity given by Eq. 4.9, (2) the $f_p(t)$ evolution is set by the value of e_*, and (3) the likelihood of a particular e_* is weighted by $p(e_*)$.

We adopt a simplified model for the SNR in the presence of chirping binaries

$$\varrho(f_p, e)^2 = 4 \int dt \frac{h_c^2(f_p(t), e=0)}{S_N(f_p(t))}(1-e(t))^{3/2} \tag{4.15}$$

shown to be a good approximation [179] to the true SNR

$$\varrho^2 = 4 \sum_n \int dt \frac{h_n^2(f_n(t))}{S_N(f_n(t))} \tag{4.16}$$

which involves a summation over the n harmonic components of the gravitational wave emission. Here, $S_N(f)$ is the noise strain and $h_c^2 = \Sigma h_n^2$ is the signal strain averaged over an orbital period. Noting that $\varrho \propto (1-e)^{3/4}$ and that $N \propto r^3$, we see that there should be an suppression to the expected number of counts with respect to the circular case given by $(1-e)^{9/4}$. Combining this suppresion term with the enhancement given by $\mathcal{F}(e)$, in Fig. 4.2 we demonstrate the overall effect to the event count with respect to circular BBHs that results from introducing eccentricity to non-chirping BBHs. We also show the suppression to the SNR for the cases of both LISA and DECIGO to demonstrate that this effect is driven by the $f_p(e)$ dynamics rather than the specific shape of the detector noise curve.

The competition of these two effects will depend on which eccentricities we assume our detectors are able to observe, which is driven by the availability of gravitational wave templates necessary to detect eccentric BBH gravitational waves. This will introduce a cut into our event count over the parameter space wherever the eccentricity is larger than is observationally possible

$$\Theta_{e_{\rm cut}}(e) = \Theta(e_{\rm cut} - e)\,. \tag{4.17}$$

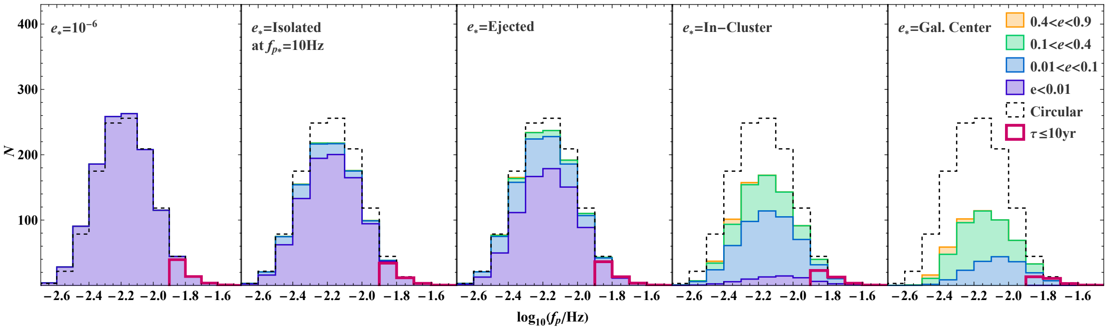

Figure 4.3: Number of observed counts in LISA for fixed, Isolated, Ejected, In-Cluster, and Galactic Center e_* distributions. Dashed line indicates expected counts for a perfectly circular distribution. While a fixed e_* distribution would exhibit a lower bound in f_p, distributions in e_* distribute counts across all frequency bins for all choices of e_{cut}. Magenta lines indicate events which will merge within 10 years and have $e_{\text{cut}} = 0.9$ - suggesting events which can have an observable Hz range terrestrial observation follow-up.

Likewise, the calculated SNR will determine whether a signal is observationally viable. We impose SNR > 8 over 10 years of LISA observation as the criteria by which to consider events observable

$$\Theta_{\text{SNR}}(r, f_p, m_c, e_*) = \Theta(\varrho(r, f_p, m_c, e_*) - 8)\,. \tag{4.18}$$

Combining these terms, the observable number count is given by

$$\begin{aligned} N_{\text{events}} = \; & \int \mathcal{R} p(m_c) p(r) p(f_p, e(f_p)) p(e_*) \\ & \times \Theta_{\text{SNR}} \Theta_{e_{\text{cut}}} dm_c dr df_p de_*\,. \end{aligned} \tag{4.19}$$

In Fig. 4.3 we show the expected number of observable events in LISA over a 10 year observation period for five choices of the BBH eccentricity distribution. In each case, we also highlight the subset of events which will merge within 10 years as an example of candidates that could potentially be followed-up with a LIGO-like detection. For reference, we compare with what the expected number of counts would be assuming a BBH population of only circular binaries, demarked by the black dashed line. In the leftmost panel, we show the result for a fixed - and relatively small - eccentricity to indicate the close matching the circular case. Progressing to the right, we consider the Isolated distribution, which is peaked at $e_* \approx 10^{-6}$ but extends to both higher and lower eccentricities. As indicated by the four different colored histograms, we begin to lose sensitivity to events as we impose progressively lower bounds on $e_{\rm cut}$, to the point that no events would be observed in the case of a Galactic Center distribution without eccentricity templates of $e_{\rm cut} \geqq 0.01$. Interestingly, in this situation, a mismatch between the LIGO merger event rate and LISA merger event rate would indicate the existence of an eccentric sub-population, despite having no direct observations in the LISA detector.

Another important consequence of transitioning from fixed e_* to a distribution is that there is no longer a single lower bound on f_p as each choice of e_* will produce a different bound. Sensitivity to the shape of the e_* distribution is contained in how the counts per frequency bin shift as we move $e_{\rm cut}$, most notably evidenced by the shifting the Galactic Center distribution counts. In general, we see that the overall counts give an indication of where the peak of the distribution sits, the difference between successive $e_{\rm cuts}$ indicating how far the distribution extends to higher eccentricities. While the difference between successive $e_{\rm cut}$ grows larger for more eccentric distributions, the overall number of counts is lower. So we see that there is a tradeoff between the signal uniqueness and statistical significance of the signal which will be expanded on in the next section.

4.4 ECCENTRICITY BREAKING OF OBSERVATIONAL DEGENERACY

Considering the Isolated and Ejected panels of Fig. 4.3, we notice that there is a very similar number of events at each frequency bin. In this case, we expect there to be a significant amount of degeneracy between these two models in the dataset. Glancing at the shape of these distributions in Fig. 4.1, this should not be surprising due to the similarity in distribution shape. It is useful to ask how such degeneracies in the LISA dataset can be broken through the use of choices in $e_{\rm cut}$. To answer this question, we consider a model for the distribution of e_* given by

$$f(e_*) = A_i f_i(e_*) + A_j f_j(e_*), \tag{4.20}$$

such that

$$A_i + A_j = 1\,. \tag{4.21}$$

To quantify the observational degeneracy between models in our dataset, we proceed to construct a Fisher information matrix for the two parameter model A_i, A_j where i, j correspond to the Isolated, Ejected, In-Cluster, and Galactic Center distributions [94]. The elements of the Fisher matrix are defined by

$$F_{ij} \equiv \sum_{f_p, k} \sqrt{N_k} \frac{\partial N_k}{\partial A_i} \frac{\partial N_k}{\partial A_j}, \tag{4.22}$$

where the index k runs across each of the f_p bins. We derive three separate F_{ij}, one for each of $e_{\mathrm{cut}} = 1.0, 0.4, 0.1$. For each of these, we consider three different fiducial choices of the ratio

$$\mathcal{R} \equiv \frac{A_i}{A_j}\,. \tag{4.23}$$

So that in total, there are nine versions F_{ij}. As a base case, we will consider how the Isolated distribution is degenerate with the other, more eccentric distributions. We show the resulting covariance contours between A_{Isolated} and A_j in Fig. 4.4.

Considering each e_{cut} as an independent dataset, we generally expect that $e_{\mathrm{cut}} = 1$, containing the must counts, would always offer the greatest constraining power. Indeed, when compared to distributions of substantially higher eccentricity than the Isolated channel this is the case for choices of $\mathcal{R}$. However, the constraint between the relatively similar Ejected and Isolated distributions exhibit the opposite behavior in the case where $\mathcal{R}$ takes on larger values.

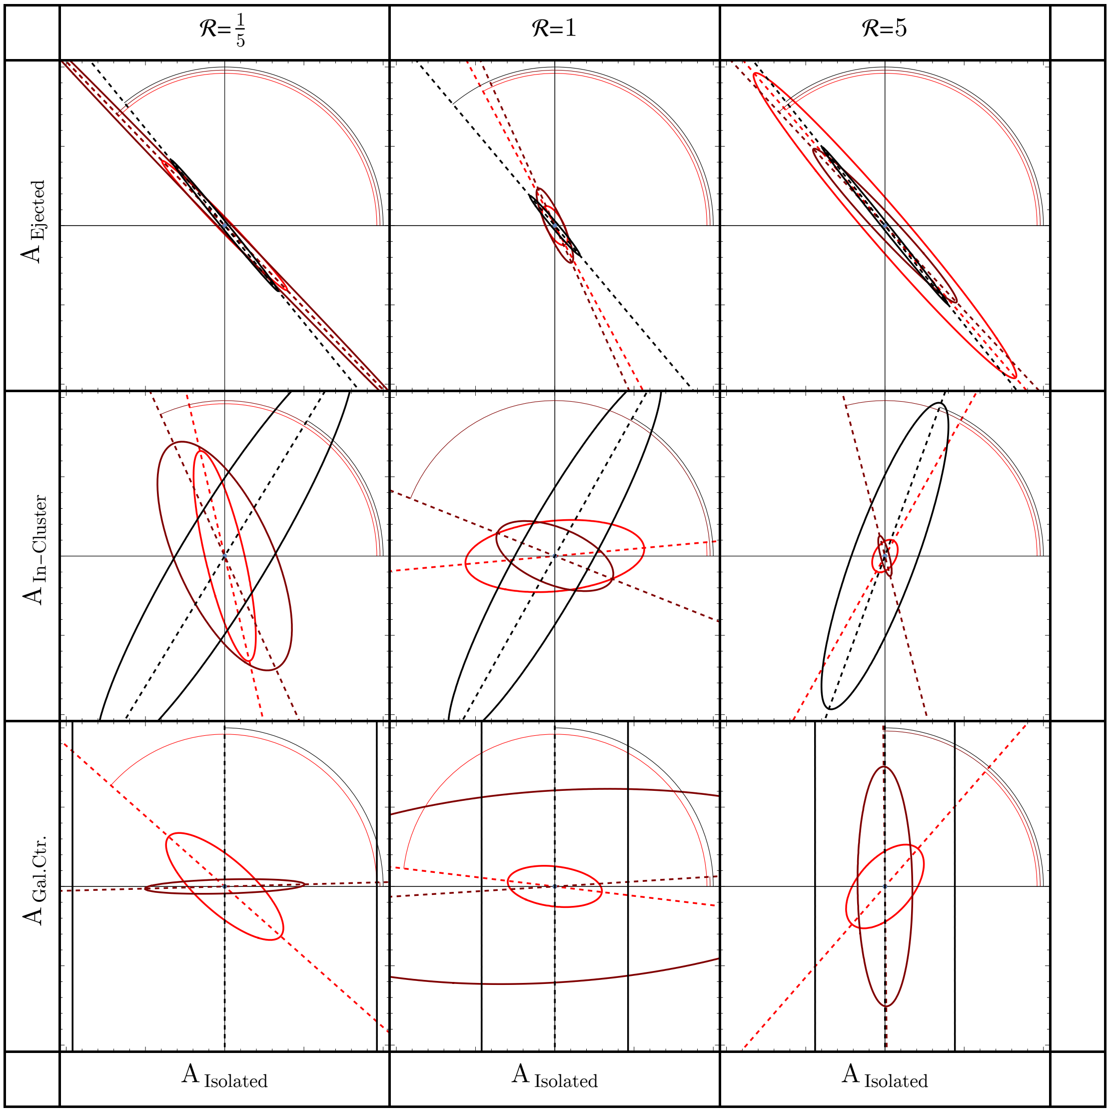

Figure 4.4: Forecasted 3σ constraint contours on coefficients of binary e_* distributions for a LISA-like experiment. A_i represents the abundance of e_* distribution i. In all cases, only two e_* distributions i, j are considered at a time such that $A_i + A_j = 1$. Three different abundance ratios $\mathcal{R} \equiv A_i/A_j$ are considered, as well as three choices of maximum observable eccentricity $e_{\text{cut}} = 0.1$ (black), 0.4 (dark red), and 1.0 (light red). All panels span ± 0.1 from the center point. Generally, ellipse size indicates constraining power, while ellipse angle indicates parameter degeneracy (e.g. the Isolated and Ejected distributions are highly degenerate).

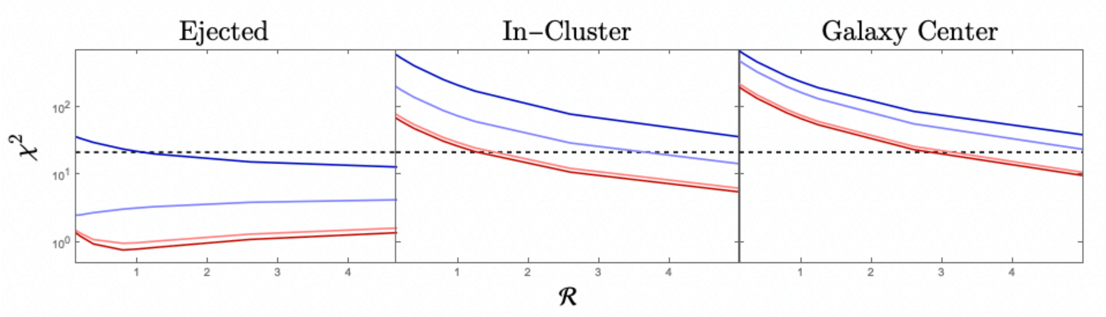

Figure 4.5: The χ^2 measured by LISA assuming a true model corresponding to the Isolated e_* distribution, but measuring the Ejected, In-Cluster, Galaxy Center e_* distributions, and assuming constraining power is systematically limited over 12 independent frequency bins. Shown are four choices of $e_{\rm cut} = 0.9$ (dark red), 0.4 (light red), 0.1 (light blue), 0.01 (dark blue). The black dashed line indicates the χ^2 value corresponding to a p-value of 0.05. Plotted for $\mathcal{R} \in [0.125, 5.0]$.

Of most interest here are cases where we see that different choices of $e_{\rm cut}$ provide powerful breaking of the dataset degeneracy. In the case of the Isolated and Ejected channels, the similarity of these two distributions ensures that they are highly degenerate - evidenced by the $\approx -45^\circ$ angle of every contour. This suggests, in the such cases, there is little information gained by pursuing higher $e_{\rm cut}$ in the data. This is sharply contrasted by the cases of the In-Cluster and Galactic Center. Here, we see substantial amounts of orthogonality in the contours - suggesting a combined constraint using data from both $e_{\rm cut}$ would provide a dramatic improvement over a single cut dataset.

4.5 Distinguishing Populations with LISA

We now wish to consider how the dataset sensitivity to the population coefficients explored in Sec. 4.4 translates to the significance to which an experiment like LISA might be able to distinguish the proportion of a BBH population deriving from different eccentricity distributions. As in Sec. 4.4, we consider the Isolated distribution as a base case. We then ask the question - how much of the BBH population needs to be derived from a different eccentricity distribution before LISA can rule out the presence of an entirely Isolated distribution with statistical significance. We will again quantify the portion of the population derived from a non-Isolated distribution by the parameter $\mathcal{R}$. We then construct a simple chi-square statistic for the dataset:

$$\chi^2 = \sum_{f_p,k} \frac{(N_{\text{Isolated}} - N_{\text{observed}})^2}{N_{\text{Isolated}}} \,. \tag{4.24}$$

For each choice of e_{cut}, we construct χ^2 as a function of $\mathcal{R}$ as shown in Fig. 4.5. We then calculate the χ^2 corresponding to a p-value of 0.05 for this data, indicated by the dashed line, with value of χ^2 above this threshold denoting regimes where LISA can significantly claim the presence of an eccentricity distribution that is not just Isolated. We see that in all cases, the most efficient metric for significantly distinguishing different populations remains the lowest $e_{\text{cut}} = 0.01$ but that $\mathcal{O}(1)$ improvements can be had by considering joint e_{cut} constraints.

4.6 Conclusions

We have shown how different choices of eccentricity distribution for a population of BBHs can produce substantially different numbers of counts per frequency bin in a gravitational wave detector like LISA. While the ability to observe mergers of higher max eccentricity is capable of strongly breaking dataset degeneracies between models of mixed distributions, the distinguishing power of detectors like LISA is still driven by data from of low eccentricity events.

5

Conclusion

This dissertation details several methods by which to infer the properties of dark sector physics exclusively through the gravitational interaction of those sectors with the rest of the Standard Model. We have shown that phenomenological models of Dark Matter which comprise a massive particle in contact with the SM at early times manifests a distinct energy scale in the cosmology known as the free streaming scale. Such Light but Massive Relic particles not only introduce a suppression to the linear matter perturbations below their free streaming scale, but additionally introduce a scale dependent feature in the non-linear growth of structure due to spherical gravitational collapse. To take advantage of the plethora of upcoming galaxy survey datasets to study DM, it is requisite that models describe the influence of DM on the relationship between halo overdensities and the linear cosmological perturbations - a relation known as the galaxy bias. We have shown that LiMR DM introduces a step like enhancement to the scale dependent halo bias which partially compensates for the suppression introduced to the linear matter power spectrum. Due to the non-linearities of structure collapse, this Growth-Induced Scale Dependent Bias feature is non-degenerate with the DM effect on the matter power spectrum. Therefore, we both learn additional information about DM by measuring the halo bias as well as infer more accurate constraints on DM properties. We demonstrate through forecasting how well the SDSS-BOSS, DESI, and Euclid galaxy surveys combined with the Planck 2018 and CMB-S4 CMB surveys will be able to constrain the mass, temperature and effective internal degrees of freedom (which can be recast as a constraint on the relic abundance) of LiMRs. Over the region which is physically allowable, we find that most masses above ≈ 1 eV are excluded at the 3σ significance. As a specific example, we then applied the LiMR analysis to the case of massive neutrinos. While we found that accounting for the GISDB for massive neutrinos induced a 1σ shift in the inferred maximum likelihood of the neutrino mass, this shift vanishes if we marginalize over the redshift dependence of the halo bias, which must be done if we do not have good information constraining this dependence. The addition of the GISDB to the case of massive neutrinos provided negligible enhancement to how well constrained the neutrino mass is about its maximum likelihood value. We then considered the case of an alternative DM phenomenology comprised of an ultralight axion field. We demonstrated that, like in the case of a LiMR, and ULA introduces an energy scale into the cosmology which is determined by its mass. In this case, the ULA sets when the field begins to oscillate about the minimum of its potential. Unlike the LiMR case, the ULA introduces a second scale which is set by is macroscopic wavelength called the Jeans Scale. These two effects combine to produce a distinct feature of ULA physics in both the linear cosmology as well as the non-linear halo bias of the same scale as the degree to which the comprise the relic DM fraction of the Universe today $\approx$ 1% level.

It is known that different dark sector theories affect the evolution and characteristics of astrophysical objects today. It is then reasonable to assume that astrophysical processes dependent on galaxy dynamics may be directly or indirectly affected by dark sector physics. One such process is the production of black holes. It is then relevant to ask to what degree we may be able to distinguish the formation channel of black holes. It is even more relevant due to the possibility of the primordial black hole production that could account for a portion of the relic dark matter. In this dissertation, we considered how different populations of black holes might be distinguished from each other using the persistent orbital eccentricity of Binary Black Hole pairs that form within the population. We demonstrated that a mHz frequency detector like LISA can strongly break degeneracies between population statistics with access to $\mathcal{O}(0.1)$ level eccentricity templates for gravitational waveforms. Though we show that the constraining power of such observations is still driven by higher sensitivity to low eccentricity events, combining observations across gravitational wave experiments of different frequency regimes enables inference of population eccentricity without having access to high eccentricity templates. In all these inquiries, we find that competitive bounds on a wide array of DM phenomenologies can be set by gravitational interactions alone and that such gravitational constraints may be sufficient to test BSM physics without analyzing the consequences of other interaction pathways.

References

[1] Abazajian, K. et al. (2019). CMB-S4 Decadal Survey APC White Paper. Bull. Am. Astron. Soc., 51(7), 209.

[2] Abazajian, K. N. et al. (2016). CMB-S4 Science Book, First Edition.

[3] Abbott, B. P. et al. (2019a). Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo. Astrophys. J. Lett., 882(2), L24.

[4] Abbott, B. P. et al. (2019b). GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. Phys. Rev. X, 9(3), 031040.

[5] Abbott, L. F. & Sikivie, P. (1983). A Cosmological Bound on the Invisible Axion. Phys. Lett. B, 120, 133–136.

[6] Abbott, R. et al. (2020). GW190521: A Binary Black Hole Merger with a Total Mass of $150 M_{\odot}$. Phys. Rev. Lett., 125(10), 101102.

[7] Abbott, R. et al. (2021). GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run.

[8] Abbott, R. et al. (2023). Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3. Phys. Rev. X, 13(1), 011048.

[9] Abbott, T. et al. (2005). The Dark Energy Survey.

[10] Abbott, T. M. C. et al. (2018). Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. Phys. Rev., D98(4), 043526.

[11] Abe, K. et al. (2011). Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam. Phys. Rev. Lett., 107, 041801.

[12] Abe, Y. et al. (2012). Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment. Phys. Rev. Lett., 108, 131801.

[13] Abel, S., Goodsell, M., Jaeckel, J., Khoze, V., & Ringwald, A. (2008). Kinetic Mixing of the Photon with Hidden U(1)s in String Phenomenology. JHEP, 07, 124.

[14] Adams, C. B. et al. (2022). Axion Dark Matter. In 2022 Snowmass Summer Study.

[15] Ageron, M. et al. (2011). ANTARES: the first undersea neutrino telescope. Nucl. Instrum. Meth., A656, 11–38.

[16] Aghamousa, A. et al. (2016). The DESI Experiment Part I: Science,Targeting, and Survey Design.

[17] Aghanim, N. et al. (2020). Planck 2018 results. VI. Cosmological parameters. Astron. Astrophys., 641, A6. [Erratum: Astron.Astrophys. 652, C4 (2021)].

[18] Agrawal, P., Cyr-Racine, F.-Y., Pinner, D., & Randall, L. (2019). Rock 'n' Roll Solutions to the Hubble Tension.

[19] Aguilar-Arevalo, A. A. et al. (2013). Improved Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment. Phys. Rev. Lett., 110, 161801.

[20] Ahn, J. K. et al. (2012). Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. Phys. Rev. Lett., 108, 191802.

[21] Aker, M. et al. (2019). Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN. Phys. Rev. Lett., 123(22), 221802.

[22] Akita, K. & Yamaguchi, M. (2020). A precision calculation of relic neutrino decoupling. JCAP, 08, 012.

[23] Alam, S. et al. (2017). The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample. Mon. Not. Roy. Astron. Soc., 470(3), 2617–2652.

[24] Alcock, C. & Paczynski, B. (1979). An evolution free test for non-zero cosmological constant. Nature, 281, 358–359.

[25] Ali-Haimoud, Y. & Bird, S. (2012). An efficient implementation of massive neutrinos in non-linear structure formation simulations. Mon. Not. Roy. Astron. Soc., 428, 3375–3389.

[26] Allison, R., Caucal, P., Calabrese, E., Dunkley, J., & Louis, T. (2015). Towards a cosmological neutrino mass detection. Phys. Rev., D92(12), 123535.

[27] Amendola, L. & Barbieri, R. (2006). Dark matter from an ultra-light pseudo-Goldsone-boson. Phys. Lett. B, 642, 192–196.

[28] Amendola, L. et al. (2018). Cosmology and fundamental physics with the Euclid satellite. Living Rev. Rel., 21(1), 2.

[29] An, F. et al. (2016). Neutrino Physics with JUNO. J. Phys., G43(3), 030401.

[30] Archidiacono, M., Hannestad, S., & Lesgourgues, J. (2020). What will it take to measure individual neutrino mass states using cosmology? JCAP, 09, 021.

[31] Arkani-Hamed, N. & Weiner, N. (2008). LHC Signals for a SuperUnified Theory of Dark Matter. JHEP, 12, 104.

[32] Arvanitaki, A., Dimopoulos, S., Dubovsky, S., Kaloper, N., & March-Russell, J. (2010). String Axiverse. Phys. Rev. D, 81, 123530.

[33] Asgari, M., Taylor, A., Joachimi, B., & Kitching, T. D. (2018). Flat-SkyPseudo-Cls analysis for weak gravitational lensing. Mon. Not. Roy. Astron. Soc., 479(1), 454–477.

[34] Ayres, D. S. et al. (2004). NOvA: Proposal to Build a 30 Kiloton Off-Axis Detector to Study $\nu_\mu \to \nu_e$ Oscillations in the NuMI Beamline.

[35] Baker, C. A. et al. (2006). An Improved experimental limit on the electric dipole moment of the neutron. Phys. Rev. Lett., 97, 131801.

[36] Banerjee, A., Jain, B., Dalal, N., & Shelton, J. (2018). Tests of Neutrino and Dark Radiation Models from Galaxy and CMB surveys. JCAP, 01, 022.

[37] Bashinsky, S. & Seljak, U. (2004a). Neutrino perturbations in CMB anisotropy and matter clustering. Phys. Rev., D69, 083002.

[38] Bashinsky, S. & Seljak, U. c. v. (2004b). Signatures of relativistic neutrinos in cmb anisotropy and matter clustering. Phys. Rev. D, 69, 083002.

[39] Bauer, J. B., Marsh, D. J. E., Hložek, R., Padmanabhan, H., & Laguë, A. (2020). Intensity Mapping as a Probe of Axion Dark Matter. Mon. Not. Roy. Astron. Soc., 500(3), 3162–3177.

[40] Baumann, D., Green, D., Meyers, J., & Wallisch, B. (2016). Phases of New Physics in the CMB. JCAP, 1601, 007.

[41] Baumann, D., Green, D., & Wallisch, B. (2018). Searching for light relics with large-scale structure. JCAP, 1808, 029.

[42] Baumann, D., Green, D., & Zaldarriaga, M. (2017). Phases of New Physics in the BAO Spectrum. JCAP, 1711, 007.

[43] Baumann, D. D., Beutler, F., Flauger, R., Green, D. R., Slosar, A. z., Vargas-Magaña, M., Wallisch, B., & Yèche, C. (2019). First constraint on the neutrino-induced phase shift in the spectrum of baryon acoustic oscillations. Nature Phys., 15, 465–469.

[44] Baur, J., Palanque-Delabrouille, N., Yèche, C., Boyarsky, A., Ruchayskiy, O., Armengaud, E., & Lesgourgues, J. (2017). Constraints from Ly-α forests on non-thermal dark matter including resonantly-produced sterile neutrinos. JCAP, 1712(12), 013.

[45] Benakli, K., Chen, Y., Dudas, E., & Mambrini, Y. (2017). Minimal model of gravitino dark matter. Phys. Rev., D95(9), 095002.

[46] Beranek, T., Merkel, H., & Vanderhaeghen, M. (2013). Theoretical framework to analyze searches for hidden light gauge bosons in electron scattering fixed target experiments. Phys. Rev., D88, 015032.

[47] Biagetti, M., Desjacques, V., Kehagias, A., & Riotto, A. (2014). Nonlocal halo bias with and without massive neutrinos. Phys. Rev., D90(4), 045022.

[48] Bird, S., Ali-Haïmoud, Y., Feng, Y., & Liu, J. (2018). An Efficient and Accurate Hybrid Method for Simulating Non-Linear Neutrino Structure. Mon. Not. Roy. Astron. Soc., 481(2), 1486–1500.

[49] Blas, D., Lesgourgues, J., & Tram, T. (2011). The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes. JCAP, 1107, 034.

[50] Boyarsky, A., Lesgourgues, J., Ruchayskiy, O., & Viel, M. (2009). Lyman-alpha constraints on warm and on warm-plus-cold dark matter models. JCAP, 0905, 012.

[51] Boyle, A. & Komatsu, E. (2018). Deconstructing the neutrino mass constraint from galaxy redshift surveys. JCAP, 1803, 035.

[52] Brinckmann, T., Hooper, D. C., Archidiacono, M., Lesgourgues, J., & Sprenger, T. (2019). The promising future of a robust cosmological neutrino mass measurement. JCAP, 01, 059.

[53] Brinckmann, T. & Lesgourgues, J. (2019). MontePython 3: boosted MCMC sampler and other features. Phys. Dark Univ., 24, 100260.

[54] Brust, C., Kaplan, D. E., & Walters, M. T. (2013). New Light Species and the CMB. JHEP, 12, 058.

[55] Bull, P., Ferreira, P. G., Patel, P., & Santos, M. G. (2015). Late-time cosmology with 21cm intensity mapping experiments. Astrophys. J., 803(1), 21.

[56] Butler, J. N. et al. (2023). Report of the 2021 U.S. Community Study on the Future of Particle Physics (Snowmass 2021).

[57] Castorina, E., Carbone, C., Bel, J., Sefusatti, E., & Dolag, K. (2015). DEMNUni: The clustering of large-scale structures in the presence of massive neutrinos. JCAP, 1507(07), 043.

[58] Castorina, E., Sefusatti, E., Sheth, R. K., Villaescusa-Navarro, F., & Viel, M. (2014). Cosmology with massive neutrinos II: on the universality of the halo mass function and bias. JCAP, 1402, 049.

[59] Chan, H. Y. J., Ferreira, E. G. M., May, S., Hayashi, K., & Chiba, M. (2022). The diversity of core–halo structure in the fuzzy dark matter model. Mon. Not. Roy. Astron. Soc., 511(1), 943–952.

[60] Cheung, C., Elor, G., & Hall, L. (2011). Gravitino Freeze-In. Phys. Rev. D, 84, 115021.

[61] Cheung, K. & Yuan, T.-C. (2007). Hidden fermion as milli-charged dark matter in Stueckelberg Z- prime model. JHEP, 03, 120.

[62] Chiang, C.-T., Hu, W., Li, Y., & Loverde, M. (2018). Scale-dependent bias and bispectrum in neutrino separate universe simulations. Phys. Rev. D, 97(12), 123526.

[63] Chiang, C.-T., LoVerde, M., & Villaescusa-Navarro, F. (2019). First detection of scale-dependent linear halo bias in *N*-body simulations with massive neutrinos. Phys. Rev. Lett., 122(4), 041302.

[64] Cooray, A. R. (2004). Non-linear galaxy power spectrum and cosmological parameters. Mon. Not. Roy. Astron. Soc., 348, 250–260.

[65] Costanzi, M., Villaescusa-Navarro, F., Viel, M., Xia, J.-Q., Borgani, S., Castorina, E., & Sefusatti, E. (2013). Cosmology with massive neutrinos III: the halo mass function andan application to galaxy clusters. JCAP, 1312, 012.

[66] Covi, L., Kim, J. E., & Roszkowski, L. (1999). Axinos as cold dark matter. Phys. Rev. Lett., 82, 4180–4183.

[67] Crocce, M., Fosalba, P., Castander, F. J., & Gaztanaga, E. (2010). Simulating the Universe with MICE: The abundance of massive clusters. Mon. Not. Roy. Astron. Soc., 403, 1353–1367.

[68] Cuesta, A. J., Niro, V., & Verde, L. (2016). Neutrino mass limits: robust information from the power spectrum of galaxy surveys. Phys. Dark Univ., 13, 77–86.

[69] Cyburt, R. H., Fields, B. D., Olive, K. A., & Yeh, T.-H. (2016). Big Bang Nucleosynthesis: 2015. Rev. Mod. Phys., 88, 015004.

[70] Dawson, K. S., Schlegel, D. J., Ahn, C. P., Anderson, S. F., Aubourg, E., Bailey, S., Barkhouser, R. H., Bautista, J. E., Beifiori, A., Berlind, A. A., & et al. (2012). The baryon oscillation spectroscopic survey of sdss-iii. The Astronomical Journal, 145(1), 10.

[71] De Bernardis, F., Kitching, T. D., Heavens, A., & Melchiorri, A. (2009). Determining the Neutrino Mass Hierarchy with Cosmology. Phys. Rev. D, 80, 123509.

[72] de Salas, P. F. & Pastor, S. (2016). Relic neutrino decoupling with flavour oscillations revisited. JCAP, 07, 051.

[73] Dentler, M., Hernández-Cabezudo, A., Kopp, J., Machado, P. A. N., Maltoni, M., Martinez-Soler, I., & Schwetz, T. (2018). Updated Global Analysis of Neutrino Oscillations in the Presence of eV-Scale Sterile Neutrinos. JHEP, 08, 010.

[74] Dentler, M., Marsh, D. J. E., Hložek, R., Laguë, A., Rogers, K. K., & Grin, D. (2021). Fuzzy Dark Matter and the Dark Energy Survey Year 1 Data. Mon. Not. Roy. Astron. Soc., 515(4), 5646–5664.

[75] DePorzio, N., Xu, W. L., Muñoz, J. B., & Dvorkin, C. (2021). Finding eV-scale light relics with cosmological observables. Phys. Rev. D, 103(2), 023504.

[76] Deser, S. & Zumino, B. (1977). Broken Supersymmetry and Supergravity. Phys. Rev. Lett., 38, 1433–1436.

[77] Despali, G., Giocoli, C., Angulo, R. E., Tormen, G., Sheth, R. K., Baso, G., & Moscardini, L. (2016). The universality of the virial halo mass function and models for non-universality of other halo definitions. Mon. Not. Roy. Astron. Soc., 456(3), 2486–2504.

[78] Dicus, D. A., Kolb, E. W., Gleeson, A. M., Sudarshan, E. C. G., Teplitz, V. L., & Turner, M. S. (1982). Primordial nucleosynthesis including radiative, coulomb, and finite-temperature corrections to weak rates. Phys. Rev. D, 26, 2694–2706.

[79] Diehl, J. & Weller, J. (2021). Constraining ultra-light axions with galaxy cluster number counts. JCAP, 08, 004.

[80] Dine, M. & Fischler, W. (1983). The Not So Harmless Axion. Phys. Lett. B, 120, 137–141.

[81] Dodelson, S., Heitmann, K., Hirata, C., Honscheid, K., Roodman, A., Seljak, U., Slosar, A., & Trodden, M. (2016). Cosmic Visions Dark Energy: Science.

[82] Dvorkin, C. et al. (2019a). Neutrino Mass from Cosmology: Probing Physics Beyond the Standard Model.

[83] Dvorkin, C., Lin, T., & Schutz, K. (2019b). Making dark matter out of light: freeze-in from plasma effects. Phys. Rev. D, 99(11), 115009.

[84] Essig, R., Schuster, P., & Toro, N. (2009). Probing Dark Forces and Light Hidden Sectors at Low-Energy e+e- Colliders. Phys. Rev., D80, 015003.

[85] Farren, G. S., Grin, D., Jaffe, A. H., Hložek, R., & Marsh, D. J. E. (2022). Ultralight axions and the kinetic Sunyaev-Zel'dovich effect. Phys. Rev. D, 105(6), 063513.

[86] Feldman, D., Liu, Z., & Nath, P. (2007). The Stueckelberg Z-prime Extension with Kinetic Mixing and Milli-Charged Dark Matter From the Hidden Sector. Phys. Rev., D75, 115001.

[87] Fields, B. D., Olive, K. A., Yeh, T.-H., & Young, C. (2020). Big-Bang Nucleosynthesis after Planck. JCAP, 03, 010. [Erratum: JCAP 11, E02 (2020)].

[88] Fogli, G. L., Lisi, E., Marrone, A., Montanino, D., Palazzo, A., & Rotunno, A. M. (2012). Global analysis of neutrino masses, mixings and phases: entering the era of leptonic CP violation searches. Phys. Rev., D86, 013012.

[89] Font-Ribera, A., McDonald, P., Mostek, N., Reid, B. A., Seo, H.-J., & Slosar, A. (2014). DESI and other dark energy experiments in the era of neutrino mass measurements. JCAP, 05, 023.

[90] Franciolini, G., Cotesta, R., Loutrel, N., Berti, E., Pani, P., & Riotto, A. (2022). How to assess the primordial origin of single gravitational-wave events with mass, spin, eccentricity, and deformability measurements. Phys. Rev. D, 105(6), 063510.

[91] Fry, J. N. & Gaztanaga, E. (1993). Biasing and hierarchical statistics in large scale structure. Astrophys. J., 413, 447–452.

[92] Fukuda, S. et al. (2000). Tau neutrinos favored over sterile neutrinos in atmospheric muon-neutrino oscillations. Phys. Rev. Lett., 85, 3999–4003.

[93] Fukuda, Y. et al. (1999). Constraints on neutrino oscillation parameters from the measurement of day night solar neutrino fluxes at Super-Kamiokande. Phys. Rev. Lett., 82, 1810–1814.

[94] Gair, J. R., Antonelli, A., & Barbieri, R. (2022). A Fisher matrix for gravitational-wave population inference. Mon. Not. Roy. Astron. Soc., 519(2), 2736–2753.

[95] Galli, S., Benabed, K., Bouchet, F., Cardoso, J.-F., Elsner, F., Hivon, E., Mangilli, A., Prunet, S., & Wandelt, B. (2014). CMB Polarization can constrain cosmology better than CMB temperature. Phys. Rev., D90(6), 063504.

[96] Giusarma, E., de Putter, R., Ho, S., & Mena, O. (2013). Constraints on neutrino masses from Planck and Galaxy Clustering data. Phys. Rev. D, 88(6), 063515.

[97] Goldberg, H. & Hall, L. J. (1986). A New Candidate for Dark Matter. Phys. Lett., B174, 151.

[98] Gonzalez-Garcia, M. C., Maltoni, M., Salvado, J., & Schwetz, T. (2012). Global fit to three neutrino mixing: critical look at present precision. JHEP, 12, 123.

[99] Gorghetto, M., Hardy, E., & Nicolaescu, H. (2021). Observing invisible axions with gravitational waves. JCAP, 06, 034.

[100] Green, D. et al. (2019). Messengers from the Early Universe: Cosmic Neutrinos and Other Light Relics. Bull. Am. Astron. Soc., 51(7), 159.

[101] Green, M. B., Schwarz, J. H., & Witten, E. (1988). SUPERSTRING THEORY. VOL. 1: INTRODUCTION. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 2012.

[102] Hahn, C., Beutler, F., Sinha, M., Berlind, A., Ho, S., & Hogg, D. W. (2019). Likelihood Non-Gaussianity in Large-Scale Structure Analyses. Mon. Not. Roy. Astron. Soc., 485(2), 2956–2969.

[103] Hall, A. C. & Challinor, A. (2012). Probing the neutrino mass hierarchy with CMB weak lensing. Mon. Not. Roy. Astron. Soc., 425, 1170–1184.

[104] Hannestad, S. & Schwetz, T. (2016). Cosmology and the neutrino mass ordering. JCAP, 1611(11), 035.

[105] Hertzberg, M. P., Tegmark, M., & Wilczek, F. (2008). Axion Cosmology and the Energy Scale of Inflation. Phys. Rev. D, 78, 083507.

[106] Hirata, C. M. & Seljak, U. (2003). Reconstruction of lensing from the cosmic microwave background polarization. Phys. Rev., D68, 083002.

[107] Hlozek, R., Grin, D., Marsh, D. J. E., & Ferreira, P. G. (2015). A search for ultralight axions using precision cosmological data. Phys. Rev. D, 91(10), 103512.

[108] Hlozek, R., Marsh, D. J. E., & Grin, D. (2018). Using the Full Power of the Cosmic Microwave Background to Probe Axion Dark Matter. Mon. Not. Roy. Astron. Soc., 476(3), 3063–3085.

[109] Hook, A., McGehee, R., & Murayama, H. (2018). Cosmologically Viable Low-energy Supersymmetry Breaking. Phys. Rev. D, 98(11), 115036.

[110] Hook, A. & Murayama, H. (2015). Low-energy Supersymmetry Breaking Without the Gravitino Problem. Phys. Rev. D, 92(1), 015004.

[111] Hotinli, S. C., Marsh, D. J. E., & Kamionkowski, M. (2021). Probing ultra-light axions with the 21-cm Signal during Cosmic Dawn. Phys. Rev. D, 106(4), 043529.

[112] Hou, Z., Keisler, R., Knox, L., Millea, M., & Reichardt, C. (2013). How Massless Neutrinos Affect the Cosmic Microwave Background Damping Tail. Phys. Rev. D, 87, 083008.

[113] Hu, W. (1998). Structure formation with generalized dark matter. Astrophys. J., 506, 485–494.

[114] Hu, W., Barkana, R., & Gruzinov, A. (2000). Cold and fuzzy dark matter. Phys. Rev. Lett., 85, 1158–1161.

[115] Hui, L. & Parfrey, K. P. (2008). The Evolution of Bias: Generalized. Phys. Rev. D, 77, 043527.

[116] Hwang, J.-c. & Noh, H. (2009). Axion as a Cold Dark Matter candidate. Phys. Lett. B, 680, 1–3.

[117] Ivezić, v. et al. (2019). LSST: from Science Drivers to Reference Design and Anticipated Data Products. Astrophys. J., 873(2), 111.

[118] Jenkins, A., Frenk, C. S., White, S. D. M., Colberg, J. M., Cole, S., Evrard, A. E., Couchman, H. M. P., & Yoshida, N. (2001). The Mass function of dark matter halos. Mon. Not. Roy. Astron. Soc., 321, 372.

[119] Jimenez, R., Kitching, T., Pena-Garay, C., & Verde, L. (2010). Can we measure the neutrino mass hierarchy in the sky? JCAP, 05, 035.

[120] Kaiser, N. (1987). Clustering in real space and in redshift space. Mon. Not. Roy. Astron. Soc., 227, 1–27.

[121] Kajita, T. (2016). Nobel Lecture: Discovery of atmospheric neutrino oscillations. Rev. Mod. Phys., 88(3), 030501.

[122] Kamionkowski, M., Kosowsky, A., & Stebbins, A. (1997). Statistics of cosmic microwave background polarization. Phys. Rev. D, 55, 7368–7388.

[123] Kaplinghat, M., Knox, L., & Song, Y.-S. (2003). Determining neutrino mass from the CMB alone. Phys. Rev. Lett., 91, 241301.

[124] Kulkarni, M. & Ostriker, J. P. (2021). What is the halo mass function in a fuzzy dark matter cosmology? Mon. Not. Roy. Astron. Soc., 510(1), 1425–1430.

[125] Laguë, A., Bond, J. R., Hložek, R., Rogers, K. K., Marsh, D. J. E., & Grin, D. (2022). Constraining ultralight axions with galaxy surveys. JCAP, 01(01), 049.

[126] Larsen, P., Challinor, A., Sherwin, B. D., & Mak, D. (2016). Demonstration of cosmic microwave background delensing using the cosmic infrared background. Phys. Rev. Lett., 117(15), 151102.

[127] Lattanzi, M. & Gerbino, M. (2018). Status of neutrino properties and future prospects - Cosmological and astrophysical constraints. Front.in Phys., 5, 70.

[128] Lemos, P., Challinor, A., & Efstathiou, G. (2017). The effect of Limber and flat-sky approximations on galaxy weak lensing. JCAP, 1705(05), 014.

[129] Lesgourgues, J. & Pastor, S. (2006). Massive neutrinos and cosmology. Phys. Rept., 429, 307–379.

[130] Lesgourgues, J., Pastor, S., & Perotto, L. (2004). Probing neutrino masses with future galaxy redshift surveys. Phys. Rev. D, 70, 045016.

[131] Linder, E. V. (2005). Cosmic growth history and expansion history. Phys. Rev., D72, 043529.

[132] Loeb, A. & Zaldarriaga, M. (2004). Measuring the small - scale power spectrum of cosmic density fluctuations through 21 cm tomography prior to the epoch of structure formation. Phys. Rev. Lett., 92, 211301.

[133] LoVerde, M. (2014a). Halo bias in mixed dark matter cosmologies. Phys. Rev., D90(8), 083530.

[134] LoVerde, M. (2014b). Spherical collapse in $\nu\Lambda$CDM. Phys. Rev., D90(8), 083518.

[135] LoVerde, M. (2016). Neutrino mass without cosmic variance. Phys. Rev. D, 93(10), 103526.

[136] LoVerde, M. & Zaldarriaga, M. (2014). Neutrino clustering around spherical dark matter halos. Phys. Rev. D, 89(6), 063502.

[137] Maltoni, M., Schwetz, T., Tortola, M. A., & Valle, J. W. F. (2004). Status of global fits to neutrino oscillations. New J. Phys., 6, 122.

[138] Maltoni, M., Schwetz, T., & Valle, J. W. F. (2003). Combining first KamLAND results with solar neutrino data. Phys. Rev., D67, 093003.

[139] Mangano, G., Miele, G., Pastor, S., Pinto, T., Pisanti, O., & Serpico, P. D. (2005). Relic neutrino decoupling including flavor oscillations. Nucl. Phys., B729, 221–234.

[140] Marsh, D. J. E. (2016). Axion Cosmology. Phys. Rept., 643, 1–79.

[141] Marsh, D. J. E. & Ferreira, P. G. (2010). Ultra-Light Scalar Fields and the Growth of Structure in the Universe. Phys. Rev. D, 82, 103528.

[142] Marsh, D. J. E. & Hoof, S. (2023). Astrophysical Searches and Constraints.

[143] Marsh, D. J. E. & Silk, J. (2014). A Model For Halo Formation With Axion Mixed Dark Matter. Mon. Not. Roy. Astron. Soc., 437(3), 2652–2663.

[144] Martin, S. P. (1998). A Supersymmetry primer. Adv. Ser. Direct. High Energy Phys., 18, 1–98.

[145] May, S. & Springel, V. (2021). Structure formation in large-volume cosmological simulations of fuzzy dark matter: impact of the non-linear dynamics. Mon. Not. Roy. Astron. Soc., 506(2), 2603–2618.

[146] McDonald, A. B. (2016). Nobel Lecture: The Sudbury Neutrino Observatory: Observation of flavor change for solar neutrinos. Rev. Mod. Phys., 88(3), 030502.

[147] McIsaac, C., Hoy, C., & Harry, I. (2023). A search technique to observe precessing compact binary mergers in the advanced detector era.

[148] Mention, G., Fechner, M., Lasserre, T., Mueller, T. A., Lhuillier, D., Cribier, M., & Letourneau, A. (2011). The Reactor Antineutrino Anomaly. Phys. Rev., D83, 073006.

[149] Modi, C., Castorina, E., & Seljak, U. (2017). Halo bias in Lagrangian Space: Estimators and theoretical predictions. Mon. Not. Roy. Astron. Soc., 472(4), 3959–3970.

[150] Moradinezhad Dizgah, A., Lee, H., Muñoz, J. B., & Dvorkin, C. (2018). Galaxy Bispectrum from Massive Spinning Particles. JCAP, 05, 013.

[151] Moroi, T., Murayama, H., & Yamaguchi, M. (1993). Cosmological constraints on the light stable gravitino. Phys. Lett. B, 303, 289–294.

[152] Muñoz, J. B. & Dvorkin, C. (2018). Efficient Computation of Galaxy Bias with Neutrinos and Other Relics. Phys. Rev. D, 98(4), 043503.

[153] Muñoz, J. B., Grin, D., Dai, L., Kamionkowski, M., & Kovetz, E. D. (2016). Search for Compensated Isocurvature Perturbations with Planck Power Spectra. Phys. Rev., D93(4), 043008.

[154] Muñoz, J. B., Kovetz, E. D., Raccanelli, A., Kamionkowski, M., & Silk, J. (2017). Towards a measurement of the spectral runnings. JCAP, 05, 032.

[155] National Academies of Sciences, Engineering, and Medicine (2021). Pathways to Discovery in Astronomy and Astrophysics for the 2020s. Washington, DC: The National Academies Press.

[156] Nishizawa, A., Berti, E., Klein, A., & Sesana, A. (2016). eLISA eccentricity measurements as tracers of binary black hole formation. Phys. Rev. D, 94(6), 064020.

[157] Nishizawa, A., Sesana, A., Berti, E., & Klein, A. (2017). Constraining stellar binary black hole formation scenarios with eLISA eccentricity measurements. Mon. Not. Roy. Astron. Soc., 465(4), 4375–4380.

[158] Noh, H., Hwang, J.-c., & Park, C.-G. (2017). Axion as a cold dark matter candidate: Proof to fully nonlinear order. Astrophys. J., 846(1), 1.

[159] Nunokawa, H., Parke, S. J., & Valle, J. W. F. (2008). CP Violation and Neutrino Oscillations. Prog. Part. Nucl. Phys., 60, 338–402.

[160] Okamoto, T. & Hu, W. (2003). CMB lensing reconstruction on the full sky. Phys. Rev., D67, 083002.

[161] Osato, K., Sekiguchi, T., Shirasaki, M., Kamada, A., & Yoshida, N. (2016). Cosmological Constraint on the Light Gravitino Mass from CMB Lensing and Cosmic Shear. JCAP, 1606, 004.

[162] Osipowicz, A. et al. (2001). KATRIN: A Next generation tritium beta decay experiment with sub-eV sensitivity for the electron neutrino mass. Letter of intent.

[163] Pace, F., Sakr, Z., & Tutusaus, I. (2020). Spherical collapse in Generalized Dark Matter models. Phys. Rev. D, 102(4), 043512.

[164] Palanque-Delabrouille, N. et al. (2015). Neutrino masses and cosmology with Lyman-alpha forest power spectrum. JCAP, 11, 011.

[165] Parfrey, K., Hui, L., & Sheth, R. K. (2011). Scale-dependent halo bias from scale-dependent growth. Phys. Rev. D, 83, 063511.

[166] Park, C.-G., Hwang, J.-c., & Noh, H. (2012). Axion as a cold dark matter candidate: low-mass case. Phys. Rev. D, 86, 083535.

[167] Peccei, R. D. & Quinn, H. R. (1977). CP Conservation in the Presence of Instantons. Phys. Rev. Lett., 38, 1440–1443.

[168] Peebles, P. J. E. (1966). Primordial Helium Abundance and the Primordial Fireball. 2. Astrophys. J., 146, 542–552.

[169] Peters, P. C. (1964). Gravitational Radiation and the Motion of Two Point Masses. Phys. Rev., 136, B1224–B1232.

[170] Poulin, V., Smith, T. L., Grin, D., Karwal, T., & Kamionkowski, M. (2018). Cosmological implications of ultralight axionlike fields. Phys. Rev. D, 98(8), 083525.

[171] Poulin, V., Smith, T. L., Karwal, T., & Kamionkowski, M. (2019). Early Dark Energy Can Resolve The Hubble Tension. Phys. Rev. Lett., 122(22), 221301.

[172] Preskill, J., Wise, M. B., & Wilczek, F. (1983). Cosmology of the Invisible Axion. Phys. Lett. B, 120, 127–132.

[173] Press, W. H. & Schechter, P. (1974). Formation of galaxies and clusters of galaxies by selfsimilar gravitational condensation. Astrophys. J., 187, 425–438.

[174] Pritchard, J. R. & Pierpaoli, E. (2008). Constraining massive neutrinos using cosmological 21 cm observations. Phys. Rev. D, 78, 065009.

[175] Raccanelli, A., Verde, L., & Villaescusa-Navarro, F. (2019). Biases from neutrino bias: to worry or not to worry? Mon. Not. Roy. Astron. Soc., 483(1), 734–743.

[176] Randall, L., Shelest, A., & Xianyu, Z.-Z. (2022). An Efficient Signal-to-noise Approximation for Eccentric Inspiraling Binaries. Astrophys. J., 924(2), 102.

[177] Randall, L. & Xianyu, Z.-Z. (2018a). An Analytical Portrait of Binary Mergers in Hierarchical Triple Systems. Astrophys. J., 864(2), 134.

[178] Randall, L. & Xianyu, Z.-Z. (2018b). Induced Ellipticity for Inspiraling Binary Systems. Astrophys. J., 853(1), 93.

[179] Randall, L. & Xianyu, Z.-Z. (2021). Eccentricity without Measuring Eccentricity: Discriminating among Stellar Mass Black Hole Binary Formation Channels. Astrophys. J., 914(2), 75.

[180] Reece, M. (2019). Photon Masses in the Landscape and the Swampland. JHEP, 07, 181.

[181] Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., Kremer, K., Rasio, F. A., Samsing, J., Ye, C. S., & Zevin, M. (2018). Post-Newtonian Dynamics in Dense Star Clusters: Formation, Masses, and Merger Rates of Highly-Eccentric Black Hole Binaries. Phys. Rev. D, 98(12), 123005.

[182] Rogers, K. K. & Peiris, H. V. (2021a). General framework for cosmological dark matter bounds using N-body simulations. Phys. Rev. D, 103(4), 043526.

[183] Rogers, K. K. & Peiris, H. V. (2021b). Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest. Phys. Rev. Lett., 126(7), 071302.

[184] Rossi, G., Yèche, C., Palanque-Delabrouille, N., & Lesgourgues, J. (2015). Constraints on dark radiation from cosmological probes. Phys. Rev. D, 92(6), 063505.

[185] Sarkar, A. K., Pandey, K. L., & Sethi, S. K. (2021). Using the redshift evolution of the Lyman-α effective opacity as a probe of dark matter models. JCAP, 10, 077.

[186] Schwabe, B. & Niemeyer, J. C. (2022). Deep Zoom-In Simulation of a Fuzzy Dark Matter Galactic Halo. Phys. Rev. Lett., 128(18), 181301.

[187] Senatore, L. (2015). Bias in the Effective Field Theory of Large Scale Structures. JCAP, 11(11), 007.

[188] Senatore, L. & Zaldarriaga, M. (2017). The Effective Field Theory of Large-Scale Structure in the presence of Massive Neutrinos.

[189] Sheth, R. K. & Tormen, G. (1999). Large scale bias and the peak background split. Mon. Not. Roy. Astron. Soc., 308, 119.

[190] Smith, T. L., Poulin, V., Bernal, J. L., Boddy, K. K., Kamionkowski, M., & Murgia, R. (2021). Early dark energy is not excluded by current large-scale structure data. Phys. Rev. D, 103(12), 123542.

[191] Sprenger, T., Archidiacono, M., Brinckmann, T., Clesse, S., & Lesgourgues, J. (2019). Cosmology in the era of Euclid and the Square Kilometre Array. JCAP, 1902, 047.

[192] Svrcek, P. & Witten, E. (2006). Axions In String Theory. JHEP, 06, 051.

[193] 't Hooft, G. (1976). Symmetry Breaking Through Bell-Jackiw Anomalies. Phys. Rev. Lett., 37, 8–11.

[194] Tegmark, M. (1997). Measuring cosmological parameters with galaxy surveys. Phys. Rev. Lett., 79, 3806–3809.

[195] Tinker, J. L., Kravtsov, A. V., Klypin, A., Abazajian, K., Warren, M. S., Yepes, G., Gottlober, S., & Holz, D. E. (2008). Toward a halo mass function for precision cosmology: The Limits of universality. Astrophys. J., 688, 709–728.

[196] Turner, M. S. (1987). Thermal Production of Not SO Invisible Axions in the Early Universe. Phys. Rev. Lett., 59, 2489. [Erratum: Phys.Rev.Lett. 60, 1101 (1988)].

[197] Vagnozzi, S., Brinckmann, T., Archidiacono, M., Freese, K., Gerbino, M., Lesgourgues, J., & Sprenger, T. (2018). Bias due to neutrinos must not uncorrect'd go. JCAP, 1809(09), 001.

[198] Vagnozzi, S., Giusarma, E., Mena, O., Freese, K., Gerbino, M., Ho, S., & Lattanzi, M. (2017). Unveiling ν secrets with cosmological data: neutrino masses and mass hierarchy. Phys. Rev., D96(12), 123503.

[199] Viel, M., Lesgourgues, J., Haehnelt, M. G., Matarrese, S., & Riotto, A. (2005). Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman-alpha forest. Phys. Rev. D, 71, 063534.

[200] Villaescusa-Navarro, F., Marulli, F., Viel, M., Branchini, E., Castorina, E., Sefusatti, E., & Saito, S. (2014). Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies. JCAP, 1403, 011.

[201] Weinberg, D. H., Bullock, J. S., Governato, F., Kuzio de Naray, R., & Peter, A. H. G. (2015). Cold dark matter: controversies on small scales. Proc. Nat. Acad. Sci., 112, 12249–12255.

[202] Weinberg, S. (1978). A New Light Boson? Phys. Rev. Lett., 40, 223–226.

[203] Wilczek, F. (1978). Problem of Strong *P* and *T* Invariance in the Presence of Instantons. Phys. Rev. Lett., 40, 279–282.

[204] Witten, E. (1984). Some Properties of O(32) Superstrings. Phys. Lett. B, 149, 351–356.

[205] Xu, W. L., DePorzio, N., Muñoz, J. B., & Dvorkin, C. (2021). Accurately Weighing Neutrinos with Cosmological Surveys. Phys. Rev. D, 103(2), 023503.

[206] Zaldarriaga, M. & Seljak, U. (1997). An all sky analysis of polarization in the microwave background. Phys. Rev. D, 55, 1830–1840.

[207] Zarrouk, P. et al. (2018). The clustering of the SDSS-IV extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample: measurement of the growth rate of structure from the anisotropic correlation function between redshift 0.8 and 2.2. Mon. Not. Roy. Astron. Soc., 477(2), 1639–1663.

[208] Zehavi, I. et al. (2002). Galaxy Clustering in early SDSS Redshift Data. Astrophys. J., 571, 172–190.